COST Action CA 16228

EUROPEAN NETWORK FOR GAME THEORY (GAMENET)

PROGRAMME and ABSTRACTS

GAMENET Conference
Management Committee Meeting
Core Group Meeting

Institute of Information Theory and Automation
The Czech Academy of Sciences
Pod Vodárenskou věží 4, 182 00 Prague 8
November 21-22, 2019

Organised by
Institute of Information Theory and Automation
The Czech Academy of Sciences

Local Organising Committee:

T. V. Guy (chair)
J. Homolová
F. Hůla
M. Kárný
P. H. Neuner
M. Ruman (webmaster)

Acknowledgement
COST Action CA 16228
Ministry of Education, Youth and Sports of the Czech Republic, project LTC18075
The Czech Academy of Sciences
DataSentics s.r.o

Dear participant,

It is our great pleasure to welcome you to Prague for the GAMENET Conference and Management Committee Meeting.

Prague is well known for its historical monuments, classical architecture and rich cultural heritage. Natural sciences have a long tradition in Prague and many famous scientists like Bernard Bolzano, Tycho Brahe, Johannes Kepler, Christian Doppler, Albert Einstein, Jaroslav Heyrovsky worked here.

We hope that the intensive scientific programme will stimulate spark discussions that will continue far beyond the given limited time of the conference and encourage further collaborations. We hope that you will enjoy this event and wish you a pleasant stay in the city of Prague. We cordially thank you for your contribution to the success of this gathering.

With kind regards,

Local Organising Committee

PROGRAMME

THURSDAY, NOVEMBER 21		
8:30--9:00		Registration
SESSION I		**Chair: Mathias Staudigl**
9:00—9:30		Welcome and Opening by MC Chair
9:30—9:40		WG1 Chair: Report on the progress
9:40—9:50		WG2 Chair: Report on the progress
9:50—10:00		WG4 Chair: Report on the progress
10:00—10:10		WG3 Chair: Report on the progress
10:10—10:40	**Coffee Break**	
SESSION II		**Chair: Luiz DaSilva**
10:40—11:30		**Invited Talk: Catherine Rainer** *Solving two-state Markov games with incomplete information on one side*
11:30—12:20		**Invited Talk: Penelope Hernandez** *How Bayesian persuasion can help reduce illegal parking and other socially undesirable behavior*
12:20—12:30	**Event Photo**	All participants
12:30—14:00	**Lunch Break**	
SESSION III		**Chair: Panayotis Mertikopoulos**
14:00—14:50		**Invited Talk: Edith Elkind** *Hedonic diversity games*
14:50—15:00		Newcomer Presentation, Miklós Krész (SI)
15:00—16:00		Poster Spotlights
16:00—17:00	**Coffee Break and Poster Session**	
17:00—18:00		Core Group Meeting
19:30—22:00		**Conference Dinner at Restaurant “PROFESNÍ DŮM”** (Malostranské nám. 25, tram stop Malostranská)

FRIDAY, NOVEMBER 22		
SESSION IV		**Chair: Marco Scarsini**
9:00—9:50		**Invited Talk: Ágnes Cseh** *Pareto optimal coalitions of fixed size*
9:50—10:40		**Invited Talk: Michal Feldman** *Auction design under interdependent values*
10:40—10:50		Closing Remarks
10:50—11:20	**Coffee Break**	
11:20—13:00		MC Meeting
14:30—16:30		**Guided Prague Tour** (mandatory booking)

INVITED TALKS

Catherine Rainer, Université de Brest

Title: Solving Two-state Markov Games with Incomplete Information on One Side

joint work with Galit Ashkenazi-Golan and Eilon Solan

Abstract: We study the optimal use of information in Markov games with incomplete information on one side and two states. We provide a finite-stage algorithm for calculating the limit value as the gap between stages goes to 0, and an optimal strategy for the informed player in the limiting game in continuous time. This limiting strategy induces an $\ep$-optimal strategy for the informed player, provided the gap between stages is small.

