Sharing of knowledge and preferences among imperfect Bayesian decision makers

Miroslav Kárný, Tatiana V.Guy
Department of Adaptive Systems
Institute of Information Theory and Automation
P.O.Box 18, Prague 182 08, Czech Republic
school@utia.cas.cz, guy@ieee.org

Abstract

Bayesian decision theory provides a strong theoretical basis for a single-participant decision making under uncertainty, that can be extended to multi-participant problems. However Bayesian decision theory assumes unlimited abilities of a participant to probabilistically model participant's environment and to optimise decision-making strategy.

The paper proposes a methodology for sharing of knowledge and strategies among participants, that helps to overcome the non-realistic assumption on participants' unlimited abilities.

1 Introduction

Dynamic decision making (DM) with uncertain information and limited resources is a *dynamic interaction* of a *participant* (decision maker) with its *environment* (part of the World). During the interaction the participant selects among available actions while aiming to reach its DM goals. These goals implicitly express the participant's preferences with respect to the future behavior of the closed loop formed of participant and the environment. The knowledge available to the participant comprise: the knowledge gained from the environment in response to the participants actions (observations); the knowledge associated with the participant's decisions (DM strategy) and knowledge considered by the participant (prior knowledge of the environment). The participant's limited cognitive and acting resources characterise participants *imperfectness*.

Unlike many other approaches to DM, Bayesian decision theory with its solid axiomatic basis proposes a systematic treatment of considered DM problem under uncertainty: given a complete probabilistic description of the environment and participant's preferences, the optimal strategy can be found explicitly. The assumption on availability of the descriptions is however quite restrictive as the participant operates (at most) with a part of domain-specific knowledge rising from its interaction with the environment and preferences expressed in domain-specific terms. The limited cognition resources of the imperfect participant prevent it to make the proper inferences from this limited and uncertain knowledge and to transfer it onto the relevant probabilistic descriptions. This calls for reliable and effective *knowledge elicitation* as well as the *preferences elicitation*.

Despite the former problem has been repeatedly treated, a variety of knowledge elicitation methods proposed [16] heavily depends on the quality of (also imperfect) domain experts. This makes it difficult and costly to select the proper method in a particular case. Theoretical and algorithmic support of preference elicitation problem remains to be a fundamental problem and no ready solution exists. One of the promising approaches considers preferences elicitation as an independent DM problem, which optimises the amount of elicitation effort/time (cost of elicitation) with respect to a gain provided by elicited preferences information (decision quality) [17].

This paper concerns a multi-participant DM problem with several imperfect selfish participants. Each of them solves own DM task given by its DM goal. The participants may cooperate and/or compete to achieve their personal goals. The participants may also be engaged in collaborative DM, i.e. they may have a common goal. The decentralised settings suggests the imperfect participant consider how its imperfect neighbours will behave in order to achieve their DM goals efficiently. Practically it means the participant should model knowledge and preferences of its neighbours, which contradicts participant's limited resources.

The paper proposes a way how to share the knowledge and preferences between imperfect selfish participants. The proposed approach formulates and solve the task as a specific DM problem. The solution does not force the participants to increase complexity of their models of environment and preferences while allowing them to handle partially incompatible, fragmental knowledge pieces.

2 Elements of Bayesian Decision Making

A participant selects and uses DM strategy $\mathsf{S} \in \mathbf{S} \neq \emptyset$ that maps available knowledge $K \in \boldsymbol{K}$ on the participant's actions $A \in \boldsymbol{A} \neq \emptyset$ to reach a preferred closed-loop behavior $B = (D, X, A)$, where $D \in \boldsymbol{D}$ denotes observed data and $X \in \boldsymbol{X}$ denotes considered but unobserved entities (internals). DM under uncertainty arises when the participant cannot uniquely assign B to S. Under the widely acceptable conditions [12, 1], the rationally chosen strategy ${}^{\mathsf{I}}\mathsf{S}$ minimizes the expectation E_{S} of a real-valued performance index I_{S}

$$ {}^{\mathsf{I}}\mathsf{S} \in \mathrm{Arg}\min_{\mathsf{S}\in\mathbf{S}} \mathsf{E}_{\mathsf{S}}[\mathsf{I}_{\mathsf{S}}] = \int_{\boldsymbol{B}} \mathsf{I}_{\mathsf{S}}(B)\, \mathsf{F}_{\mathsf{S}}(B) \mu(dB). \tag{1} $$

