
Game theoretic learning using the imprecise Dirichlet model

Erik Quaeghebeur & Gert de Cooman

{Erik.Quaeghebeur,Gert.deCooman}@UGent.be

*SYSTeMS research group
Ghent University, Belgium*

My promoter, myself and my research

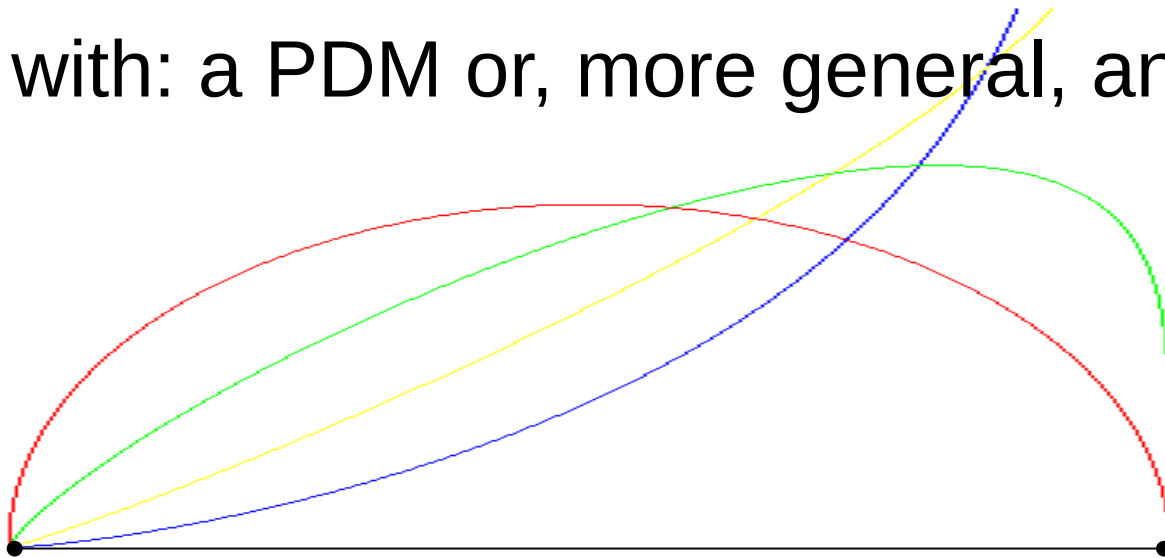
- Presented research: master's thesis cont'd
 - PhD-research started last fall
 - Current research: Using the IDM for learning in Markov models
 - Research interests: the IDM and its applications, learning models
 - Research detour: imprecise central moments
-

Two players in strict competition, but hey, it's only a game

- Yourself and one opponent
 - His loss, your gain (and vice-versa)
 - Playing: choosing a strategy
 - Afterwards: the pay-off, positive or negative
 - Strategies: from pure to mixed
-
- The expected payoff
-


Model your uncertainty, take an IDM

- You, the player, think/suppose that your *opponent plays an unknown, fixed, strategy*
- Why uncertainty in the model: to allow you to make an informed strategy choice yourself
- Model with: a PDM or, more general, an IDM



Updating the IDM

- Gathering information: observing the pure strategies your opponent plays
- Update your IDM with the gathered observations



To play, we need to pick an optimal strategy

- Optimal: maximise immediate expected pay-off, perhaps minimise risk (limit losses)
 - Use IDM and pay-off function to order the gambles
 - One optimal strategy or a set of optimal strategies (partial order)
 - Optimal set: no further choice, but an arbitrary one
-

One opponent, one game, playing over and over again

- Equilibrium of a game: special couple of strategies, if only you change your strategy, you'll get less (idem for your opponent)
 - In some cases, for a special type of equilibrium, convergence to such an equilibrium is guaranteed
 - In all cases, if the played strategies converge, it is to an equilibrium
-

Conclusions

- What we did: generalise a learning model, replacing PDM by IDM, complete ordering of strategies by a partial ordering
 - The resulting learning model has similar properties with regard to convergence to equilibria
 - We obtain a more complex, but also more expressive learning model
-

Questions: fire away

Slide reference:

- 1) Presenting myself and my research
 - 2) Defining games: strategies and pay-off
 - 3) Modelling uncertainty: IDM
 - 4) Updating an IDM
 - 5) Optimal strategies
 - 6) Convergence to equilibria
 - 7) Conclusions
-