Penélope Hernández, University of Valencia

Title: How Bayesian Persuasion Can Help Reduce Illegal Parking and Other Socially Undesirable Behavior

joint work with Zvika Neeman

Abstract:
We consider the question of how best to allocate enforcement resources across different locations with the goal of deterring unwanted behaviour. We rely on “Bayesian persuasion” to improve deterrence. Our approach is distinguished by the following five features: (1) we consider a problem in which the principal has to allocate resources and then send messages (persuade) rather than just persuade. (2) Messages are received by drivers in n different neighborhoods, so persuasion is with respect to multiple audiences. (3) The problem is a “constrained convexification” rather than just a convexification problem, where the constraints are due to resource and probability restrictions. This implies that convexification may be partial rather than complete as is usually the case in Bayesian persuasion models. (4) Even though the basic problem is not linear, we show that it can be cast as a linear programming problem. Finally, (5) we characterize the number of messages needed in order to obtain the optimal solution, and describe conditions under which it is possible to explicitly solve the problem with only two messages.

Edith Elkind, University of Oxford

Title: Hedonic Diversity Games

based on joint work with Ayumi Igarashi, Robert Bredereck and Niclas Boehmer

Abstract: We consider a coalition formation setting where each agent belongs to one of the two types, and agents' preferences over coalitions are determined by the fraction of the agents of their own type in each coalition. This setting differs from the well-studied Schelling's model in that some agents may prefer homogeneous coalitions, while others may prefer to be members of a diverse group, or a group that mostly consists of agents of the other type. We model this setting as a hedonic game and investigate the existence of stable outcomes using hedonic games solution concepts, such as Nash stability, individual stability and core stability. In particular, we establish that every hedonic diversity game admits an individually stable solution, and such solutions can be computed in polynomial time.

Ágnes Cseh, Institute of Economics, Hungarian Academy of Sciences

Title: Pareto Optimal Coalitions of Fixed Size

joint work with Tamás Fleiner and Petra Harján

Abstract: We tackle the problem of partitioning players into groups of fixed size, such as allocating eligible students to shared dormitory rooms. Each student submits preferences over the other individual students. We study several settings, which differ in the size of the rooms to be filled, the orderedness or completeness of the preferences, and the way of calculating the value of a coalition---based on the best or worst roommate in the coalition. In all cases, we determine the complexity of deciding the existence, and then finding a Pareto optimal assignment, and the complexity of verifying Pareto optimality for a given assignment.

Michal Feldman, Tel-Aviv University

Title: Auction Design Under Interdependent Values

Abstract: We study combinatorial auctions with interdependent valuations. In such settings, every agent has private signal, and every agent has a valuation function that depends on the private signals of all the agents. Interdependent valuations capture settings where agents lack information to determine their own valuations. Examples include auctions for artwork or oil drilling rights. For single item auctions and assume some restrictive conditions (the so-called single-crossing condition), full welfare can be achieved. However, in general, there are strong impossibility results on welfare maximization in the interdependent setting. This is in contrast to settings where agents are aware of their own valuations, where the optimal welfare can always be obtained by an incentive compatible mechanism. Motivated by these impossibility results, we study welfare maximization for interdependent valuations through the lens of approximation. We introduce two valuation properties that enable positive results. The first is a relaxed, parameterized version of single crossing; the second is a submodularity condition over the signals. We obtain a host of approximation guarantees under these two notions for various scenarios.

LIST OF POSTER PRESENTATIONS AND ABSTRACTS

The posters will be on boards from 15:00 till 18:30, Thursday.
Presenting authors are requested to be ready to present their posters during the poster spotlights session (15:00-16:00) and the poster session (16:00-17:00) on Thursday.

Poster spotlights are meant to be strictly brief (4 minutes) presentations of posters.