In (1), F_{S} is a Radon-Nikodým derivative (rnd) – with respect to a product, strategy-independent, measure μ – of a probabilistic measure describing the loop formed of the participant's environment and the DM strategy S. The closed-loop model can be factorised as $\mathsf{F}_{\mathsf{S}} = \mathsf{M} \times \mathsf{S}$ where the rnd M is a participant's model of its environment. The real-valued strategy-dependent mapping I_{S} specifies DM preferences while respecting participant's risk attitude. The further exposition replaces the traditional choice of I_{S} by defining an *ideal closed-loop model*, determined by the rnd ${}^{\mathsf{I}}\mathsf{F}$ and interpreted as the closed-loop model with the strategy ${}^{\mathsf{I}}\mathsf{S}$ (1), i.e., ${}^{\mathsf{I}}\mathsf{F} = \mathsf{F}_{{}^{\mathsf{I}}\mathsf{S}}$.

With ${}^{\mathsf{I}}\mathsf{F}$ chosen, the optimal strategy ${}^{O}\mathsf{S}$ is defined as a minimizer of the Kullback-Leibler divergence (KLD) $\mathsf{D}(\cdot||\cdot)$, [9], of the closed-loop model F_{S} on its ideal counterpart ${}^{\mathsf{I}}\mathsf{F}$

$$ {}^{O}\mathsf{S} \in \mathrm{Arg}\min_{\mathsf{S}\in\mathbf{S}} \mathsf{D}(\mathsf{F}_{\mathsf{S}}||\,{}^{\mathsf{I}}\mathsf{F}) = \int_{\boldsymbol{B}} \mathsf{F}_{\mathsf{S}}(B) \ln\left(\frac{\mathsf{F}_{\mathsf{S}}(B)}{{}^{\mathsf{I}}\mathsf{F}(B)}\right) \mu(dB), \;\; \mathsf{F}_{\mathsf{S}}(B) = \mathsf{M}(B)\mathsf{S}(B). \tag{2} $$

This fully probabilistic design (FPD) [6] is a dense extension of traditional designs (1), [7], so that we can focus on it further on (for a related learning perspective, see [14]). The rnds in (2) are implicitly conditioned on the knowledge available for the design. Use of FPD relies on the ability to:

- specify quantitatively DM elements, i.e.,
  - $\star$ knowledge $\boldsymbol{K}$, data $\boldsymbol{D}$, internals' $\boldsymbol{X}$ and actions' $\boldsymbol{A}$ sets on which the involved functions act,
  - $\star$ environment model given by the rnd M and used in the design (2),
  - $\star$ set of admissible strategies $\mathbf{S}$ among which the optimal strategy ${}^{O}\mathsf{S}$ (2) is searched for,
  - $\star$ an ideal closed-loop model ${}^{\mathsf{I}}\mathsf{F} = {}^{\mathsf{I}}\mathsf{M}\,{}^{\mathsf{I}}\mathsf{S}$ specifying DM preferences, constraints and risk attitude;
- evaluate the strategy ${}^{O}\mathsf{S}$, i.e., store, integrate and optimise functions in (2), and finally apply ${}^{O}\mathsf{S}$.

3 Support of Imperfect Bayesian participant

Support of computational and implementation aspects of DM is developing permanently. Conceptual complexity boundaries are being crossed by a broad use of distributed solutions. This allows us to focus on the neglected support of creation of the DM elements (the items marked by $\star$, see Section 2).

The selection of respective variables forming the considered behavior can be addressed by Bayesian testing of hypotheses, e.g. [1], or can be avoided via Bayesian averaging [11]. The choice of ranges of respective variables is less supported. A use of environment models having supports with unknown bounds seems to be an appropriate approach, e.g. [10].

