Fair and Explainable Collective Decision

Monday, March 30, 2026 - Friday, April 10, 2026

Institut Pascal

Scientific Program

Tutorials
Dominik Peters (CNRS, LAMSADE, Université Paris Dauphine)
Participatory budgeting
Bettina Klaus (University of Lausanne)

Matching under preferences

This tutorial introduces the theory and applications of matching and allocation problems under preferences, a central topic in market design and microeconomic theory. The tutorial is divided into two parts.
The first part presents the foundations of matching theory, focusing on two-sided matching markets and school choice. We introduce the basic matching model, stability as a central solution concept, and the Gale–Shapley deferred acceptance algorithm. We discuss optimal stable matchings, incentives and manipulation, and the trade-offs between stability, fairness, and efficiency. The tutorial then turns to school choice mechanisms, comparing the Boston mechanism, deferred acceptance, and the top trading cycles mechanism, and discusses real-world market design applications such as school assignment systems and labor market matching.
The second part focuses on allocation problems with indivisible goods and the top trading cycles mechanism. We review the Shapley–Scarf housing market model and key results on core allocations, strategy-proofness, and efficiency of trading cycles mechanisms. We then discuss deterministic and probabilistic allocation rules, including serial dictatorship and random serial dictatorship, and present recent research on efficiency notions such as pairwise efficiency, Pareto efficiency, and monotonicity in deterministic and probabilistic allocation problems.
Overall, the tutorial provides an overview of both classical and recent results in matching theory and object allocation, with an emphasis on mechanism design, incentives, and efficiency, and highlights connections between theory and practical market design applications.
Invited talks
Péter Biró (KRTK, Corvinus University of Budapest)

Optimality and fairness in kidney exchange programmes

In this talk we survey some recent developments on the OR and Game Theoretical aspects of kidney exchange programmes (KEPs). These programmes have been established in most of the Western countries in the last two decades to facilitate the exchange of living kidney donors for those recipients who have willing but immunologically incompatible donors. In the first part, we describe the European practices including the hierarchic optimisation used in the matching runs of the KEPs for computing optimal solutions. These IP-based methods have been implemented in the simulator software developed in the ENCKEP COST Action (2016-2021) and subsequently incorporated in the KEPSOFT software that provides a common IT-platform for European applications. In the second part, we describe some alternative solution concepts based on the (strong, Wako, and weak) core of the underlying cooperative game. We explain how the usage of these classical solution concepts can incentivise the recipients to register valuable donors or multiple donors in KEPs by the so-called respecting improvement property. Finally, in the third part, we present new results on international KEPs, where credit-based compensation schemes have been proposed to balance out the benefits of the countries when merging their national pools, where the reference target values are coming from classical solution concepts of TU matching games.
Jiehua Chen (TU Wien)

Fair Congested Assignments: From Aversion to Appreciation

The congested assignment problem concerns assigning agents to resources where agents have preferences over both the resources and their congestion levels. We examine two complementary preference models: congestion-averse (agents prefer lower congestion) and congestion-appreciating (agents prefer higher congestion). Despite their apparent symmetry, these models exhibit fundamentally different structural and computational behavior.
I will explore several fairness concepts--non-wastefulness, envy-freeness, Nash stability, top-guarantee, and competitiveness--where top-guarantee and competitiveness were recently introduced by Bogomolnaia and Moulin. The implication structure among these concepts differs markedly between the two settings, with some implications even reversing direction.
For the averse model, I will present a polynomial-time network-flow approach for determining the existence competitive assignments, along with NP-hardness and parameterized complexity results for two optimization variants: finding an envy-free and top-guarantee assignment, and finding a maximally competitive assignment. For the appreciating model, I will present a polynomial-time algorithm for determining the existence of top-guaranteed assignments and discuss an NP-hard variant: finding a top-k-guarantee assignment that minimizes k, i.e., minimizing the largest rank of any agent. The talk will conclude with open problems, most notably the conjecture that competitive assignments always exist under congestion-appreciation.
Ulle Endriss (University of Amsterdam)

Axiomatic Engineering

In economic theory we often use axioms to characterise mechanisms of interest or to argue why a normatively appealing mechanism might not be available. In this talk, I will present an 'algorithmic' perspective on axioms, where we think of axioms as formal objects we can reason about with the help of computers. I will illustrate this way of thinking with examples such as (i) automating parts of the process of proving theorems involving axioms and (ii) automatically constructing explanations for the decisions suggested by an economic mechanism.
Sébastien Konieczny (CNRS, CRIL, Université d'Artois)

Taking correct decisions together: Belief Merging and Truth Tracking

We will discuss belief merging operators, which aim to resolve conflicts between information coming from different sources. This information can consist of either beliefs or goals. Thus, belief merging provides a way to derive a collective decision from individual opinions expressed in propositional logic. We will begin with a brief introduction to the topic by recalling the properties expected of these operators and the main families of merging operators. We will then show that these operators can be used to track the truth in different settings. Finally, we will show how to improve truth-tracking performance in belief merging and voting methods by analyzing individual opinions and estimating the reliability of the sources.
Jan Maly (WU Wien)