No.	Authors	Title
P-02	**Y. Tsodikovich**, E. Lehrer (IL)	*Stochastic Revision Opportunities in Markov Decision Problems*
P-03	D. Gilo, A. Porat, **Y. Tsodikovich** (IL, FR)	*Addiction to a Network*
P-04	**R. Bou Rouphael**, M. Le Treust (FR)	*Impact of the Private Observation in Persuasion Game*
P-05	**J. Syrovátková** (CZ)	*Finding Automaton Maximizing Score in Prisoner's Dilemma Tournament*
P-06	**F. Hůla**, T.V. Guy, M. Kárný (CZ)	*On Inspection of Bayesian Knowledge Sharing*
P-07	**F. Rolenec**, M. Kárný (CZ)	*Learning with Forgetting Preserving a Guaranteed Knowledge*
P-08	**E. Macault**, M. Scarsini, T. Tomala (FR, IT)	*Learning in Dynamic Routing Games with Symmetric Incomplete Information*
P-09	**P. Kocourek** (CZ)	*Revealing Private Information in a Patent Race*
P-10	**M. Ruman**, T.V. Guy, M. Kárný (CZ)	*Preference Elicitation for Markov Decision Processes within Fully Probabilistic Design Framework*
P-11	**D. Gagliardi**, G. Russo (IRL)	*Constrained Fully Probabilistic Control*
P-12	**J. Forcan**, M. Mikalacki (SRB)	*Winning Fast When Both Maker and Breaker Are Walkers*
P-13	H. Echzell, P. Lenzner, **L. Molitor**, T. Friedrich, M. Pappik, F. Schöne, F. Sommer and D. Stangl (DE)	*Convergence and Hardness of Strategic Schelling Segregation*

P-02

Stochastic revision opportunities in Markov decision problems

Yevgeny Tsodikovich[1] · Ehud Lehrer[1]

Published in Annals of Operations Research, Aug. 2019, Vol. 279, Issue 1–2, pp 251–270, https://doi.org/10.1007/s10479-019-03252-9.

Abstract
We extend Markov Decision Processes to situations where the actions are binding and cannot be changed in every period. Instead, the decision maker can revise her actions at random times. We consider two slightly different models. In the first, the revision opportunity appears at a specific stage at which the decision maker can change her action, but is lost if not used. The action taken then remains constant until the next revision opportunity comes up. In the second model, the revision opportunity remains open and can be used at any time after it appears. Only when the action is changed, it becomes binding again for another random period. We compare between different stochastic revision processes and characterize when one is always preferred to another.

Keywords Markov decision process · Stochastic dominance · Commitment · Exogenous timing

JEL Classification C41 · C73 · D81

The authors wish to thank Eilon Solan, Galit Golan-Ashkenazi, David Lagziel, Bar Light, Ilan Nehama and two anonymous refrees of the *Annals of Operations Research* for their highly valuable comments. We acknowledge the support of the Israel Science Foundation, Grants #963/15 and #2510/17.

✉ Yevgeny Tsodikovich
yevgenyt@post.tau.ac.il

1 School of Mathematical Sciences, Tel-Aviv University, 69978 Tel-Aviv, Israel

P-03

Addiction to a Network

David Gilo
Buchmann Faculty of Law, Tel Aviv University

Ariel Porat
President, Tel Aviv University

Yevgeny Tsodikovich
Aix-Marseille School of Economics (AMSE)

yevgets@gmail.com

In this paper we study why a well informed and rational consumer starts consuming a product that harms him. In our model, the utility from consuming the product depends not only on the consumer's own current and past consumption, but also on the aggregate past consumption of other consumers. This network of consumers exerts peer pressure on the individual consumer that reduces the consumer's marginal utility from consuming an alternative product. Consumers in the network differ from one another by the expected harm they suffer from consumption. The model also captures cases where other consumers' consumption directly harms the individual consumer. For example, a non-smoker might be exposed to passive smoking by his peers. Such a consumer may start smoking himself, so as to at least enjoy the benefits of smoking.

The paper characterizes optimal consumption paths and equilibrium behavior and defines conditions for individual consumers to start consuming. We examine when a consumer consumes the harmful product due to the existence of the network, while he is better off pursuing a consumption path on his own, without the network. We distinguish between consumption paths in which the consumer consumes the harmful product only because other consumers have been consuming it (a "social consumer") and consumers who have also become addicted to the harmful product ("addicts"). Accordingly, "rehabilitation" (reinterpreted as disconnecting the consumer from the network) is effective when it occurs during the time the consumer is merely a "social consumer", while it is ineffective when it occurs too late and the consumer has already become addicted.