When addressing a construction of the environment model and of the ideal rnd, the FPD allows the unified treatment due to the common probabilistic language. The subsequent discussion made in terms of general rnds is applicable to the environment model, the ideal rnd as well as to the factors in them (e.g. to prior rnd describing internals).

Design of a complete set of DM elements from the knowledge available to an imperfect participant is a specific decision task and as such is formulated and solved. We call this task *supporting DM* to distinguish from the original DM (supported DM). The supporting DM tasks are formulated via FPD. To simplify reading, elements of the original DM are denoted by small-letters counterparts of the general notation. Specifically, $b \in \boldsymbol{b}$ is the original closed-loop behavior and the rnd

$$\mathsf{f} \in \mathbf{f} \subset \mathbf{f}_{\triangle} = \left\{ \mathsf{f}(b) : \mathsf{f}(b) \geq 0, \int_{b \in \boldsymbol{b}} \mathsf{f}(b)\mu(db) = 1 \right\} \tag{3}$$

concerns the original DM task. To suppress technicalities, finite cardinality $|\boldsymbol{b}|$ of the behavior set $\boldsymbol{b} = \{b_1, \ldots, b_{|\boldsymbol{b}|}\}$ is assumed. Then, μ is a counting measure and the rnd f is a finite-dimensional vector belonging to the simplex (3). General validity of obtained results is conjectured.

Approximation of Rnds: Often, a rnd f constructed from the available knowledge is too complex to be treated by an imperfect participant and has to be approximated by a rnd $\hat{\mathsf{f}} \in \hat{\mathbf{f}} \subset \mathbf{f}_{\triangle}$ (3). The approximation is supporting DM problem with an action $A = \hat{\mathsf{f}}$, knowledge $K = \mathsf{f}$ of the approximated pdf. The corresponding closed-loop model, conditioned on the knowledge available for design of the optimal approximation strategy ${}^{O}\mathsf{S}(\hat{\mathsf{f}}|\mathsf{f})$, is $\mathsf{F}(b, \hat{\mathsf{f}}|\mathsf{f}) = \mathsf{f}(b)\mathsf{S}(\hat{\mathsf{f}}|\mathsf{f})$ and uses the fact that the rnd f models b. The ideal closed-loop model is specified as ${}^{\mathsf{I}}\mathsf{F}(b, \hat{\mathsf{f}}|\mathsf{f}) = \hat{\mathsf{f}}(b)\mathsf{S}(\hat{\mathsf{f}}|\mathsf{f})$. The choice of the first factor means that the approximating rnd is to describe ideally the original behavior. The second factor expresses a lack of additional requirements on the constructed $\mathsf{S}(\hat{\mathsf{f}}|\mathsf{f})$: the strategy resulting from the design is accepted as the ideal one. This type of choice was called leave-to-the-fate option, [5]. With this choice, the KLD (2) is linear in $\mathsf{S}(\hat{\mathsf{f}}|\mathsf{f})$ and reaches its minimum for the deterministic strategy providing the optimal approximation

$${}^{O}\hat{\mathsf{f}} \in \operatorname{Arg}\min_{\hat{\mathsf{f}} \in \hat{\mathbf{f}}} \mathsf{D}(\mathsf{f}||\hat{\mathsf{f}}). \tag{4}$$

This approximation principle was justified, using rather different arguments, in [2].