Fairness, games and power in participatory budgeting

Many of the most prominent fairness axioms in the participatory budgeting literature are, sometimes implicitly, based on the idea that groups of voters cannot "leave" with their share of the budget to achieve a better outcome. In this talk, I will explain how we can formally model many of these axioms and some prominent voting rules as equilibria in a relatively simple game - at least in the unit-cost case. I will then sketch which of these results carry over to the more general participatory budgeting setting and, more interestingly, which don't. Subsequently, I will focus on a specific type of equilibrium, the core, and explain why it might be empty or, if it is not empty, lead to undesirable

outcomes due to power imbalances between the voters, highlighting an important potential drawback of equilibria-based fairness notions.
Nicolas Maudet (LIP6, Sorbonne Université)

Fair division of indivisible goods: whose fairness are we talking about?

The problem of fair division of indivisible goods has recently attracted a lot of interest in computational social choice, with impressive algorithmic progresses, and a large variety of original notions introduced. However, it is still unclear which fairness notion is the most appropriate -- and who should decide that in the first place, on which basis. I will reflect on those questions, discussing some recent results in the field, and pointing connections to participatory budgeting and matching.
Piotr Skowron (University of Warsaw)

Stable-priceable outcomes in participatory budgeting

A number of rules for participatory budgeting and committee elections are inspired by market-based ideas: voters are initially endowed with virtual budgets, which they can use to “purchase” candidates they support. This principle underlies, in particular, the Method of Equal Shares and Phragmén’s rule. We explore this perspective further, showing how market-based reasoning can be used to derive stronger notions of proportionality and to guide the design of new voting rules.
Short talks
Martin Bullinger (University of Bristol)

The power of matching for welfare maximization in hedonic games

In this talk, I will review results concerned with welfare approximation in hedonic games. A recurrent theme is that, despite hedonic games being a model of coalition formation where large groups may form, matching algorithms yield (asymptotically) optimal approximation algorithms.
We will experience this paradigm in the prominent game classes of additively separable and fractional hedonic games and for the following algorithmic settings:
1.) The standard offline model.
2.) The standard online model.
3.) The online model where algorithms may dissolve coalitions.
Piotr Faliszewski (AGH University of Kraków)

The Games That Project Proposers Play

In this talk I will present results from three recent papers about games that arise in participatory budgeting. Each of these games regards project proposers and include the following strategic issues:
- choosing costs of the projects,
- choosing which of several possible projects to submit,
- choosing whether to merge some projects.
Issues captured by these games indeed appear in real-life and proposers have to make strategic choices. We will discuss the existence of stable solutions to the games, the complexity of finding them, and results from analysis of PB data from Pabulib with respect to them.
Matthieu Hervouin

Online Algorithms for Participatory Budgeting

Participatory Budgeting (PB), and in particular, proportionality in PB, have received significant attention from the Computational Social Choice community in recent years. This led to the discovery of voting rules like the Method of Equal Shares, which selects outcomes in a fair and representative fashion. However, these rules assume that the set of possible projects is given in advance. In this paper, we propose a new framework, which we call online PB, in which projects are revealed one by one, and a binding funding decision has to be made on the spot.
We study several classical fairness axioms from the offline PB literature in this online setting, namely priceability, local fair share, and the most prominent axioms of justified representation, JR, PJR, and EJR. We show that the first two properties are always satisfiable but incompatible in the online setting and find tight approximations for the justified representation axioms.
Additionally, we show in experiments that, in practice, a simple greedy online PB rule produces outcomes that are nearly as fair as the outcomes produced by the Method of Equal Shares (without completion) in offline PB. We also show that if a completion were possible (which is not the case in online PB), our simple greedy rule could come close to the performance of MES with completion.
Šimon Schierreich (AGH University of Krakow)

Participatory Budgeting Project Strength via Candidate Control

We study the complexity of candidate control in participatory budgeting elections. The goal of constructive candidate control is to ensure that a given candidate wins by either adding or deleting candidates from the election (in the destructive setting, the goal is to prevent a given candidate from winning). We show that such control problems are NP-hard to solve for many participatory budgeting voting rules, including Phragmén and Equal-Shares, but there are natural cases with polynomial-time algorithms. We also argue that control by deleting candidates is a useful tool for assessing the performance (or, strength) of initially losing projects, and we support this view with experiments on real-life PB instances.

Joint work with Piotr Faliszewski, Łukasz Janeczko, Dušan Knop, Jan Pokorný, Mateusz Słuszniak, and Krzysztof Sornat.
Ildiko Schlotter (ELTE Centre for Economic and Regional Studies)