Keywords

Dynamic programming, peer pressure, network effect, optimal consumption

P-04

Impact of the Private Observation in Persuasion Game

Rony Bou Rouphael, Maël Le Treust

ETIS UMR 8051, Université Paris Seine, Université Cergy-Pontoise, ENSEA, CNRS, 6 avenue du Ponceau, 95014 Cergy-Pontoise CEDEX, France

rony.bou-rouphael@ensea.fr, mael.le-treust@ensea.fr

In [1], Kamenica-Gentzkow investigate a persuasion game in which the sender observes the realization of a state variable and commits to some signalling mechanism, then the receiver chooses the best-reply action corresponding to its posterior belief. We consider the repeated version of the persuasion game with two variations: 1) the communication may be constrained by a limited channel capacity and messages distorted by some source of noise, as in [2], 2) the receiver privately observes a signal correlated to the state, as in the source coding problem of Wyner-Ziv in [3]. The optimal solution to this setting relies on a specific concavification involving an auxiliary utility function for the sender, as in [4]. In this work, we examine the impact of the receiver's private signal on the sender's utility. In order to do so, we investigate an example involving a binary state, a binary private signal and a binary receiver's actions. We identify the optimal splitting of the receiver's beliefs satisfying the channel capacity constraint, and we compute the sender's optimal utility value, with and without private signal. Varying the parameters such as the prior, the precision of the private signal and the channel capacity, we aim at determining which settings are more favorable to the sender.

Keywords

Bayesian persuasion, Communication channel, Noisy channel, Posterior beliefs, Optimal strategies, State information, Concavification.

References

[1] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *American Economic Review*, vol. 101, pp. 2590 – 2615, 2011.

[2] M. Le Treust and T. Tomala, *Journal of Economic Theory [Online], Available: https://doi.org/10.1016/j.jet.2019.104940.*

[3] A. D. Wyner and J. Ziv, "The rate-distortion function for source coding with side information at the decoder," *IEEE Transactions on Information Theory*, vol. 22, no. 1, pp. 1–11, Jan. 1976.

[4] M. Le Treust and T. Tomala, "Information-theoretic limits of strategic communication," *Preliminary Draft, https://arxiv.org/abs/1807.05147, 2018.*, 2018.

Acknowledgement

The authors gratefully acknowledge the support of the Paris Seine Initiative, of the ETIS laboratory and of the UFR-ST of University of Cergy-Pontoise.

P-05

Finding automaton maximizing score in prisoner's dilemma tournament

Jana Syrovátková
Computer Science Institute of Charles University

syrovatkova@vso-praha.eu

Prisoner's dilemma is a well-known concept. Robert Axelrod initiated the study of tournaments of finite automata playing repeated prisoner's dilemma. Similar tournaments have appeared afterwards. The strategy is written in the form of a finite automaton, the number of rounds is given by a random geometric distribution. In the tournament every automaton plays with every other one. Our goal was to create an automaton that will maximize the score earned in the tournament for a given set of opponents.

In the first stage, we convert the opponent's finite automaton into an oriented graph describing the possible iterations of repeated play. In the second phase, we search the walk in this graph that maximizes the score. To play against one automaton, the best combination consists of a path and a cycle sharing one common vertex. Finding the best automaton can be done in time $O(n3)$, where n is the number of states of the opponent's automaton.

When playing against multiple automata, a graph with states from the Cartesian product of the states of the original automata is created. The states are also merged into groups according to the automata responses during the game. Within them, the solution is again the combination of path and cycle. If n is the number of states of the opponent's automaton with the most states and m the product of the number of states of all the automata, the search can be performed in time $O(mn)$.

Keywords

finite automata, prisoner's dilemma

Acknowledgement

In cooperation with my tutor Robert Sámal

P-06

On the Inspection of Bayesian Knowledge Sharing

František Hůla, Tatiana Valentina Guy, Miroslav Kárný
The Czech Academy of Sciences, Institute of Information Theory and Automation
Adaptive Systems Department

{hula,guy,school}@utia.cas.cz

Retailers nowadays face a great challenge in providing relevant personalised offers to their customers on the fly. Contemporary machine learning algorithms struggle with the processing of *big data* in real time. Thus there is an urge to develop systems capable of online learning. Bayes' rule emerges as an auspicious contender in this endeavor for it possesses online learning capabilities. This work extends Bayes' rule by employing knowledge sharing provided by a model which involves greater amounts of data whilst cannot learn online.