Minimum KLD Principle: The approximation discussed above requires complete knowledge of the approximated pdf f. Here, the considered knowledge about the approximated f is more vague

$$K: \quad \mathsf{f} \in \mathbf{f} \subset \mathbf{f}_{\triangle} \text{ (3) and a rnd } \mathsf{f}_0 \in \mathbf{f}_{\triangle} \text{ is an available prior (flat) guess of } \mathsf{f}. \tag{5}$$

The incomplete knowledge of f implies that this rnd belongs to unobserved entities (internals) $X = \mathsf{f}$ of the supporting DM problem. The corresponding action is a probabilistic estimate of f, i.e., $A = \mathsf{F}(\mathsf{f}) \in \mathbf{F}$ with

$$\mathbf{F} = \{\mathsf{F}(\mathsf{f}) : \text{rnds with support on } \mathbf{f} \subset \mathbf{f}_{\triangle}\}. \tag{6}$$

The behavior consists of $B = (b, \mathsf{f}, A)$. The corresponding closed-loop model is then

$$\mathsf{F}(b, \mathsf{f}, A) = \mathsf{F}(b|\mathsf{f}, A)\mathsf{F}(\mathsf{f}|A)\mathsf{F}(A) = \mathsf{f}(b)\, A\, \mathsf{S}(A|K), \tag{7}$$

where meaning of f as a model of $b \in \boldsymbol{b}$ and the definition of A are exploited. The chosen ideal rnd

$${}^{\mathsf{I}}\mathsf{F}(b, \mathsf{f}, A) = {}^{\mathsf{I}}\mathsf{F}(b|\mathsf{f}, A)\, {}^{\mathsf{I}}\mathsf{F}(\mathsf{f}|A)\, {}^{\mathsf{I}}\mathsf{F}(A) = \mathsf{f}_0(b)\, A\, \mathsf{S}(A|K), \tag{8}$$

represents preferences of this supporting DM task: i) the rnd $\mathsf{f}_0(b)$ is the best (prior) guess of the model of the original behavior, see (5), ii) both action and strategy are left to their fate as no prior preferences exist among them . Consequently, the action and strategy enter the optimized KLD linearly. It implies that the optimal strategy and the action are deterministic with a full mass on

$${}^{O}\mathsf{f} \in \operatorname{Arg}\min_{\mathsf{f} \in \mathbf{f}} \mathsf{D}(\mathsf{f}||\mathsf{f}_0). \tag{9}$$

The result (9) coincides with the minimum KLD principle and reduces to the maximum entropy principle if f_0 is a uniform rnd. It was axiomatically justified in [13] for the set $\mathbf{f}$ specified by given values of given linear functionals on $\mathbf{f}_{\triangle}$ (3). The following alternative to the knowledge (5)

$$K: \mathsf{f} \in \mathbf{f} \subset \mathbf{f}_{\triangle} \text{ and a rnd } A_0 = \mathsf{F}_0(\mathsf{f}) \in \mathbf{F}, \text{ is available prior (flat) guess of } A \tag{10}$$

does not change the closed-loop model (7) but leads to the ideal rnd, differing from (8),

$$^{\mathsf{I}}\mathsf{F}(b, \mathsf{f}, A) = {}^{\mathsf{I}}\mathsf{F}(b|\mathsf{f}, A)\,{}^{\mathsf{I}}\mathsf{F}(\mathsf{f}|A)\,{}^{\mathsf{I}}\mathsf{F}(A) = \mathsf{f}(b) A_0 \mathsf{S}(A). \tag{11}$$

It respects that f models b, takes $A_0 = \mathsf{F}_0(\mathsf{f})$ as the best prior guess of $A = \mathsf{F}(\mathsf{f}|A)$ and leaves the strategy $\mathsf{S}(A)$ to its fate. The resulting choice of strategy generalises the minimum KLD principle

$$^{O}\mathsf{F} \in \operatorname{Arg}\min_{\mathsf{F}\in\mathbf{F}} \int_{(\boldsymbol{b},\mathbf{f})} \mathsf{f}(b)\mathsf{F}(\mathsf{f}) \ln\left(\frac{\mathsf{F}(\mathsf{f})}{\mathsf{F}_0(\mathsf{f})}\right) \mu(d(b,\mathsf{f})), \ \ \mathbf{F} = \text{rnds acting on } \mathbf{f}, \text{ see (6)}. \tag{12}$$

Use of (4), (9) and (12) for constructing DM elements is exemplified below.

