

COOPERATIVE DECISION MAKING WITHOUT FACILITATOR

Miroslav Kárný, Jan Kracík, Tatiana V. Guy

** Adaptive Systems Department
Institute of Information Theory and Automation
Academy of Sciences of the Czech Republic
P. O. Box 18, 182 08 Prague, Czech Republic*

Abstract:

Estimation, learning, pattern recognition, diagnostics, fault detection and adaptive control are prominent examples of dynamic decision making under uncertainty. Under rather general conditions, they can be cast into a common theoretical framework labelled as Bayesian decision making. Richness of the practically developed variants stems from: (i) domain-specific models used; (ii) adopted approximations fighting with limited perceiving and evaluation abilities of the involved decision-making units, called here participants. While modelling is a well-developed art, the item (ii) still lacks a systematic theoretical framework. This paper provides a promising direction that could become a basis of such framework. It can be characterized as multiple-participant decision making exploiting Bayesian participants equipped with tools for sharing their knowledge and harmonizing their aims and restrictions with their neighbors. Intentional avoiding of the negotiation facilitator makes the solution fully scalable.

Keywords: Bayesian decision making, multiple participants, learning, cooperation, distributed decision making

1. INTRODUCTION

Decision making (DM), or problem solving, takes place almost in every activity concerning the World (biology, technics/industry, human, social). A particular decision maker (*participant*) forms and applies his DM strategy in order to meet his goals with respect to the behavior of his *environment*, which may include other participants, too. The participant uses available observations and model of his environment as well as quantitative description of his goals and restrictions. Such DM template inherently implies a number of controversies, when participant's DM elements (goals,

observations, models) are even partially incompatible with others. If the participants' environments are overlapping, controversies are inevitable and only a cooperation can overcome them.

Here a normative multi-participant version of DM with participants acting as *Bayesian participants* is considered. Bayesian paradigm, Berger (1985); Bertsekas (2001); Kárný et al. (2005), is proven to serve as a well-grounded DM framework. Under rather general conditions, it consists of Bayesian filtering of unobserved, but considered quantities (called *internals*), and dynamic programming serving to design of an optimal DM strategy. Description of the formal solution is relatively straightforward, but its practical application is difficult due to: (i) non-trivial transformation of

¹ This research has been partially supported by AV ČR 1ET 100 750 401, MŠMT ČR 2C06001.

knowledge, decision goals and restrictions into an appropriate probabilistic language, and (ii) always limited ability of participants to perceive and evaluate. While the first item can be often solved by a well-developed, domain-specific modelling, the second one is the application bottleneck, which can be overcome by a cooperative DM. The desired efficient approach, however, must not rely on a non-realistic facilitator with unlimited capabilities. The paper contributes to the problem solution and tries to provide the missed part of Bayesian DM theory. The proposed cooperative DM respects the following circumstances:

- Real participant with limited perceiving and evaluation abilities is equipped with, more or less standard, Bayesian methodology of constructing and applying feasible DM strategies.
- A participant has a relatively small set of neighbors he can cooperate with, i.e., the overall cooperation load remains within the ability of a real participant.
- The cooperation rests on two basic activities: *sharing* and *harmonization* of the DM elements. A tool set supporting these activities is provided to a participant.

The paper summarizes the state of the art in the outlined direction. Predominantly, it focuses on the steps, where either existing solutions are felt insufficient or insufficiently justified and points out directions of further systematic research. Specifically, Section 2 recalls a fully probabilistic version of the Bayesian DM, Kárný and Guy (2006), for a single participant. Section 3 outlines considered cooperation ways. Section 4 summarizes the available tools for knowledge sharing and the tools for harmonizing DM goals and restrictions. Critical aspects of the proposed scheme are discussed in Section 5. Section 6 closes the paper.

2. FULLY PROBABILISTIC DECISION MAKING

The data d available to a participant consists of observation Δ he made on his environment and his action a ; at discrete time t , $d_t \equiv (\Delta_t, a_t)$ (\equiv means *defining equality*). To reach his DM goal with respect to his environment, the participant selects sequence of informationally causal rules R_t , which forms his DM strategy $R \equiv R^T \equiv (R_1, \dots, R_T)$. Each rule $R_t \in R_t^*$ (x^* marks *set of x*) maps the available data $d^{t-1} \equiv (d_1, \dots, d_{t-1})$ on the action a_t , i.e., $R_t : d^{t-1} \rightarrow a_t^*$.