Specifically, a Bayesian agent accumulates knowledge for a prediction of consequences $c \in \mathsf{c}$ by using a parametric model $\mathsf{m}(c|\theta)$, $\theta \in \boldsymbol{\theta}$ and a chosen probability density (pd). The posterior pd $\mathsf{p}(\theta|K)$ on $\boldsymbol{\theta}$ quantifies the knowledge K consisting of already observed consequences c and prior knowledge. Bayes' rule is characterised by the updating formula $\mathsf{p}(\theta|c, K) \propto \mathsf{m}(c|\theta)\mathsf{p}(\theta|K)$. One can generalise [1, 2, 3] Bayes' rule as

$$\mathsf{p}(\theta|\mathsf{f}, K) \propto \mathsf{p}(\theta|K) \exp\left[\tau \int_{\mathsf{c}} \mathsf{f}(c) \ln(\mathsf{m}(c|\theta)) \mathrm{d}c\right], \tag{1}$$

where the predictive pd $\mathsf{f}(c)$ is supplied either by the more data-rich agent or results from Bayesian learning based on a different parametric model. Hence, rule (1) reduces to Bayes' rule if $\mathsf{f}(c)$ is dominant in the vicinity of an observed consequence and $\tau = 1$. The weight $\tau \geq 0$ expresses confidence in the knowledge produced by f. Extensive simulations of formula (1) are carried out.

Equation (1) is adapted to a Markov parametric model with discrete-valued data. The predictive pd $\mathsf{f}(c)$ is obtained by learning a different parametric model predicting consequences from a richer data set. The obtained results shall be put forth.

Keywords

Bayesian estimation, Bayesian prediction, knowledge sharing

References

[1] J. Kracík and M. Kárný, "Merging of data knowledge in Bayesian estimation," in *Proc. of the Second Int. Conference on Informatics in Control, Automation and Robotics*, J. Filipe, J. A. Cetto, and J. L. Ferrier, Eds. Barcelona: INSTICC, 2005, pp. 229–232.

[2] M. Kárný, J. Andrýsek, A. Bodini, T. Guy, J. Kracík, and F. Ruggeri, "How to exploit external model of data for parameter estimation?" *Int. J. of Adaptive Control and Signal Processing*, vol. 20, no. 1, pp. 41–50, 2006.

[3] A. Quinn, M. Kárný, and T. Guy, "Fully probabilistic design of hierarchical Bayesian models," *Inf. Sci.*, vol. 369, pp. 532–547, 2016.

Acknowledgement

The research has been supported by MSMT LTC18075 and by EU-COST Action CA16228.

P-07

Learning With Forgetting Preserving a Guaranteed Knowledge

Filip Rolenec, Miroslav Kárný
Institute of Information Theory and Automation of the CAS

rolenfil@fjfi.cvut.cz

This research concerns an advanced task of information processing in the problem of Bayesian parameter estimation. It focuses on Markov decision processes (MDPs) with an unknown time-variant transition probability function without explicit model of its evolution. The problem is approached by recursive Bayesian estimation with forgetting. The current methods of data-dependent forgetting, such as exponential [1], directional [2] and stabilised [3] are presented and discussed. With a help of minimum Kullback-Leibler principle and a theory of new hypothesis' probability assignment [4], a new generalisation of the stabilised forgetting is proposed.

The newly presented forgetting algorithm allows the agent to: 1) store any number of probability distributions, whose information is combined in one, in some sense, optimal distribution and 2) declare some distribution (or distributions) as *guaranteed knowledge*, which – by definition – is not forgotten.

The algorithm consists of initialization and an estimation-update loop whose each iteration consist of three steps. In the first step, assuming some data was observed, a new distribution is constructed as Bayesian update of the *guaranteed knowledge* on this data. In the second step another distribution is created as an optimal representative of the currently hold set of distributions. The final step of the algorithm concerns the discarding procedure, which needs to be used when the number of currently held distributions is larger than the allowed memory limit of the agent. This procdure is conducted with regard to the declared *guaranteed knowledge* and the assigned beliefs to each of the kept distributions.

Experiments show, that the newly developed forgetting algorithm is robust with regard to the poorly informative data. The performed simple experiments, however, do not prove the superiority of the newly developed method in comparison to the current methods, namely directional and stabilised forgetting. This conclusion motivates the first step of the future research as testing the newly developed algorithm on MDPs with a more complex model of parameter evolution. The second step was recognized as testing the new algorithm on MDPs with continuous action and state sets. The study of a dependency of the prediciton accuracy on the number of preserved parameter distributions and a possible time-variant declaration of *guaranteed knowledge* are other future research paths.