Extension of Non-Probabilistic Knowledge: Prior non-probabilistic knowledge can be often expressed by restricting $\mathbf{f}_\triangle$ to the set $\mathbf{f}$ in (5) for which functionals $\mathsf{E}_\mathsf{f}[\phi_\kappa] = \int_{\boldsymbol{b}} \phi_\kappa(b)\mathsf{f}(b)\mu(db) = 0$, $\kappa \in \boldsymbol{\kappa} = \{1, 2, \ldots, |\boldsymbol{\kappa}|\},\ |\boldsymbol{\kappa}| < \infty$. Indeed, participants often exploit deterministic models resulting from the first principles and domain-specific knowledge. They are mostly expressed by a set of equations $\phi_\kappa(b) = \epsilon_\kappa(b)$, where $\epsilon_\kappa(b)$ is a modelling error. Constraints $\mathsf{E}_\mathsf{f}[\phi_\kappa] = 0$ then simply express an expectation that modelling error is unbiased. Known ranges $\boldsymbol{\epsilon}_\kappa(b)$ of errors can modelled in the same way. It suffices to take $\phi_\kappa(b) =$ indicator of $\boldsymbol{\epsilon}_\kappa(b)$. If the expectation that modelling errors are out of this range is too high, the second moments serve well for error characterization. Known ranges of variables forming behavior can be respected via range indicators or the second moments similarly as modeling errors. After specifying the set $\mathbf{f}$, minimum KLD principle (9) is applied and possibly followed by an approximation of the obtained $\mathsf{f} = {}^{O}\mathsf{f}$ by a feasible $\hat{\mathsf{f}}$ according to (4). The needed prior guess f_0 is chosen as a soft delimitation of the support $\boldsymbol{b}$ of the involved rnds. An algorithmic implementation well supports imperfect participants, e.g. [4].

Combination of Incompletely Compatible Rnds: The set $\mathbf{f}$ specified by conditions $\mathsf{E}_\mathsf{f}[\phi_\kappa] = 0$, $\forall \kappa \in \boldsymbol{\kappa}$, can be empty as the processed knowledge pieces are incompletely compatible. Then, a meaningful solution of (9) does not exist. By considering various "compatible" subsets of these conditions, say considering them individually, we get a collection of different rnds $\mathsf{f}_\kappa \in \mathbf{f}_\triangle$ (3) that have to be combined into a single representant $\hat{\mathsf{f}}$. This is a prototype of DM task that has to be resolved when supporting imperfect participants. Especially, it is believed to be an efficient tool for solving, otherwise extremely hard, problems of de-centralized decision making [3]. The representant $\hat{\mathsf{f}}$ is found via the generalized KLD principle (12). The behavior $b \in \boldsymbol{b}$ is assumed to be described by an unknown rnd $\mathsf{f} \in \mathbf{f} \subset \mathbf{f}_\triangle$, where $\mathbf{f}$ is delimited by the specifying (10)

$$\begin{aligned} K: \quad & \mathsf{E}_\mathsf{F}[\mathsf{D}(\mathsf{f}_\kappa||\mathsf{f})] \le \beta_\kappa < \infty,\ \kappa \in \boldsymbol{\kappa} = \{1, \ldots, |\boldsymbol{\kappa}|\},\ |\boldsymbol{\kappa}| < \infty, \\ & \mathsf{F}_0(\mathsf{f}) = \text{ prior (flat) guess of the action },\ \text{see (6), and } \mathsf{f}_\kappa(b) \text{ are given rnds in } \mathbf{f}_\triangle. \end{aligned} \tag{13}$$

The constraints on the expected KLD of f_κ on f state that the rnd $\mathsf{f} = \hat{\mathsf{f}}$ is acceptable compromise with respect to a given f_κ if it is its good approximation, cf. (4). Under constraints (13), the optimal action ${}^{O}\mathsf{F} \in \mathbf{F}$, cf. (6), defined by (12), minimizes Kuhn-Tucker functional, given by multipliers λ_κ,

$$^{O}\mathsf{F} \in \operatorname{Arg}\min_{\mathsf{F}\in\mathbf{F}} \int_{(\boldsymbol{b},\mathbf{f})} \mathsf{F}(\mathsf{f}) \left[\mathsf{f}(b)\ln\left(\frac{\mathsf{F}(\mathsf{f})}{\mathsf{F}_0(\mathsf{f})}\right) + \sum_{\kappa\in\boldsymbol{\kappa}} \lambda_\kappa \mathsf{f}_\kappa(b) \ln\left(\frac{\mathsf{f}_\kappa(b)}{\mathsf{f}(b)}\right)\right] \mu(d(b,\mathsf{f})). \tag{14}$$