Stable hypergraph matching in unimodular hypergraphs

We study the NP-hard Stable Hypergraph Matching (SHM) problem and its generalization allowing capacities, the Stable Hypergraph b-Matching (SHbM) problem, and investigate their computational properties under various structural constraints. Our study is motivated by the fact that Scarf's Lemma together with a result of Lovász guarantees the existence of a stable matching whenever the underlying hypergraph is normal. Furthermore, if the hypergraph is unimodular (i.e., its incidence matrix is totally unimodular), then even a stable b-matching is guaranteed to exist. However, no polynomial-time algorithm is known for finding a stable matching or b-matching in unimodular hypergraphs.
We identify subclasses of unimodular hypergraphs where SHM and SHbM are tractable such as laminar hypergraphs or so-called subpath hypergraphs with bounded-size hyperedges; for the latter case, even a maximum-weight stable b-matching can be found efficiently. We complement our algorithms by showing that optimizing over stable matchings is NP-hard even in laminar hypergraphs. As a practically important special case of SHbM for unimodular hypergraphs, we investigate a tripartite stable matching problem with students, schools, and companies as agents, called the University Dual Admission problem, which models real-world scenarios in higher

education admissions.
Finally, we examine a superclass of subpath hypergraphs that are normal but not necessarily unimodular, namely subtree hypergraphs where hyperedges correspond to subtrees of a tree. We establish that for such hypergraphs, stable matchings can be found in polynomial time but, in the setting with capacities, finding a stable b-matching is NP-hard.
Krzysztof Sornat (AGH University of Kraków)

Algorithms for Structured Elections under Thiele Voting Rules

We study the computational complexity of winner determination problems in approval-based committee elections under Thiele voting rules. These form a class of rules parameterized by a fixed weight vector that specifies how a voter's satisfaction depends on the number of approved candidates elected. We first analyze the structure of optimal solutions based on the sets of voters who approve each candidate---that is, how voters' approval ballots induce dependencies between candidates---revealing constraints on a winning committee under any fixed Thiele voting rule. Using this, we design FPT algorithms for Proportional Approval Voting (PAV) and other Thiele rules on a natural restricted domain known as the Voter Interval domain---that is, after a suitable ordering of voters, each candidate is approved by a consecutive interval of voters. In particular, we show that every Thiele rule on Voter Interval is FPT with respect to a parameter for which the problem is NP-hard on general instances, even when the parameter takes constant values. Our results advance the understanding of the computational complexity of PAV on Voter Interval instances, which remains one of the central open questions in this area. We further resolve two open questions from the literature on PAV (and other Thiele voting rules) by providing a polynomial-time algorithm for instances where each candidate is approved by at most two voters, and an FPT algorithm parameterized by the total score of a winning committee.

It is a joint work with Alexandra Lassota. Appeared at AAAI-26.
Tomasz Wąs (University of Oxford)

Agreement, Diversity, and Polarization Indices for Approval Elections

An index is a function that, given an election, outputs a value between 0 and 1, indicating the extent to which this election has a particular feature. We seek indices that capture agreement, diversity, and polarization among voters in approval elections, and that are normalized with respect to saturation. By the latter we mean that if two elections differ by the fraction of candidates approved by an average voter, but otherwise are of similar nature, then they should have similar index values. We propose several indices, analyze their properties, and use them to (a) derive a new map of approval elections, and (b) show similarities and differences between various real-life elections from Pabulib, Preflib and other sources.
Makoto Yokoo (Kyushu University)

Fairness and efficiency trade-off in two-sided matching under hereditary constraints

The theory of two-sided matching has been extensively developed and applied to many real-life application domains. As the theory has been applied to increasingly diverse types of environments, researchers and practitioners have encountered various forms of distributional constraints. As a mechanism can handle a more general class of constraints, we can assign students more flexibly to colleges to increase students' welfare. However, it turns out that there exists a trade-off between students' welfare (efficiency) and fairness (which means no student has justified envy). In this work, we establish new weaker fairness / efficiency requirements and clarify the boundary on whether

these requirements are compatible under general hereditary constraints.

Posters

Matthew Casey (Northwestern University): Streamlining Equal Shares

Ritu Dutta (Indian Institute of Technology Bombay): Is the Condorcet winner always a fair and compelling choice? A dilemma!

Luca Kreisel (Hasso Plattner Institute, University of Potsdam): Explanation Systems for Approval-Based Multiwinner Voting

Georgios Kalantzis (University of Edinburgh): Approximate-EFX Allocations with Ordinal and Limited Cardinal Information

Huynh Le Nhat Linh (GATE Lyon - Saint Etienne): Consistency in Collective Decision-Making under Uuncertainty: An Axiomatic Approach

Karl Jochen Micheel (CentraleSupélec, Université Paris-Saclay): Proportionality Variations in Repeated Fair Division

Zein Pishbin (University of Paris Dauphine-PSL): Fairness in the Multi-Secretary Problem

Andrew Poile (University of Southampton): School choice with transportation

Soumyajit Pyne (TIFR, Mumbai): Nearly Tight Bounds on Approximate Equilibria in Spatial Competition on the Line

Francesco Sabatino (CentraleSupélec, Université Paris-Saclay): ExFaMa: A Platform To Explain Fair Matchings

Drew Springham (King’s College London): Multiwinner Voting with d-Dimensional Euclidean Preferences under Incomplete Information

Markus Utke (TU Eindhoven): City Sampling for Citizens’ Assemblies

Juan Ignacio Zambrano (Université Toulouse Capitole): Privacy-Aware Predictions in Participatory Budgeting

Cathy Zeng (University of Rochester): Multi-Period Student Assignment