Keywords

Bayesian Estimation, Directional Forgetting, Exponential Forgetting, Guaranteed Knowledge, Markov Decision Process, Minimum Cross-entropy Principle, Stabilized Forgetting.

References

[1] V. Peterka, "Bayesian approach to system identification," in *in Trends and Progress in System Identification, P. Eykhoff, Ed.* Pergamon Press, 1981, pp. 239–304.

[2] R. Kulhavý and M. Kárný, "Tracking of slowly varying parameters by directional forgetting," in *Preprints of the 9th IFAC World Congress.* Budapest: IFAC, 1984, vol. X, pp. 178–183.

[3] R. KULHAVY and M. B. ZARROP, "On a general concept of forgetting," *International Journal of Control*, vol. 58, no. 4, pp. 905–924, 1993.

[4] M. Kárný, "On assigning probabilities to new hypotheses," *IEEE Signal Processing Letters*, 2019, submitted.

Acknowledgement

I would like to thank Grant MSMT LTC18075 for the support of my work and this field of study in general.

P-08

Learning in Dynamic Routing Games with Symmetric Incomplete Information

Emilien Macault[1], Marco Scarsini[2], Tristan Tomala[1]
1:HEC Paris, 2:LUISS

emilien.macault@hec.edu, tomala@hec.fr, marco.scarsini@luiss.it

We introduce a model of dynamic routing games under symmetric incomplete information. It consists in a routing game in which cost functions are determined by an unknown state of the world, as in [1] and [2], to the difference that the game is played repeatedly in discrete time. At each stage, a new mass of players routes over the network, publicly revealing its realized equilibrium costs. Our objective is to study how information aggregates according to the equilibrium dynamics and to which extent a centralized information system can learn about the state of the world. We define several forms of learning – whether agents eventually learn the state of the world or act as in full information – and provide simple examples showing that in the general case, with a non-atomic set of players, learning may fail and routing may be inefficient even in the looser sense. This contrasts with the atomic case, in which a folk theorem ensures players can learn the game parameters. In a non-atomic setup, learning cannot be ensured unless there is an additional source of randomness to incentivize exploration of the network. We show that this role can be fulfilled by an exogenous source of randomness in the game. We first explore the case of a variable and unbounded demand size. Our main result proves that under some network topology condition as in [3] and costs unboundedness, a variable and unbounded demand is sufficient to ensure learning. This result can be adapted to infinite state spaces, whether costs are observed edge by edge or path by path. We additionally provide examples to show these conditions are tight. In a second variant of our model, we connect our work with the behavioural learning literature *à la* Smith and Sorensen [4] and show that if instead of having random demand size, costs are observed with some unbounded noise, then learning may occur under some conditions on the distribution of noises.

Keywords

Routing games, Repeated games, Incomplete information, Rational learning

References

[1] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar, "Informational braess' paradox: The effect of information on traffic congestion," *Oper. Res.*, vol. 66, no. 4, pp. 893–917, Aug. 2018. [Online]. Available: https://doi.org/10.1287/opre.2017.1712

[2] M. Wu, S. Amin, and A. E. Ozdaglar, "Value of information systems in routing games," 2018.

[3] I. Milchtaich, "Network topology and the efficiency of equilibrium," *Games and Economic Behavior*, vol. 57, no. 2, pp. 321–346, 2006.

[4] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.

Acknowledgement

This work has been partly supported by COST Action CA16228 European Network for Game Theory.

P-09

Revealing Private Information in a Patent Race*

Pavel Kocourek†
New York University

Job Market Paper

November, 2018

Abstract

In this paper I study dynamic and strategic aspects of R&D rivalry. I consider a patent race in which the first firm to make two consecutive breakthroughs wins the prize. A breakthrough arrives with instantaneous probability equal to the firm's R&D effort level, and its arrival is observed privately. A firm varies its effort as it updates its belief about the rival's progress. I find that a firm drops its effort over time until its first breakthrough arrives, in which case the effort jumps up and keeps increasing until one of the players patents. Further, I investigate whether a firm would want to reveal success in order to discourage its rival. I find that a firm never reveals if its rival has, and is first to reveal when a breakthrough is hard to achieve. When breakthroughs arrive quickly the firm prefers secrecy to revelation. For intermediate levels of research difficulty firm's revelation behavior entails randomization or delay. Interestingly, when there are more than two players, equilibrium always entails revelation.