It gives ${}^{O}\mathsf{F}(\mathsf{f}) \propto \mathsf{F}_0(\mathsf{f}) \prod_{b\in\boldsymbol{b}} \mathsf{f}(b)^{\rho(b)}$ with $\rho(b) = \sum_{\kappa\in\boldsymbol{\kappa}} \lambda_\kappa \mathsf{f}_\kappa(b),\ \lambda_\kappa \ge 0$ respect inequalities in (13).

For the conjugated prior guess of Dirichlet form $\mathsf{F}_0(\mathsf{f}) \propto \prod_{b\in\boldsymbol{b}} \mathsf{f}(b)^{\nu_0(b)-1}$ with $\nu_0(b) > 0$, $\int_{\boldsymbol{b}} \nu_0(b)\mu(db) < \infty$, the rnd ${}^{O}\mathsf{F}(\mathsf{f})$ (10) is also a Dirichlet rnd given by $\nu(b) = \nu_0(b) + \rho(b)$. This rnd has expected value, which is a "point" representant of incompletely compatible rnds f_κ, $\kappa \in \boldsymbol{\kappa}$,

$$\hat{\mathsf{f}}(b) = \mathsf{E}_{\,{}^{O}\mathsf{F}}[\mathsf{f}(b)] = \frac{\nu_0(b) + \sum_{\kappa\in\boldsymbol{\kappa}} \lambda_k \mathsf{f}_\kappa(b)}{\int_{\boldsymbol{b}} \nu_0(b)\mu(db) + \sum_{\kappa\in\boldsymbol{\kappa}} \lambda_\kappa} = \text{affine combination of the merged rnds.} \tag{15}$$

Extension of Fragmental Rnds: The combination of rnds (15) provides invaluable tool for sharing knowledge/preferences among participants indexed by $\kappa \in \boldsymbol{\kappa}$ [8]. Due to their imperfection, they may provide only a conditional (marginal) version of $\mathsf{f}_\kappa(b)$, i.e.,

$$b = (u_\kappa, m_\kappa, k_\kappa) = \text{behavior part (uninteresting for, modelled by, known to) the } \kappa\text{th rnd provider.} \tag{16}$$

Applicability of merging (15) requires an extension of rnds $\mathsf{f}_\kappa(m_\kappa|k_\kappa)$ provided by the members of knowledge/preference sharing group. As $\hat{\mathsf{f}}(b)$ (15) is the best compromise, the extensions ${}^{e}\mathsf{f}_\kappa(b)$ of $\mathsf{f}_\kappa(m_\kappa|k_\kappa)$ should be its best approximations. According to (4), they have to minimise the KLD of $\hat{\mathsf{f}}$. It gives unique extensions ${}^{e}\mathsf{f}_\kappa(b) = \hat{\mathsf{f}}(u_\kappa|m_\kappa, k_\kappa)\mathsf{f}_\kappa(m_\kappa|k_\kappa)\hat{\mathsf{f}}(k_\kappa)$. Inserting these extensions into (15), we get an equation for systematic combination of fragmental incompletely specified knowledge/preferences originating from different sources or different participants

$$\hat{\mathsf{f}}(b) = \mathsf{E}_{\mathcal{O}\mathsf{F}}[\mathsf{f}(b)] = \frac{\nu_0(b) + \sum_{\kappa\in\boldsymbol{\kappa}} \lambda_k \hat{\mathsf{f}}(u_\kappa|m_\kappa, k_\kappa)\mathsf{f}_\kappa(m_\kappa|k_\kappa)\hat{\mathsf{f}}(k_\kappa)}{\int_{\boldsymbol{b}} \nu_0(b)\mu(db) + \sum_{\kappa\in\boldsymbol{\kappa}} \lambda_\kappa}.$$

The merger $\hat{\mathsf{f}}(b)$ can be either exploited externally [15] or projected back to domains of respective imperfect participants offering them corrected but understandable knowledge: they are not bothered by the u_κ-th part of the behaviour, see (16).

4 Open Problems in Construction of DM Elements

The paper proposes a methodology for sharing of knowledge and strategies among participants, that helps to overcome the non-realistic assumption on participants' unlimited abilities.

In addition to non-trivial conversion into algorithms, the future work should find answers on the following questions.
- Is the outlined approach the "best" one?
- What is extent of ambiguity in our assumptions and tools?
- What are relationships to alternative approaches?
- What will be an emergent behavior of a network of interacting participants using our tool set?
- Can we use our approach for modelling natural/societal systems?
- Can be Bayesian DM enriched so that approximations (projections) become its inherent part?

Acknowledgments The research has been supported by MŠMT ČR 1M0572, GAČR 102/08/0567.

References

[1] J.O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag, New York, 1985.

[2] J. M. Bernardo. Expected information as expected utility. *The Annals of Statistics*, 7(3):686–690, 1979.

[3] D.S. Bernstein, C. Amato, E.A. Hansen, and S. Zilberstein. Policy iteration for decentralized control of markov decision processes. *Journal of Artificial Intelligence Research*, 34:89 –132, 2009.

[4] M. Kárný, A. Bodini, T.V. Guy, J. Kracík, P. Nedoma, and F. Ruggeri. Fully probabilistic knowledge expression and incorporation. *Int. J. of Adaptive Control and Signal Processing*, 2010. submitted.

[5] M. Kárný, J. Böhm, T. V. Guy, L. Jirsa, I. Nagy, P. Nedoma, and L. Tesař. *Optimized Bayesian Dynamic Advising: Theory and Algorithms*. Springer, London, 2005.

[6] M. Kárný and T. V. Guy. Fully probabilistic control design. *Systems & Control Letters*, 55(4): 259–265, 2006.

[7] M. Kárný and T.V. Guy. Preference elicitation in fully probabilistic design of decision strategies. In *Proceedings of the 49th IEEE Conference on Decision and Control*. IEEE, 2010.

[8] M. Kárný, T.V. Guy, A. Bodini, and F. Ruggeri. Cooperation via sharing of probabilistic elements. *International Journal of Computational Intelligence Studies*, II(2):139–162, 2009.

[9] S. Kullback and R. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–87, 1951.

[10] L. Pavelková. *Estimation of Models with Uniform Innovations and its Application on Traffic Data*. PhD thesis, Czech Technical University in Prague, Faculty of Transportation Sciences, December 2008.

[11] A.E. Raftery, M. Kárný, and P. Ettler. Online prediction under model uncertainty via dynamic model averaging: Application to a cold rolling mill. *Technometrics*, pages 52–66, 2010.

[12] L.J. Savage. *Foundations of Statistics*. Wiley, New York, 1954.

[13] J. Shore and R. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, 26(1):26–37, 1980.

[14] N. Tishby and D. Ploani. Information theory of decisions and actions. In Vassilis, Hussain, and Taylor, editors, *Perception-reason-action cycle: Models, algorithms and systems*. Springer, 2010.

[15] D.H. Wolpert and N. Kulkarni. Game-theoretic management of interacting adaptive systems. In *Proc. 2008 NASA/ESA Conference on Adaptive Hardware and Systems*. 2008.

[16] Cooke, N. J. Varieties of knowledge elicitation techniques. *International Journal of Human-Computer Studies*, Vol. 41, pages 801–849, 1994.

[17] Chajewska, U., Koller, D. and Parr, R. Making rational decisions using adaptive utility elicitation In *Proc.of the Seventeenth National Conference on Artificial Intelligence.* pages 363-369, 2000.