The sequence of actions the participant selects is $a^T \equiv (a_1, \dots, a_T)$ with individual actions $a_t \in a_t^* \equiv$ the set of admissible values applied at discrete time $t \in t^* \equiv \{1, \dots, T\}$, $T < \infty$.

Quality of the selected DM strategy is judged by a loss function $Z : \mathcal{Q}^* \rightarrow [-\infty, \infty]$ that orders *a posteriori closed-loop behaviors* $\mathcal{Q} \equiv (d^T, x^T)$ consisting of data d^T and internals x^T .

DM always faces uncertainty as his environment is not isolated part of the world and internals are unknown to the participant. Thus, his strategies are *a priori* ordered via the loss function only when considering a functional dependence of the behavior on DM strategy *and* uncertainty. This does not allow him to select the best DM strategy: influence of uncertainty on the ordering has to be removed.

Anybody wants to avoid selecting the DM strategy, which can be *a posteriori* outperformed irrespectively of uncertainties. To achieve this, the *mathematical expectation* $E[\cdot] \equiv \int \cdot f(\mathcal{Q}) d\mathcal{Q}$ should be used for prior elimination of uncertainty. The weighting function $f(\mathcal{Q})$ of the expectation functional is *probability density function* (pdf) describing the closed decision loop for the inspected DM strategy. Often, the expected loss $E[Z]$ is taken as the prior measure of quality. Generally, however, the attitude of the participant to losses is modified by the probability assigned to realized behavior, i.e., the optimal DM strategy ${}^O R$ is defined

$${}^O R \in \text{Arg} \inf_{R \in R^*} \int \mathcal{U}(Z(\mathcal{Q}), f(\mathcal{Q})) f(\mathcal{Q}) d\mathcal{Q} \quad (1)$$

with a function \mathcal{U} increasing in the first argument. Hereafter, $\int \cdot d\mathcal{Q}$ denotes *definite multivariate integral* over the domain of integrand.

The uncertain behavior \mathcal{Q} enters the function \mathcal{U} via the loss function Z . The postulate that otherwise the behavior enters it at most via $f(\mathcal{Q})$ reflects the assumption that the risk attitude depends only on probability that \mathcal{Q} will occur.

Neither \mathcal{U} nor Z are usually fully determined by the partial ordering of behaviors they express. It is reasonable to decrease this ambiguity and use an alternative expression of this ordering. In that sense, the combined effect of functions \mathcal{U} , Z can be expressed by introducing the notion of the *ideal pdf* ${}^I f(\mathcal{Q})$ as the pdf $f(\mathcal{Q})$ achieved for the DM strategy ${}^O R$ optimal with respect to them.

In the rest of the paper, we assume that the ideal pdf is unique. This assumption is restrictive. It has been partially removed under specific circumstances, Kárný et al. (2005), but a general solution without it is to be elaborated.

Assuming that \mathcal{U} has continuous derivative with respect to the second argument and requiring

$$\mathcal{U}(Z(\mathcal{Q}), {}^I f(\mathcal{Q})) = 0 \quad (2)$$

we get, cf. Bernardo (1979), the Kullback-Leibler divergence (KLD), Kullback and Leibler (1951), as the optimized functional

$$\mathcal{U}(Z(\mathcal{Q}), f(\mathcal{Q})) = \ln \left(\frac{f(\mathcal{Q})}{If(\mathcal{Q})} \right) \Leftrightarrow \quad (3)$$

$$\mathbb{E}[\mathcal{U}(Z, f)] = D(f||If) \equiv \int \ln \left(\frac{f(\mathcal{Q})}{If(\mathcal{Q})} \right) f(\mathcal{Q}) d\mathcal{Q}.$$

Specification of the DM goals and restrictions (support of $f(\mathcal{Q})$ giving the finite KLD is included in the support of $If(\mathcal{Q})$) via the *ideal pdf* $If(\mathcal{Q})$ and minimization of the KLD is called *fully probabilistic design (FPD)*, Kárný and Guy (2006).

Respecting the structure of the behavior $\mathcal{Q} = (d^T, x^T)$, the pdf $f(\mathcal{Q})$ describing closed loop, formed by the environment and by the DM strategy, can be factorized and interpreted as follows

$$f(\mathcal{Q}) = \left. \begin{aligned} &= \prod_{t \in t^*} \underbrace{f(\Delta_t | a_t, d^{t-1}, x^t)}_{\text{observation model}} \\ &\times \underbrace{f(x_t | a_t, d^{t-1}, x^{t-1})}_{\text{evolution model of internals}} \\ &\times \underbrace{f(a_t | d^{t-1})}_{\text{environment model}} \end{aligned} \right\} \quad (4)$$

informationally causal optimized DM strategy

Having the ideal pdf factorized in the way mimic to (4), the solution of FPD is described by the following pair of propositions proven in Kárný and Guy (2006). For presentation simplicity, d^{t-1} is omitted in condition and internals are assumed to have character of the state, i.e., x_t replaces x^t .

Proposition 2.1. (Stochastic filtering). Let a prior pdf $f(x_0)$ be given. Then, the pdf $f(x_t | d^t)$, describing *state estimate*, and the pdf $f(x_t | a_t, d^{t-1})$, providing *state prediction*, evolve

$$\begin{aligned} &f(x_t | a_t, d^{t-1}) \\ &= \int f(x_t | a_t, x_{t-1}) f(x_{t-1} | d^{t-1}) dx_{t-1} \\ f(x_t | d^t) &= \frac{f(\Delta_t | a_t, x_t) f(x_t | a_t, d^{t-1})}{\underbrace{\int f(\Delta_t | a_t, x_t) f(x_t | a_t, d^{t-1}) dx_t}_{f(\Delta_t | a_t, d^{t-1})}}. \end{aligned}$$

The filtering does not depend on the used informationally causal DM strategy $\{f(a_t | d^{t-1})\}_{t \in t^*}$ but on the generated actions only.

Proposition 2.2. (Solution of the FPD). The optimal DM strategy minimizing $D(f||If)$ is randomized one given by the pdfs

$$\begin{aligned} Of(a_t | d^{t-1}) &= If(a_t | d^{t-1}) \frac{\exp[-\omega(a_t, d^{t-1})]}{\gamma(d^{t-1})} \\ \gamma(d^{t-1}) &\equiv \int If(a_t | d^{t-1}) \exp[-\omega(a_t, d^{t-1})] da_t. \end{aligned}$$

Starting with $\gamma(d^T) \equiv 1$, functions $\omega(a_t, d^{t-1})$ are generated recursively for $t = T, T-1, \dots, 1$ in the backward manner, as follows

$$\begin{aligned} &\omega(u_t, d^{t-1}) \\ &\equiv \int \Omega(u_t, d^{t-1}, x_{t-1}) f(x_{t-1} | d^{t-1}) dx_{t-1} \end{aligned}$$

with $f(x_t | d^t)$ given by Proposition 2.1 and

$$\begin{aligned} &\Omega(a_t, d^{t-1}, x_{t-1}) \\ &\equiv \int f(\Delta_t | a_t, x_t) f(x_t | a_t, x_{t-1}) \times \\ &\ln \left(\frac{f(\Delta_t | a_t, x_t) f(x_t | a_t, x_{t-1})}{\gamma(d^t) If(\Delta_t | a_t, x_t) If(x_t | a_t, x_{t-1})} \right) d\Delta_t dx_t. \end{aligned}$$

The presented FPD has the following features:

- The minimizing DM strategy is found explicitly. This simplifies approximate solution of this version of dynamic programming.
- An equivalent FPD can be constructed to any Bayesian DM given by the pair \mathcal{U} and Z .
- *All DM elements* (the environment model, randomized strategies and their ideal counterparts) are described probabilistically.

3. COOPERATION

A participant acts, i.e., transforms the chosen DM elements into his optimal DM strategy and applies it. He acts irrespectively whether the set of his neighbors is empty or not. However, the explicit awareness of neighbors opens a way to a better performance via cooperation, i.e., via exchange and modification of the DM elements of neighbors.

The discussed methodology takes participants, labelled by p , as neighbors iff behaviors \mathcal{Q}_p of closed loops they consider overlap. Due to the limited abilities of a participant $p = \pi$, the *set of his neighbors* p_π^* has a few members. The methodology cares about the neighbors who are willing to exchange DM elements and possibly modify personal ones. The modification, which is discussed below, induces re-design of the optimal DM strategy and, hopefully, leads to an improved DM performance of neighboring participants.

“Pure” types of participants are:

Selfish participant accepts the DM elements offered by his neighbor, however, uses them irrespectively of DM goals of the neighbor.

Cooperative participant negotiates with his neighbor the degree of mutual influence, i.e., cooperating participants come to a common degree of acceptance of the offered DM elements.

Dominating participant enforces his DM elements to other(s). It creates a kind of, possibly softly, centralized DM. Applicability of this most efficient cooperation reaches soon robustness, communication and computation barriers.

A more extensive discussion of this classification can be found, in Kárný and Guy (2004).

4. COOPERATION TOOLS

Knowledge about behavior \mathcal{Q}_p of the closed loop formed by a participant p and its environment is expressed by the pdf $f(\mathcal{Q}_p|p)$. According to (4), it is product of the environment model with factors $f(\Delta_{t;p}|a_{t;p}, x_{t;p}, p)$, $f(x_{t;p}|a_{t;p}, x_{t-1;p}, p)$ and the DM strategy with a generic factor $f(a_{t;p}|d_p^{t-1}, p)$.

The ways of enriching this knowledge as well as modification of its ideal counterpart by using DM elements offered by neighbors are inspected here.

4.1 Knowledge offered by neighbors

Let us consider a fixed participant π with neighbors labelled by $p \in p_\pi^*$. The participant π delimits the group of his neighbors as those participants about whose behaviors he is aware of and who have a common part C_p with his behavior \mathcal{Q}_π . For a $p \in p_\pi^*$, let us decompose $\mathcal{Q}_\pi = (\bar{C}_p, C_p)$, where \bar{C}_p is available for the participant π but it is unavailable for the participant p . Similarly, $\mathcal{Q}_p = (\bar{C}_\pi, C_p)$, where \bar{C}_π is the part of the behavior \mathcal{Q}_p unavailable to the participant π .

The neighboring participant p offers his knowledge by providing the pdf $f(\mathcal{Q}_p|p) = f(\bar{C}_p, C_p|p)$. The accepting participant π is obviously able to exploit at most the marginal pdf $f_p \equiv f(C_p|p)$ characterizing the common part C_p of behaviors \mathcal{Q}_π and \mathcal{Q}_p . To establish a cooperation, the participant π needs a tool that allows him to use all offered pdfs $f_p \equiv f(C_p|p)$ and modify $f(\mathcal{Q}_\pi|\pi)$ by the set of marginal pdfs

$$f_\pi^* \equiv \{f_p \equiv f(C_p|p), \mathcal{Q}_\pi = (\bar{C}_p, C_p), p \in p_\pi^*\}. \quad (5)$$

It is useful to include $f_\pi \equiv f(\mathcal{Q}_\pi|\pi)$ into the set f_π^* . This is possible with $\mathcal{Q}_\pi \equiv C_\pi$ and void \bar{C}_π .

Knowledge exploitation in decentralized DM is always poorer than in the centralized DM. Thus, for fixed strategies of neighbors, there is a global closed-loop pdf ${}^G f \equiv {}^G f(\mathcal{Q}_\pi)$ describing completely knowledge of the considered group of neighbors. The respective pdfs in f_π^* (5) can be interpreted as noisy (not fully compatible) projections of the global pdf ${}^G f(\mathcal{Q}_\pi)$. The knowledge sharing can be interpreted as a correction of $f(\mathcal{Q}_\pi|\pi)$, a noisy projection of the global pdf ${}^G f(\mathcal{Q}_\pi)$ to the knowledge domain of the π th participant, by the pdfs in the set f_π^* . Basic ways are outlined below.

The inspected corrections need to extend the processed marginal pdfs $f(C_p|p)$ on \mathcal{Q}_π . It is done via a chain rule $\tilde{f}(\bar{C}_p|C_p)f(C_p|p)$ with a suitably chosen conditional pdf $\tilde{f}(\bar{C}_p|C_p)$.

4.2 Approximation-based sharing

The global, group describing, pdf ${}^G f = {}^G f(\mathcal{Q}_\pi)$ is unknown and we approximate it by the pdf denoted $f(\mathcal{Q}_\pi|f_\pi^*)$ and defined

$$\begin{aligned} f(\mathcal{Q}_\pi|f_\pi^*) &\in \text{Arg} \min_{\{\hat{f} \equiv f(\mathcal{Q}_\pi)\}} \mathbb{E} \left[\text{D} \left({}^G f \middle| \middle| \hat{f} \right) \right] \\ &= \text{Arg} \min_{\{\hat{f}\}} \mathbb{E} \left[\text{D} \left({}^G f \middle| \middle| \hat{f} \right) \middle| \hat{f}, f_\pi^* \right], \end{aligned} \quad (6)$$

where the expectation \mathbb{E} and its conditional version $\mathbb{E} \left[\cdot \middle| \hat{f}, f_\pi^* \right]$ are taken over the uncertain global pdfs ${}^G f \in {}^G f^* \equiv$ the set of possible objective global pdfs. The equality (6) follows from the basic lemma of DM, Astrom (1970), Kárný et al. (2005). It can be exploited if we manage to relate the optimized pdf \hat{f} and (meta)data f_π^* (5) to the unknown global pdf ${}^G f$. It is done heuristically.

First, we extend the respective marginal pdfs $f(C_p|p)$ on \mathcal{Q}_π using the appropriate conditional pdfs derived from the pdfs $\hat{f}(\mathcal{Q}_\pi)$ over which the optimization runs, i.e., the extensions have the form $\hat{f}(\bar{C}_p|C_p)f(C_p|p)$.

Let us split the set ${}^G f^*$ of possible global pdfs ${}^G f$ to non-overlapping subsets ${}^G f_p^*$, $p \in p_\pi^*$. Then,

$$\begin{aligned} &\mathbb{E} \left[\text{D} \left({}^G f \middle| \middle| \hat{f} \right) \middle| \hat{f}, f_\pi^* \right] \\ &= \mathbb{E} \left[\sum_{p \in p_\pi^*} \chi_p ({}^G f) \text{D} \left({}^G f \middle| \middle| \hat{f} \right) \middle| \hat{f}, f_\pi^* \right], \end{aligned} \quad (7)$$

where $\chi_p(\cdot)$ is indicator of the set ${}^G f_p^*$. Sets ${}^G f_p^*$ are selected so that on them the approximation

$$\begin{aligned} &\mathbb{E} \left[\text{D} \left({}^G f \middle| \middle| \hat{f} \right) \middle| \hat{f}, {}^G f_p^* \right] \\ &\approx \text{D} \left(\hat{f}(\bar{C}_p|C_p)f(C_p|p) \middle| \middle| \hat{f} \right) \\ &= \int f(C_p|p) \ln \left(\frac{f(C_p|p)}{\hat{f}(C_p)} \right) dC_p \end{aligned} \quad (8)$$

causes the smallest error when we minimize the functional resulting from this approximation instead of the functional (6). With the approximation (8), the minimization task (6) gets the form

$$\begin{aligned} f(\mathcal{Q}_\pi|f_\pi^*) &\in \text{Arg} \min_{\{\hat{f}\}} \\ &\sum_{p \in p_\pi^*} \alpha_p(\hat{f}) \int f(C_p|p) \ln \left(\frac{f(C_p|p)}{\hat{f}(C_p)} \right) dC_p \end{aligned} \quad (9)$$

where $\alpha_p = \alpha_p(\hat{f})$ is the probability of the set ${}^G f_p^*$. Let us neglect dependence of these probabilities on the optimized pdf $\hat{f} = \hat{f}(\mathcal{Q}_\pi)$. Then, it is easy to find that the pdf $f(\mathcal{Q}_\pi|f_\pi^*)$ optimally

approximating the unknown global pdf ${}^G f(\mathcal{Q}_\pi)$ in the sense (9) has to meet the necessary condition

$$\sum_{p \in p_\pi^*} \alpha_p \frac{f(C_p|p)}{f(C_p|f_\pi^*)} = \text{constant}. \quad (10)$$

Explicit solution of (10) is rarely available. For instance, if $\pi = 2$, $p_2^* \equiv \{1, 2\}$, then it reads

$$f(C_2|f_1, f_2) \equiv f(\mathcal{Q}_2|f_1, f_2) = f(\bar{C}_1|C_1, \pi = 2) \\ \times [\alpha_1 f(C_1|p = 1) + (1 - \alpha_1) f(C_1|p = 2)].$$

Generally, (10) has to be solved numerically. An efficient iterative procedure with a guaranteed convergence was proposed in Kracík (2005).

The approximation (8) is motivated by the fact that the particular models $f(C_p|p)$ are usually obtained by Bayesian learning and thus they have “tendency” to minimize the KLD to the global pdf ${}^G f(C_p)$, Kárný et al. (2005). This also motivates the choice of the factors $\hat{f}(\bar{C}_p|C_p)$ for extending respective marginal pdfs $f(C_p|p)$ to the pdf describing the behavior \mathcal{Q}_π .

Note that the participant’s type, Section 3, determines the selection of probabilistic weights α_p .

4.3 Alternative solution

The assumption that $\alpha_p, p \in p_\pi^*$, do not depend on the optimized \hat{f} can be rather rough. Thus, it is worthwhile to consider alternative approximations. For instance, the processed marginal pdfs can be extended to describe whole \mathcal{Q}_π^* by the conditional pdf of $f(\mathcal{Q}_\pi|\pi)$, i.e., $f(\bar{C}_p|C_p, \pi) f(C_p|p)$. This leads to the re-definition of sets ${}^G f_p^*$ so that

$$\mathbb{E} \left[D \left({}^G f \middle| \middle| \hat{f} \right) \middle| \hat{f}, {}^G f_p^* \right] \quad (11) \\ \approx D \left(f(\bar{C}_p|C_p, \pi) f(C_p|p) \middle| \middle| \hat{f}(\mathcal{Q}_p) \right)$$

causes the smallest error when we minimize the functional resulting from this approximation instead of the functional (6). In this case, it seems that the assumption that $\alpha_p, p \in p_\pi^*$, do not depend on the optimized argument is more realistic. Moreover, the optimal approximation can be found explicitly as the mixture of pdfs

$$f(\mathcal{Q}_\pi|f_\pi^*) = \sum_{p \in p_\pi^*} \alpha_p f(\bar{C}_p|C_p, \pi) f(C_p|p).$$

This is an, often recommended, arithmetic polling of pdfs. Note that solution of the same approximation task, under the same conditions but with the alternative version of the KLD, gives the geometric polling. It declines, however, from the well-justified version of the KLD, Bernardo (1979).

The participant’s type, Section 3, determines the way of selecting the probabilistic weights α_p .

4.4 Learning option

To simplify discussion in this paragraph, let us assume that the behavior \mathcal{Q}_π of the participant π is fully included in behaviors of his neighbors, i.e., $C_p = \mathcal{Q}_\pi, p \in p_\pi^*$. The pdfs $f_p \in f_\pi^*$ represent (meta)data, which should be processed by the Bayes rule in order to learn the unknown global pdf ${}^G f \equiv \Theta \equiv$ unknown parameter of the problem; model of the behavior \mathcal{Q}_π is parameterized by Θ . Thus, the participant π has at disposal $f(\mathcal{Q}_\pi|\Theta, \pi)$ and as Bayesian DM assigns prior pdf $f(\Theta|\pi)$ to it. For a finite dimensional Θ , the papers Kracík and Kárný (2005); Kárný et al. (2006) proposed the following extension of the Bayes rule that modifies the prior pdf $f(\Theta|\pi)$ by probabilistic information f_π^* about \mathcal{Q}_π

$$f(\Theta|f_p, \pi) \propto \underbrace{f(\Theta|\pi)}_{\text{prior pdf chosen by } \pi} \times \quad (12) \\ \exp \left[\sum_{p \in p_\pi^*} \nu_p \int \underbrace{f(\mathcal{Q}_\pi|p)}_{\text{exploited } f_p} \times \underbrace{\ln(f(\mathcal{Q}_\pi|\Theta, \pi))}_{\text{relates } \mathcal{Q}_\pi \text{ to } {}^G f = \Theta} d\mathcal{Q}_\pi \right],$$

where $\nu_p \geq 0, p \in p_\pi^*$, control the impact of the processed knowledge. The type of participant, Section 3, determines the way of their selection.

4.5 Harmonizing ideal pdfs

Harmonization concerns combination of the neighbors’ ideal pdfs. Formally, the same combination rules as in knowledge sharing could be used. However, the global pdf has to be interpreted as a group ideal and the learning option makes less sense in this case. Thus, the approximation-based options remain to be only candidates.

5. DISCUSSION

Sharing and harmonizing tools outlined in Section 4, seem to exhaust practically the range of meaningful alternatives. Naturally, the problem of combination of pdfs has been addressed repeatedly in various contexts. The fact that the combined marginal pdfs are incompatible is the major specificity of the addressed problem. This rules out sophisticated techniques based on copulas, Nelsen (1999), and their extensions. Techniques related to Bayesian networks, Jensen (2001); Jiroušek (2003), seem to be unsuitable for the same reason.

Concerning to sharing of knowledge, the composition based on learning is conceptually the

most satisfactory as it fits the basic Bayesian paradigm. Still, there are technical problems related to it. Specifically, the application of the Bayes rule assumes implicitly existence of the model $f(f_p|\Theta)$. Its correct definition requires employment of heavy machinery of stochastic processes. Definite results are unavailable.

Concerning to approximation-based solutions. The appropriate choice of weights is critical. While for knowledge sharing they can be taken as unknown parameters and estimated in Bayesian way, their choice for harmonization is unsolved: a formalized way of their negotiation is missing.

The outlined techniques face common problems. Specifically,

- The compound pdf is often out of an easy-to-handle class of its constituents. Thus, either subsequent projection to this class is needed or the optimization has to be performed within it.
- The operations are performed on marginal pdfs derived from the conditional pdfs creating environment model and description of the DM strategy. This can be quite demanding.
- Any change of a DM element induces changes of others. For instance, modification of models or ideal pdfs implies the need to re-design the DM strategy. This change should be reported to neighbors. Consequently, the sharing of knowledge and harmonizing the aims becomes an iterative process. Its organization and termination are generally non-trivial.
- Locality is the main good feature of the considered compositions. Global model or ideal for neighbors are formed either implicitly or at most temporarily: the result is projected back to respective participants. The cooperation can run independently and in parallel within non-overlapping groups of neighbors. Cooperation of overlapping groups influence each other, hopefully without bad consequences. No formal analysis is now available and various types of global behaviors may emerge.

6. EPILOGUE

Adaptive systems have been recognized as systems fighting with the complexity of DM problems by using local (in time and data spaces) approximations around behavior realizations, The methodology outlined in this paper is a natural and quite powerful tool extending capabilities of adaptive systems. Among others, it opens the promising way of combining highly sophisticated adaptive algorithms that otherwise can solve only small-scale problems.

REFERENCES

- K.J. Astrom. *Introduction to Stochastic Control*. Academic Press, New York, 1970.
- J.O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag, New York, 1985.
- J. M. Bernardo. Expected information as expected utility. *The Annals of Statistics*, 7(3): 686–690, 1979.
- D.P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Nashua, US, 2001. 2nd edition.
- F.V. Jensen. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York, 2001.
- R. Jiroušek. On experimental system for multidimensional model development MUDIN. *Neural Network World*, (5):513–520, 2003.
- M. Kárný. Adaptive systems: Local approximators? In *Preprints of the IFAC Workshop on Adaptive Systems in Control and Signal Processing*, pages 129–134, Glasgow, August 1998. IFAC.
- M. Kárný, J. Andryšek, A. Bodini, T. V. Guy, J. Kracík, and F. Ruggeri. How to exploit external model of data for parameter estimation? *International Journal of Adaptive Control and Signal Processing*, 20(1):41–50, 2006.
- 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.
- M. Kárný and T. V. Guy. Fully probabilistic control design. *Systems & Control Letters*, 55(4):259–265, 2006.
- M. Kárný and T.V. Guy. On dynamic decision-making scenarios with multiple participants. In J. Andryšek, M. Kárný, and J. Kracík, editors, *Multiple Participant Decision Making*, pages 17–28, Adelaide, May 2004. Advanced Knowledge International.
- J. Kracík. Composition of probability density functions based on minimization of Kullback-Leibler divergence. In *Proceedings of the 5th International PhD Workshop on Systems and Control - Young Generation Viewpoint*, pages 1–4, Budapest, September 2005. Hungarian Academy of Sciences.
- J. Kracík and M. Kárný. Merging of data knowledge in Bayesian estimation. In J. Filipe, J. A. Cetto, and J. L. Ferrier, editors, *Proceedings of the Second International Conference on Informatics in Control, Automation and Robotics*, pages 229–232, Barcelona, September 2005. INSTICC.
- S. Kullback and R. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–87, 1951.
- R.B. Nelsen. *An Introduction to Copulas*. Springer, New York, 1999.