*I am indebted to Ennio Stacchetti for his invaluable guidance and support throughout this project. My most sincere gratitude to Boyan Jovanovic, Erik Madsen, Jaroslav Borovička, Dilip Abreu, Sylvain Chassang, and Rumen Konstadinov for continued discussions and suggestions which helped improve this paper greatly. I am also grateful to seminar participants at NYU for their comments and suggestions. All mistakes are my own.

†Department of Economics, NYU. Email: pk1050@nyu.edu

P-10

Preference Elicitation for Markov Decision Processes within Fully Probabilistic Design Framework

M. Ruman, T. V. Guy, M. Kárný

Institute of Information Theory and Automation, AS CR

ruman@utia.cas.cz, guy@utia.cas.cz, school@utia.cas.cz

Recently, a preference elicitation for fully probabilistic design (FPD) of decision strategies has been proposed. The adopted FPD framework converts the preference elicitation to constructing an ideal probability. This ideal probability descrics the desired closed-loop behaviour formed by all involved random variables. The choice of this ideal closed-loop model has been reduced to a minimisation task over a set of prospective ideals. The possibility that this set may be empty was left aside. Usual internal inconsistencies of the formulated preferences make this case frequent. The current paper solves with this problem. It applies the solution to Markov decision processes with discrete-valued closed-loop behaviours, which are quite vulnerable to inconsistencies. The result is broadly applicable as FPD is provably a proper dense extension of standard Bayesian decision making.

Keywords

Fully Probabilistic Design, Preference Elicitation, Markov Decision Processes, Kullback Leibler Divergence

Acknowledgement

The research has been supported by MSMT project LTC18075 and COST Action CA16228.

P-11

Constrained fully probabilistic control

Davide Gagliardi, Giovanni Russo
University College of Dublin

davide.gagliardi@ucd.ie, giovanni.russo1@ucd.ie

The work we present with this poster proposes a modification of the *Fully Probabilistic Design* (FPD) algorithm that has been pioneered by Prof. Kárný and his group, see e.g. [1], [2], [3] and references therein. In essence, the key idea of the FPD is that of controlling a system of interest, modeled via a probability density function (pdf), so that the Kullback-Leibler divergence (KLD) between the closed-loop system pdf and an ideal pdf is minimized. In turn, the ideal pdf models the desired evolution for the system.

In this context, we consider the problem of minimizing the KLD between the system and the ideal pdfs subject to additional constraints on the control variable. This is a problem of practical interest in many fields where the ideal distribution is given (i.e. it cannot be changed) and does not take into account the effective capabilities of the system under control. In an unconstrained setting, as shown in e.g. [4], small parametric changes can have dramatic effects on the closed-loop system, even destabilizing a system that otherwise would be stable. This can be instead avoided if the variables involved in the problem are constrained. Another case can be represented by a system whose behaviour should be optimized through a different performance index (see e.g.[5]), than the one used as a current ideal reference.

We present here a closed analytical solution for constrained problem. Interestingly, our solution keeps the same structure as the classic FPD solution and, when there are no constraints, the two solutions coincide.

Finally, our theoretical results are complemented with simulations.

Keywords

Uncertain systems; Optimal control; Fully Probabilistic Design

References

[1] M. Kárný, "Towards fully probabilistic control design," *Automatica*, vol. 32, no. 12, pp. 1719–1722, 1996.

[2] M. Kárný and T. V. Guy, "Fully probabilistic control design," *Systems & Control Letters*, vol. 55, no. 4, pp. 259–265, 2006.

[3] M. Kárný and T. Kroupa, "Axiomatisation of fully probabilistic design," *Inf. Sci.*, vol. 186, no. 1, pp. 105–113, 2012.

[4] B. G. Peguerole and G. Russo, "Sensitivity and safety of fully probabilistic control," in *2019 European Control Conference*, 2019.

[5] D. Gagliardi, P. Tkachenko, and L. del Re, "Outcome oriented evaluation of autonomous driving functions," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 6970–6975.

P-12

Winning fast when both Maker and Breaker are walkers

Jovana Forcan, Mirjana Mikalački
University of Novi Sad, Faculty of Sciences, Department of Mathematics and Informatics, Novi Sad, Serbia.

dmi.jovana.jankovic@student.pmf.uns.ac.rs, mirjana.mikalacki@dmi.uns.ac.rs

We study (1 : 1) WMaker–WBreaker games, a variant of Maker–Breaker positional games. Two players WMaker (taking the role of Maker) and WBreaker (taking the role of Breaker) alternately claim edges of the complete graph K_n, with the constraint that they need to choose their edges according to a walk. Espig, Frieze, Krivelevich and Pegden introduced in [1] the Walker–Breaker games, a variant of Maker–Breaker games in which only Maker is constrained to choose edges of a walk. Walker–Breaker games are studied only in [2, 1], so there are lots of questions that are still open. It is proven that the largest path Walker is able to create in (1 : 1) game on K_n has $n-2$ vertices [1]. So, Walker is not able to create a spanning structure. In the same paper [1], the authors asked the following question: *what happens if Breaker is also a walker.* Here we address this question and consider two standard games – the Connectivity game and the Hamilton cycle game. We show that WMaker can create both spanning structures. Moreover, we are curious to see how quickly WMaker can do this, i.e. what is the minimum number of moves that WMaker needs to win in both games. We prove that in the unbiased (1 : 1) WMaker–WBreaker Connectivity game on K_n, WMaker can make a spanning tree in at most $n+1$ moves. Also, we show that WMaker can make a Hamilton cycle in at most $n+6$ moves.

Keywords

positional games, Maker–Breaker games, connectivity, Hamilton cycle

References

[1] L. Espig, A. Frieze, W. Pegden, and M. Krivelevich, "Walker–Breaker games," *SIAM Journal on Discrete Mathematics*, vol. 29, no. 3, pp. 1476–1485, 2015.

[2] D. Clemens and T. Tran, "Creating cycles in Walker–Breaker games," *Discrete Mathematics*, vol. 339, no. 8, pp. 2113–2126, 2016.

P-13

Convergence and Hardness of Strategic Schelling Segregation

Hagen Echzell[2], Tobias Friedrich[1], Pascal Lenzner[1], Louise Molitor[1], Marcus Pappik[2], Friedrich Schöne[2], Fabian Sommer[2], and David Stangl[2]
Hasso Plattner Institute, University of Potsdam, Potsdam, Germany
[1]firstname.lastname@hpi.de, [2]firstname.lastname@student.hpi.de

The phenomenon of residential segregation was captured by Schelling's famous segregation model where two types of agents are placed on a grid and an agent is content with her location if the fraction of her neighbors which have the same type as her is at least τ, for some $0 < \tau < 1$. Discontent agents simply swap their location with a randomly chosen other discontent agent or jump to a random empty cell.

We analyze a generalized game-theoretic model of Schelling segregation which allows more than two agent types and more general underlying graphs modeling the residential area. For this we show that both aspects heavily influence the dynamic properties and the tractability of finding an optimal placement. We map the boundary of when improving response dynamics (IRD), i.e., the natural approach for finding equilibrium states, are guaranteed to converge. For this we prove several sharp threshold results where guaranteed IRD convergence suddenly turns into the strongest possible non-convergence result: a violation of weak acyclicity. In particular, we show such threshold results also for Schelling's original model, which is in contrast to the standard assumption in many empirical papers. Furthermore, we show that in case of convergence, IRD find an equilibrium in $\mathcal{O}(m)$ steps, where m is the number of edges in the underlying graph and show that this bound is met in empirical simulations starting from random initial agent placements. See [1] for more info.

Keywords

Schelling Segregation, Convergence of Improving Response Dynamics, Potential Games, Computational Hardness

References

[1] H. Echzell, T. Friedrich, P. Lenzner, L. Molitor, M. Pappik, F. Schöne, F. Sommer, and D. Stangl, "Convergence and hardness of strategic schelling segregation," *arXiv:1907.07513*, 2019.

NOTES: