Towards Information-Theoretic Decision Making in a *Conservative* Information Space

Vadim Indelman*

***Abstract*— We propose the conceptual idea of resorting to conservative information fusion techniques for information-theoretic decision making, aiming to address challenges involved with decision making over a high-dimensional, possibly highly-correlated, information space. Our key observation is that in certain cases, the impact of any two actions (or controls) on an appropriate utility measure, such as entropy, has the same trend regardless if using the original probability distribution function (pdf) or a conservative approximation of thereof. This observation suggests that in these cases, decision making can be performed over a conservative pdf, *instead* of the original pdf, *without* sacrificing performance. We introduce and prove this concept for the basic one-dimensional case assuming Gaussian probability distributions, and then consider its extension to a high-dimensional state space. In particular, we consider a specific conservative pdf that *decouples* the random variables in the joint pdf, admitting extremely efficient entropy computation. We then present our progress in identifying classes of problems in which information-theoretic decision making over this conservative and original pdfs produce identical results. The concept is illustrated in the context of choosing informative image observations in an aerial visual simultaneous localization and mapping scenario.**

I. INTRODUCTION

Intelligent decision-making under uncertainty is essential in numerous problem domains, including environment modeling [14], sensor selection [7], active vision [2], and active simultaneous localization and mapping (SLAM) [16]. Information-theoretic decision-making is a common approach to address these type of problems, where the objective is to find an action (or sequence of actions) that maximizes a utility measure quantifying the expected information gain.

In many cases, decision making should be performed over a high-dimensional information space that captures spatial and, perhaps also temporal relations between the random variables. For example, in active SLAM, the state represents the robot trajectories and the observed environment, while in the context of environment modeling it represents a multidimensional field which can be dynamically changing over time. Decision making in these cases is computationally expensive, as it involves evaluating the expected information gain (for different candidate actions) over a high-dimensional information space.

Furthermore, it is often required to perform information-theoretic decision making over *multiple* look-ahead steps, i.e. looking for the best action sequence. Examples of such problems include different instances of planning under uncertainty. Due to the inherent process and measurement noise, these problems are instances of stochastic control and are therefore often framed within a partially observable Markov decision process (POMDP) framework [9], which is defined over the space of probability distributions of the state space, also referred to as the belief space or information space [13]. The POMDP framework is a powerful mathematical tool, naturally representing the inherent coupling between action and information, however, unfortunately it is known to be computationally intractable [11]. Over the past decade numerous sub-optimal approaches have been developed that trade-off computational complexity with performance (see, e.g., [12], [13], [4]).

We note that another related problem is sensor selection, where one has to choose a subset of k most informative sensors. While this problem has been shown to be NP-hard [6], also here, sub-optimal approaches have been developed in recent years. These include approaches that perform greedy selection by calculating mutual information between the state and measurements or sensors (e.g. [2]), or by reformulating the problem as convex optimization [7]. However, in both cases, operation in a high-dimensional information space is computationally challenging.

In this paper we introduce a new paradigm, aiming to address the above mentioned challenges. The basic idea is to resort to conservative information fusion techniques [8] for information-theoretic decision making. These techniques have been developed to fuse different, possibly correlated, sources of information when the correlation is unknown and cannot be neglected, producing a consistent state estimation.

The motivation to use such a concept is twofold. First, a conservative approximation of the original probability distribution function (pdf) can often be more easily calculated than the original pdf. Second, this concept can drastically reduce computational complexity, as it is possible to construct extremely sparse conservative pdfs.

Our key observation is that information-theoretic decision making over a *conservative* approximation of the original pdf yields, in certain cases, the *same result* as would be obtained by using the original pdf. More precisely, in these scenarios, the impact of any two candidate actions on entropy (or other related utility measure) of the corresponding a posteriori pdfs (beliefs) has the same trend *regardless* if it is calculated based on the original or the conservative pdf.

This paper introduces and analyzes the fundamental aspects of this novel paradigm. We first (Section III) consider a one-dimensional state and prove, assuming Gaussian distributions, the proposed concept does not sacrifice performance.

We then make the first steps towards extending this concept to high-dimensional state spaces, considering the problem of a single look-ahead step information-theoretic

*Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel.

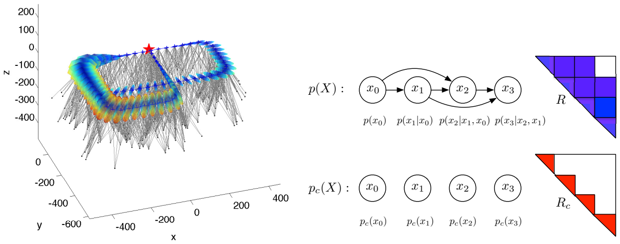

(a) (b)

Fig. 1: (a) Aerial visual SLAM. The figure shows a snapshot of robot trajectory and observed 3D points. Lines denote monocular camera observations of 3D points. Ellipsoids represent *incremental* robot pose covariances. The mark '$\star$' indicates initial robot position. (b) Bayes net and corresponding sparsity structure of square root information matrix for $p\left(X\right)$ and $p_c\left(X\right)$.

decision making (Section IV). In particular, we consider a specific conservative approximation that *decouples* the random variables in the high-dimensional state space (see, e.g., [10]). For Gaussian distributions, this conservative approximation corresponds to a *diagonal* information matrix, admitting extremely efficient entropy calculation. We present our current progress on identifying classes of problems in which information-theoretic decision making over this decoupled conservative pdf and the original pdf produce identical results. We illustrate the proposed concept in the context of informative measurement selection in aerial visual SLAM problem (see Figure 1a).

II. Problem Formulation

Consider a state x representing, for example, camera pose, navigation state, or the coordinates of a tracked target. The probability distribution function (pdf) over x_k at time t_k is given by

$$p\left(x_k|z_{0:k},u_{0:k-1}\right), \tag{1}$$

where $z_{0:k}$ and $u_{0:k-1}$ denote, respectively, all the acquired observations and performed actions (controls) by time t_k.

Given a control u_k and observation(s) z_{k+1} captured at the next time t_{k+1}, the posterior pdf over x_{k+1} is given by

$$p\left(x_{k+1}|z_{0:k+1},u_{0:k}\right)=\eta p\left(z_k|x_k\right)\cdot\int p\left(x_k|z_{0:k},u_{0:k-1}\right)\cdot p\left(x_{k+1}|x_k,u_k\right)dx_k, \tag{2}$$

where η is a normalization constant, and $p\left(x_{k+1}|x_k,u_k\right)$ and $p\left(z_k|x_k\right)$ are probabilistic motion and observation models.

When the observation model involves additional random variables, as common in vision-aided navigation and SLAM, (some of the) past states are not marginalized out in Eq. (2) and are kept estimated instead [4], [5].

Now, we consider at time t_k, the impact of some action u_k on the posterior (2) of a future state x_{k+1}. Limiting the discussion to information-theoretic decision making we define the following single look-ahead step objective function:

$$J\left(u_k\right)=\mathbb{E}_{z_{k+1}}\left[\mathcal{H}\left(p\left(x_{k+1}|z_{0:k+1},u_{0:k}\right)\right)\right], \tag{3}$$

where we use the entropy $\mathcal{H}$ as utility measure:

$$\mathcal{H}\left(p\left(x\right)\right)=-\mathbb{E}\left[\log p\left(x\right)\right]=-\int p\left(x\right)\log p\left(x\right)dx. \tag{4}$$

The expectation operator in Eq. (3) accounts for the fact that future observations z_{k+1} are unknown at the current time t_k.

In this paper we consider observation models with Gaussian noise, i.e.

$$z_i=h_i\left(x_i\right)+v_i\quad,\quad v_i\sim\mathcal{N}\left(0,\Sigma_{v_i}\right), \tag{5}$$

and a deterministic control, i.e. x_{k+1} is predicted with high accuracy given x_k and control u_k.

Since the posterior $p\left(x_{k+1}|z_{0:k+1},u_{0:k}\right)$ is Gaussian, the entropy $\mathcal{H}$ is proportional to the a posteriori information (covariance) matrix I^+:

$$\mathcal{H}\left(p\left(x_{k+1}|z_{0:k+1},u_{0:k}\right)\right)=-\frac{1}{2}\log\left[\left(2\pi e\right)^n\left|I^+_{k+1}\right|\right], \tag{6}$$

with $I^+_{k+1}=I_k+H^T\Sigma_v^{-1}H$, where I_k is the information matrix from the previous step and H is the measurement Jacobian, linearized about the current estimate of x. The above can be considered as a single iteration of Gauss-Netwon optimization (or equivalently, extended Kalman filter) and we assume it sufficiently captures the impact of action u_k on entropy [4]. Therefore, I^+_{k+1} is not a function of the unknown future observations z_{k+1}, and the objective function (3) simply becomes the entropy (6).

Having presented in detail the probability terms, we switch to concise notations. We introduce the operator $\zeta\left(.,.\right)$ to represent the pdf evolution between two consecutive time instances and compactly rewrite Eq. (2) as:

$$p\left(x_{k+1}|z_{0:k+1},u_{0:k}\right)\doteq\zeta\left(p\left(x_k|z_{0:k},u_{0:k-1}\right),u_k\right). \tag{7}$$

We also will omit from now on, the explicit conditioning on the given (and known) observations and controls, as well as the time instant t_k:

$$p\left(x\right)\doteq p\left(x_k|z_{0:k},u_{0:k-1}\right) \tag{8}$$

and write the pdf at the next time step, given control $u\doteq u_k$, as $\zeta\left(p\left(x\right),u\right)$. Furthermore, with a slight abuse of notation, we denote by $u_i\in\mathcal{U}$ the ith control, with $\mathcal{U}$ representing a discrete set of possible controls.

The fundamental problem addressed in this paper is information-theoretic decision-making, defined as identifying the most informative action $u^\star$ from a discrete set of controls $\mathcal{U}$, i.e. minimizing the entropy (6).

In the remainder of this paper we develop the concept of decision-making over a *conservative* pdf instead of the original pdf $p\left(x\right)$, beginning the discussion with a one-dimensional state $x\in\mathbb{R}$, and considering its extension to high-dimensional state spaces.

III. Decision Making Over a Conservative 1D Information Space

We denote by $p_c(x)$ a *conservative approximation* of some pdf $p(x)$, satisfying the two sufficient conditions [1]:

1) $\mathcal{H}(p(x)) \leq \mathcal{H}(p_c(x))$, and
2) Order preservation: $\forall x_i, x_j$

$$p_c(x = x_i) \leq p_c(x = x_j) \;\text{ iff }\; p(x = x_i) \leq p(x = x_j)\,.$$

The concept of decision-making in the conservative information (belief) space is formulated in the following theorem considering a one-dimensional state x and Gaussian probability distributions.

Theorem 1: For any two actions $u_1, u_2 \in \mathcal{U}$:

$$\mathcal{H}(\zeta(p(x), u_1)) \leq \mathcal{H}(\zeta(p(x), u_2))$$

if and only if

$$\mathcal{H}(\zeta(p_c(x), u_1)) \leq \mathcal{H}(\zeta(p_c(x), u_2))\,.$$

In words, the impact of any two candidate actions on the entropy of the corresponding future pdfs has the *same trend* regardless if it is calculated based on the original pdf $p(x)$, or based on the conservative pdf $p_c(x)$. Therefore, choosing the most informative action can be done by considering the conservative pdf $p_c(x)$ *instead* of the original pdf $p(x)$.

Proof: Recall the entropy definition for the Gaussian case (6). Since entropy $\mathcal{H}$ is monotonic in $|I^+|$, Theorem 1 is equivalent to

$$\left|I_1^+\right| \leq \left|I_2^+\right| \;\text{ iff }\; \left|I_{c1}^+\right| \leq \left|I_{c2}^+\right|, \tag{9}$$

where I_i^+ denotes the a posteriori information matrix due to action u_i, representing the second moment of the future pdf $\zeta(b, u_i)$, and I_c^+ refers to the same quantity for $\zeta(b_c, u_i)$.

In general, for any measurement z with observation model (5) and Jacobian $A = \frac{\partial h}{\partial x}$, the following holds

$$I^+ = I + A^T \Sigma_v^{-1} A\;,\; I_c^+ = I_c + A^T \Sigma_v^{-1} A. \tag{10}$$

Since x is a scalar random variable: $|I| \equiv I$.

Therefore, if $\left|I_1^+\right| \leq \left|I_2^+\right|$ then $A_1^T \Sigma_{v1}^{-1} A_1 \leq A_2^T \Sigma_{v2}^{-1} A_2$. Hence, $\left|I_{c1}^+\right| = I_c + A_1^T \Sigma_{v1}^{-1} A_1 \leq I_c + A_2^T \Sigma_{v2}^{-1} A_2 = \left|I_{c2}^+\right|$. The opposite direction can be proved similarly. ■

Figure 2 illustrates Theorem 1 in a basic example, considering two measurement models with varying accuracy: $\Sigma_v = 0.5^2$ and 0.2^2. The figure shows the a priori pdfs $p(x)$ and $p_c(x)$, and the a posteriori pdfs $\zeta(p(x), u)$ and $\zeta(p_c(x), u)$, choosing each time a different measurement model. In this simple case, the control u indicates which measurement model to use and therefore these a posteriori pdfs merely correspond to the pdfs $p(x|z)$ and $p_c(x|z)$, which incorporate the chosen measurement model $p(z|x)$. As seen from the figure, entropy of $p(x|z)$ is smaller for a more accurate measurement model, as expected. However, the figure shows this is also the case for $p_c(x|z)$, as stated by Theorem 1.

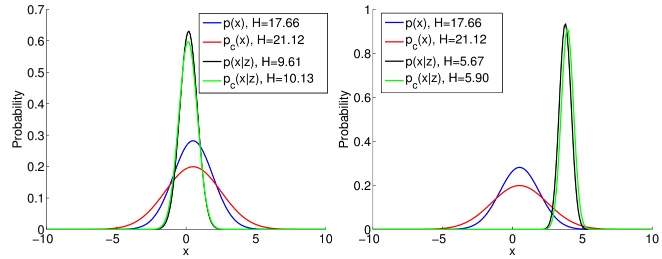

(a) (b)

Fig. 2: Example for a 1D Gaussian distribution: the figures show $p(x)$ and $p_c(x)$ and the a posteriori pdfs $p(x|z)$ and $p_c(x|z)$ for two measurement covariances (a) $\Sigma_v = 0.5^2$ and (b) $\Sigma_v = 0.2^2$. As expected, entropy is smaller for the latter. However, according to Theorem 1, this can be concluded directly from $p_c(x|z)$. Entropy values for each case are indicated in the label.

IV. Extension to High-Dimensional State Space

Many real-world problems involve a high-dimensional state space that captures spatial and, perhaps also temporal relations between random variables. For example, in SLAM, the state represents robot poses and observed environment, while in the context of environment modeling it can represent a multidimensional, possibly dynamically changing, field. Decision making in these cases is computationally expensive, as it involves evaluating entropy (or another related measure), and in particular the determinant $\left|I_i^+\right|$, for each candidate action $u_i \in \mathcal{U}$. Calculation of $\left|I_i^+\right|$ is computationally expensive due to correlations between the random variables in X, which correspond to off-diagonal terms in I^+.

In this section we make progress towards extending the proposed concept to a high-dimensional state space. First, we define the notion of a *decoupled* conservative pdf, in which all the random variables are statistically independent of each other. Decision making over this specific conservative pdf can therefore be performed very efficiently. Next, we examine if it incurs any sacrifice in performance, by analyzing whether Theorem 1, in conjunction with the decoupled conservative pdf, is valid in different scenarios of interest.

A. Decoupled Conservative PDF

While so far the exposition referred to *some* conservative approximation, here we consider the following specific approximation of the high-dimensional state X (see, e.g. [1]):

$$p_c(X) \doteq \eta \prod_i p^{w_i}(x_i)\,, \tag{11}$$

with x_i being the ith component in X, η a normalization constant, and w_i are weights such that $\sum_i w_i = 1$.

While the states in the joint pdf $p(X)$ can be *highly correlated*, the conservative pdf $p_c(X)$ essentially *decouples* these correlations. Therefore, from a computational point of view, the pdf $p_c(X)$ can be considered as *sparsifying* the original pdf $p(X)$.

For Gaussian distributions, this is equivalent to scaling the covariance of each component by a factor of $1/w$ [10], and can be considered as a special instance of covariance intersection [8], [15]. As a result, the (square root) information matrix and the covariance become *diagonal* matrices. Figure 3a illustrates this process for the two-dimensional case, considering different values of $w \equiv w_i$.

For example, considering some $p(X)$ with $X \in \mathbb{R}^{4\times 1}$ and factorization:

$$p(X) = p(x_0)\, p(x_1|x_0)\, p(x_2|x_1, x_0)\, p(x_3|x_2, x_1), \quad (12)$$

the conservative pdf $p_c(X)$ can be trivially written as $p_c(X) = p_c(x_0)\, p_c(x_1)\, p_c(x_2)\, p_c(x_3)$. Figure 1b illustrates the corresponding Bayes net and sparsity pattern of square root information matrices of $p(X)$ and $p_c(X)$.

Information-theoretic decision-making, i.e. entropy calculation for each candidate action, over this decoupled, completely sparse, conservative pdf $p_c(X)$ can be performed extremely efficiently even for large state spaces, since it *does not require accounting for any correlations* between the states in X. The main question is whether such a process will produce the same result as would be obtained with the original pdf $p(X)$.

In the next section we address this crucial question, analyzing whether Theorem 1 is valid, in conjunction with the decoupled conservative pdf (11). If this is indeed the case, decision making over the decoupled pdf (11) will not incur any sacrifice in performance, while greatly reducing computational complexity.

As will be seen next, we prove validity of this statement for specific classes of scenarios and conduct numerical study suggesting validity in more general cases. However, this preliminary numerical study also suggests Theorem 1, in conjunction with the decoupled pdf (11), is *not* always valid.

Remark: Observe the difference with conservative information fusion formulation, that can be written for any two pdfs $p_a(X)$ and $p_b(X)$ as $p_c(X) = \eta p_a^w(X)\, p_b^{1-w}(X)$, where η is a normalization constant [1]. The formulation considered herein differs in two respects: (a) it calculates a conservative approximation of one of the pdfs via Eq. (11), while leaving the other unchanged; (b) in the context of decision making, only the *trend* in entropy (or other measure) corresponding to a posteriori distributions for different pdfs $p_b(X)$ is of interest.

B. Decision-Making Over a Decoupled Conservative PDF

For a high-dimensional state X, the determinant of the information matrix $I = R^T R$ can be written as

$$|I| = \prod_i r_{ii}^2, \quad (13)$$

where R is the square root information matrix, and r_{ii} is its ith eigenvalue. An alternative formulation of Theorem 1 is therefore:

$$\prod_i \left|r_{1,ii}^+\right| \leq \prod_i \left|r_{2,ii}^+\right| \text{ iff } \prod_i \left|r_{1,cii}^+\right| \leq \prod_i \left|r_{2,cii}^+\right|, \quad (14)$$

stating the product of eigenvalues changes *consistently* between the original and conservative pdfs. Here, $r_{1,ii}^+$ represents the ith eigenvalue in R_1^+, obtained by fusing I and the appropriate observation model(s) (5) corresponding to measurement(s) that will be acquired after performing action u_1, and $r_{1,cii}^+$ is defined similarly for I_c.

We proceed by first analyzing sufficient conditions for satisfying Eq. (14), proving Theorem 1 is valid for certain family of systems, and then consider more general cases.

1) Sufficient Conditions: It is straightforward to show that Eq. (14) is satisfied when *each* of the eigenvalues changes consistently:

$$\forall i: \ \left|r_{1,ii}^+\right| \leq \left|r_{2,ii}^+\right| \text{ iff } \left|r_{1,cii}^+\right| \leq \left|r_{2,cii}^+\right|. \quad (15)$$

Recalling Eq. (13), these sufficient conditions lead to Eq. (9), and therefore Theorem 1 is satisfied.

In particular, it follows from Eq. (15) that the proposed concept is applicable to problems with state-dependent observation noise:

$$z_i = h(X') + v_i, \quad (16)$$

where $v_i \sim \mathcal{N}\left(0, \alpha_i^2 \Sigma_v\right)$, $\alpha_i \in \mathbb{R}$ some scalar, and $X' \subseteq X$ is the same (arbitrary) subset of states for all z_i.

To see that, note that any two such observation models (16) have the same Jacobian $A_1 = A_2 \equiv A \doteq \frac{\partial h}{\partial x}$. Thus, $I_1^+ = I + \alpha_1^2 A^T A$ and $I_2^+ = I + \alpha_2^2 A^T A$, and similarly for I_c as well. While α_1 and α_2 are arbitrary, $\alpha_1^2 A^T A \leq \alpha_2^2 A^T A$ or $\alpha_1^2 A^T A \geq \alpha_2^2 A^T A$, and therefore either $I_1^+ \leq I_2^+$ and $I_{c1}^+ \leq I_{c2}^+$, or the same with the opposite sign. Therefore, all cases satisfy the sufficient conditions (15). Hence, decision-making can be performed over the decoupled conservative pdf *without* any sacrifice in performance, i.e. the same decisions will be obtained regardless if using the original or the conservative pdfs.

Basic example: We illustrate this concept in a basic example (Figures 3b-3c) that involves 2 scalar states, i.e. $X \in \mathbb{R}^{2\times 1}$. Two observation models that differ in accuracy are considered. In both cases we use the Jacobian $A = \frac{\partial h}{\partial X} = \begin{bmatrix} 1 & -0.1 \end{bmatrix}$, such that the update will be applied mainly on the first component of the state; measurement covariance models $\alpha_i^2 \Sigma_v$ are 0.5^2 and 0.2^2 and the actual measurements are $z_1 = 0.3$ and $z_2 = 1.3$, respectively. In both cases, the weights are $w_i = w = 0.5$, see Eq. (11), and the calculated Σ_c is shown in Figures 3b-3c. These figures also show the a posteriori uncertainty ellipses representing Σ^+ and Σ_c^+ for each of the two measurements, as well as the entropy. As expected, since the second measurement model is more accurate, it yields smaller entropy: $\mathcal{H}_1^+ = 3.3 > \mathcal{H}_2^+ = 2.4$. However, as stated by Theorem 1, the same conclusion can be reached by looking at the entropies obtained from the conservative pdf, i.e. $\mathcal{H}_{c1}^+ = 4.1 > \mathcal{H}_{c2}^+ = 3.2$.

Two remarks are in order at this point. First, observe that while state elements in the conservative covariance Σ_c are not correlated (it is a diagonal matrix), this is *not* the case in the a posteriori covariance Σ_c^+ due to introduced correlation terms (in this case between x_1 and x_2). Second, note that the a posteriori estimates $\hat{X}^+$ and $\hat{X}_c^+$ are not identical (see, e.g., Figure 3b). These observations raise the question whether the proposed concept can be extended to decision-making with *multiple* look-ahead steps; we intend to address this question in future research.

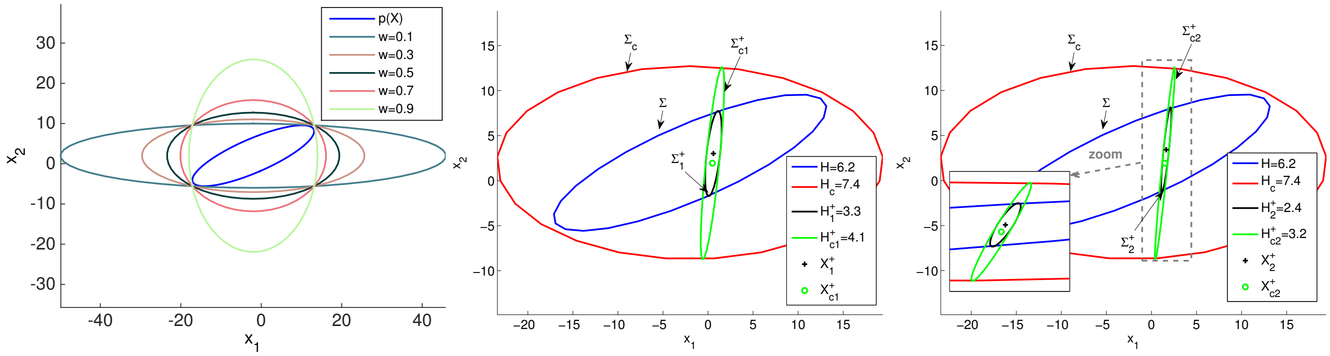

(a) (b) (c)

Fig. 3: (a) Covariances Σ and Σ_c representing, respectively, the original and conservative pdfs $p\left(X\right)$ and $p_c\left(X\right)$ for the 2D case. Different values of w are shown; in all cases the elements in X are decoupled (no correlation between x_1 and x_2), while $\mathbb{E}\left[x_1 x_2^T\right]$ is nonzero in Σ. (b)-(c): Demonstration of Theorem 1 in a 2D example for two different measurements models. A posteriori covariances and state estimates are shown for original and conservative pdfs. Entropy values are shown in legend. One can see that, in this case, $H_1^+ > H_2^+$ but also $H_{c1}^+ > H_{c2}^+$, as claimed by Theorem 1. Observe that $\hat{X}^+ \neq \hat{X}_c^+$.

2) More General Observation Models: While a formal in-depth analysis is the subject of ongoing work, in this section we provide numerical evidence suggesting the constraints (14), in conjunction with the decoupled conservative information space (11), are valid in certain additional cases as well. These include decision making involving: (a) different unary measurement models on arbitrary random variables in the state X, and (b) different binary measurement models all involving the same uncorrelated random variable.

Unary Measurement Model on Arbitrary Variable: We consider observation models

$$z_i = h_i\left(x_i\right) + v_i, \tag{17}$$

where the ith model involves an arbitrary single random variable $x_i \in X$, some observation function h_i and measurement noise $v_i \sim N\left(0, \Sigma_{vi}\right)$. Different observation models (as parametrized by i) can involve different x_i, h_i and Σ_{vi}. Such a scenario can represent, for example, sensing an environmental field at different locations, observing different 3D points, different onboard sensors or just GPS measurements obtained at different locations.

In the Appendix, we prove Theorem 1 is valid for measurement models (17), considering a two-dimensional case. Here, we demonstrate numerically it is valid also for higher dimensions, considering a toy example involving a four dimensional state $X \in \mathbb{R}^{4\times 1}$. Both $x_i \in X$ and the numerical values of Jacobian $\frac{\partial h_i}{\partial x_i}$ are randomly drawn for each i. Typical results are shown in Figure 4: The evolution of the four eigenvalues, corresponding to R^+ for each candidate action (which represents, in this case, measurement model selection), is shown in Figure 4a. Observe that sufficient conditions (15) are *not* satisfied because of the second and third eigenvalue, respectively, that does not change consistently. Nevertheless, as shown in Figure 4b, the products of eigenvalues, $det\left(I^+\right)$ and $det\left(I_c^+\right)$, do change consistently, demonstrating Theorem 1 (see Eq. (9)). A similar behavior has been observed in numerous additional runs.

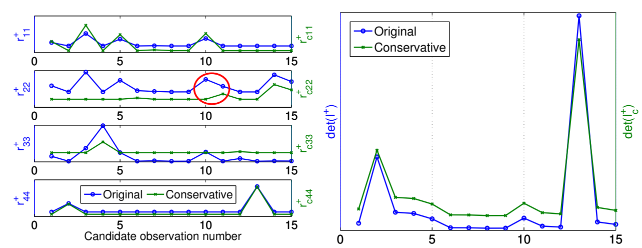

(a) (b)

Fig. 4: Observation model $p\left(z|x_i\right)$ involving an arbitrary-chosen single random variable $x_i \in X$ for each candidate measurement (see text for details). (a) evolution of eigenvalues of R^+ and R_c^+ for each candidate observation model; (b) $det\left(I^+\right)$ and $det\left(I_c^+\right)$ representing the information-theoretic objective. The two are consistent with each other, demonstrating Theorem 1 (i.e. satisfying Eq. (14)), despite one of the eigenvalues being inconsistent (red ellipse).

Binary Measurement Models with Uncorrelated State: We extend the model (17) to a binary observation model involving some variable(s) $x_i \in X$ and a *specific* variable $x \in X$ that is not correlated with other variables in X:

$$z_i = h_i\left(x, x_i\right) + v_i. \tag{18}$$

This restriction on x is required as we have observed that without it, the constraints (14) in conjunction with the decoupled pdf (11), are *not* always satisfied. We hope a formal analysis will provide insights as to whether it is possible to remedy this deficiency.

Here, we demonstrate the concept of decision-making in a conservative pdf using the model (18) while satisfying the above mentioned condition. We consider the problem of choosing most informative measurements within visual SLAM framework, where the state X comprises camera poses and observed 3D points.

The state x, corresponding to the pose of a new camera, remains uncorrelated as long as no image observations have been incorporated. Note that a prior $p(x)$ on x can still be added without introducing correlation to other states.

Given a new image, each image observation z_i involves the new camera pose x and some 3D point l_i. In this case, the observation model is given by $z_i = \pi(x, l_i) + v$, with the operator $\pi(.,.)$ representing the standard projection equations for a pinhole camera [3]. One can verify this observation model is of the same form as Eq. (18). For simplicity, we assume the camera is calibrated. Thus, the Jacobian A is given by

$$A_i \doteq \frac{\partial \pi}{\partial X} = \left[\begin{array}{cccccccc} 0 & \dots & \frac{\partial \pi}{\partial l_i} & 0 & \dots & 0 & \frac{\partial \pi}{\partial x} \end{array} \right], \quad (19)$$

where we arbitrarily assumed state X is ordered with 3D points appearing first, followed by robot poses.

Choosing the most informative observation involves calculating mutual information (MI) between the state and all the available observations [2], which requires entropy calculation over the a posteriori distribution (6). As we demonstrate next, doing so considering the *conservative* information space, instead of the original one, yields exactly the same trend, as stated by Theorem 1.

The considered scenario is shown in Figure 1a: an aerial robot travels starting from the red '$\star$' mark, while observing 3D points on the ground. The figure shows pose uncertainty that increases over time, until a loop closure (re-observation of 3D points) resets the uncertainty to prior values. Observations of 3D points are shown by gray lines; 3D points are denoted by black dots.

Figures 5a and 5b show the corresponding sparsity patterns of the information matrix I and of the conservative matrix I_c at a certain time instant. The latter is calculated via Eq. (11) and is therefore diagonal. Because I_c is diagonal, updating it with a Jacobean A_i is extremely efficient, with only 4 (block) entries involved in the calculations. Observe the number of non-zero elements: around 411k versus 2.4k elements.

Finally, Figure 5c shows the a posteriori entropy, comparing between $\log\left[\prod_i \left(r_{ii}^{+}\right)^2\right]$ with $\log\left[\prod_i \left(r_{cii}^{+}\right)^2\right]$ for all candidate observations $z_{k,j}$ in the considered two time instances. As can be seen, in both cases, using a conservative information space yields *exactly* the same trend as the original one, as stated by Theorem 1.

V. Conclusions

We introduced the conceptual idea of resorting to conservative information fusion techniques for information-theoretic decision making, aiming to address challenges involved with decision making over a high-dimensional, possibly highly-correlated, information space. This paradigm is driven by a *key observation* that, often, the impact of any two candidate actions on entropy (or another measure) has the *same trend* regardless if calculated over the original or conservative information space, suggesting decision making can be performed, in these cases, over the latter without any sacrifice in performance.

We first proved this statement *always* holds in a basic scenario of one-dimensional state and Gaussian noise probability distributions. We then focused on a high-dimensional state space, and considered a conservative probability distribution function (pdf) that *decouples* the random variables regardless of the actual correlations in the original joint pdf. This specific conservative pdf allowed extremely efficient entropy calculation even for a high-dimensional information space. Importantly, our analysis indicated decision-making over this conservative pdf *does not sacrifice performance* for several important classes of problem formulations, including informative planning, active sensing and sensor selection. We demonstrated applicability of the concept in the context of choosing informative image observations in visual simultaneous localization and mapping scenario. Future work includes further analysis and extension to multiple look-ahead steps and continuous control space.

Acknowledgments

This work was partially supported by the Technion Autonomous Systems Program (TASP).

Appendix

In this appendix we prove Theorem 1 for a unary measurement model on arbitrary variables (17), $z_i = h_i(x_i) + v_i$, considering a *two-dimensional* state $X \in \mathbb{R}^2$. Extending the proof to higher dimensions is the subject of ongoing work.

We assume the following general square root information matrices representing the original and the conservative pdfs:

$$R = \left[\begin{array}{cc} r_{11} & r_{12} \\ & r_{22} \end{array} \right] \in \mathbb{R}^{2\times 2}\ ,\ R_c = \left[\begin{array}{cc} r_{c11} & \\ & r_{c22} \end{array} \right] \in \mathbb{R}^{2\times 2},$$

Consider some two arbitrary functions h_1 and h_2, and without loss of generality assume the former is a function of x_1 while the latter is a function of x_2. We linearize these two functions about the current estimate $\hat{X}$ (which is the same for the original and the conservative case, $\hat{X} \equiv \hat{X}_c$). The resulting Jacobians, denoted respectively, by A and B are of the following general form:

$$A \doteq \frac{\partial h_1}{\partial X} = \left[\begin{array}{cc} \tilde{a}_1 & 0 \end{array} \right] \quad , \quad B \doteq \frac{\partial h_2}{\partial X} = \left[\begin{array}{cc} 0 & \tilde{b}_2 \end{array} \right], \quad (20)$$

where $\tilde{a}_1, \tilde{b}_2 \in \mathbb{R}$ are arbitrary scalars.

To prove Theorem 1 we need to show Eq. (14)

$$\prod_i \left|r_{1,ii}^{+}\right| \leq \prod_i \left|r_{2,ii}^{+}\right| \ \text{ iff } \ \prod_i \left|r_{1,cii}^{+}\right| \leq \prod_i \left|r_{2,cii}^{+}\right|, \quad (21)$$

is satisfied. Thus, first the a posteriori square root information matrices R^{+} and R_c^{+} need to be calculated for each of the two measurement models.

We start with R^{+}. In general, the a posteriori distribution, considering the ith measurement model, is given by $p_i(X|z) \propto p(X)\, p_i(z|x_i)$. It is not difficult to show that the maximum a posteriori (MAP) estimate corresponds to

$$\hat{X}^{+} = \arg\min_X \left\| \left(\begin{array}{c} R \\ \Sigma_{vi}^{-1/2} \frac{\partial h_i}{\partial X} \end{array} \right) \Delta X + \left(\begin{array}{c} 0 \\ \times \end{array} \right) \right\|^2, \quad (22)$$

where $\times$ indicates a non-zero entry. The matrix R^{+} can be calculated by performing an incremental factorization, nullifying each entry in $\Sigma_{vi}^{-1/2} \frac{\partial h_i}{\partial X}$.

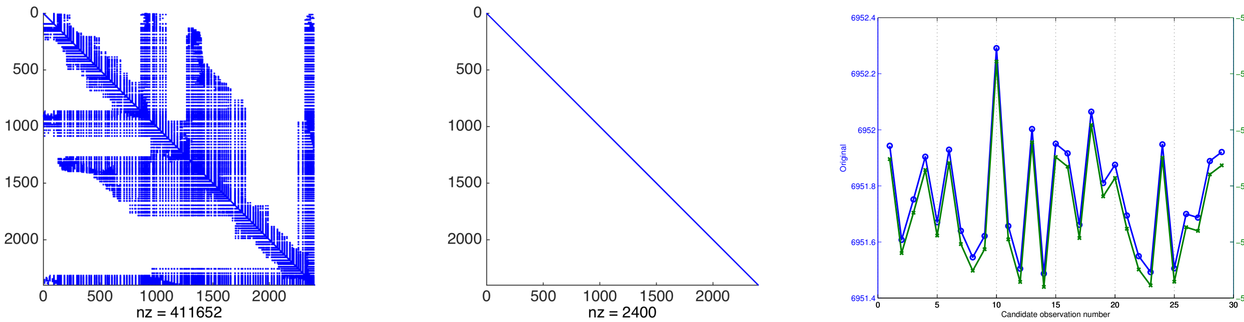

(a) (b) (c)

Fig. 5: Concept demonstration in choosing most informative image observations in aerial visual SLAM problem shown in Figure 1a. Figures (a) and (b) show sparsity patterns of the information matrix I and the conservative information matrix I_c, respectively; Figure (d) shows the a posteriori entropy calculated for each candidate image observation, using the original and the conservative information matrices. One can observe the trend is identical.

Specifically, for the two measurement models we have

$$\begin{bmatrix} r_{11} & r_{12} \\ & r_{22} \\ a_1 & 0 \end{bmatrix}, \quad \begin{bmatrix} r_{11} & r_{12} \\ & r_{22} \\ 0 & b_2 \end{bmatrix}, \tag{23}$$

where $a_1 \doteq \Sigma_{v1}^{-1/2}\tilde{a}_1$ and $b_2 \doteq \Sigma_{v2}^{-1/2}\tilde{b}_2$.

Applying a series of Givens rotations results in

$$\begin{pmatrix} R^{a+} \\ 0 \end{pmatrix} = \begin{bmatrix} r_{11}^{a+} & \times \\ & r_{22}^{a+} \\ 0 & 0 \end{bmatrix}, \begin{pmatrix} R^{b+} \\ 0 \end{pmatrix} = \begin{bmatrix} r_{11} & \times \\ & r_{22}^{b+} \\ 0 & 0 \end{bmatrix},$$

where $r_{11}^{a+} = \sqrt{r_{11}^2 + a_1^2}$, $r_{22}^{b+} = \sqrt{r_{22}^2 + b_2^2}$ and $r_{22}^{a+} = \sqrt{\frac{r_{22}^2\left(r_{11}^2+a_1^2\right)+\left(a_1 r_{12}\right)^2}{r_{11}^2+a_1^2}}$. Similarly, it is possible to show the conservative a posteriori square root information matrix assumes the following form:

$$\begin{pmatrix} R_c^{a+} \\ 0 \end{pmatrix} = \begin{bmatrix} r_{c11}^{a+} & \\ & r_{c22} \\ 0 & 0 \end{bmatrix}, \begin{pmatrix} R_c^{b+} \\ 0 \end{pmatrix} = \begin{bmatrix} r_{c11} & \\ & r_{c22}^{b+} \\ 0 & 0 \end{bmatrix},$$

with $r_{c11}^{a+} = \sqrt{r_{c11}^2 + a_1^2}$ and $r_{c22}^{b+} = \sqrt{r_{c22}^2 + b_2^2}$.

Now, we write Eq. (21) using the above expressions, starting with $\prod_i \left|r_{1,ii}^{+}\right| \leq \prod_i \left|r_{2,ii}^{+}\right|$:

$$\left(r_{11}^{a+} r_{22}^{a+}\right)^2 \leq \left(r_{11}^{b+} r_{22}^{b+}\right)^2 \Rightarrow a_1^2\left(r_{22}^2 + r_{12}^2\right) \leq b_2^2 r_{11}^2, \tag{24}$$

and for $\prod_i \left|r_{1,cii}^{+}\right| \leq \prod_i \left|r_{2,cii}^{+}\right|$:

$$\left(r_{c11}^{a+} r_{c22}\right)^2 \leq \left(r_{c11} r_{c22}^{b+}\right)^2 \Rightarrow a_1^2 r_{c22}^2 \leq b_2^2 r_{c11}^2. \tag{25}$$

Letting $k \doteq \frac{b_2}{a_1}$, we get

$$r_{22}^2 + r_{12}^2 \leq k^2 r_{11}^2 \tag{26}$$

$$r_{c22}^2 \leq k^2 r_{c11}^2. \tag{27}$$

Thus, to prove Theorem 1 we need to show an if-and-only-if relation between Eqs. (26) and (27). To do so, we recall Eq. (11) that relates between R and R_c. Recalling that X is two-dimensional and considering, for simplicity, $w_i = w = 0.5$, it is possible to show that $r_{c11}^2 = \frac{1}{2}\frac{r_{11}^2 r_{22}^2}{r_{12}^2 + r_{22}^2}$ and $r_{c22}^2 = \frac{1}{2} r_{22}^2$. Plugging these relations into Eq. (27) results in Eq. (26). This completes the proof.

References

[1] T. Bailey, S. Julier, and G. Agamennoni. On conservative fusion of information with unknown non-gaussian dependence. In *Intl. Conf. on Information Fusion, FUSION*, 2012.

[2] M. Chli and A. J. Davison. Active matching for visual tracking. *Robotics and Autonomous Systems*, 57(12):1173–1187, December 2009.

[3] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition, 2004.

[4] V. Indelman, L. Carlone, and F. Dellaert. Towards planning in generalized belief space. In *The 16th International Symposium on Robotics Research*, Singapore, December 2013.

[5] V. Indelman, L. Carlone, and F. Dellaert. Planning under uncertainty in the continuous domain: a generalized belief space approach. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.

[6] S. Jiang, R. Kumar, and H. E. Garcia. Optimal sensor selection for discrete-event systems with partial observation. *IEEE Transactions on Automatic Control*, 48(3):369–381, 2003.

[7] S. Joshi and S. Boyd. Sensor selection via convex optimization. *IEEE Transactions on Signal Processing*, 57(2):451–462, 2009.

[8] S.J. Julier and J.K. Uhlmann. A non-divergent estimation algorithm in the presence of unknown correlations. In *American Control Conference*, pages 2369–73, 1997.

[9] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.

[10] R. Mahler. Optimal/robust distributed data fusion: a unified approach. In *AeroSense 2000*, pages 128–138. International Society for Optics and Photonics, 2000.

[11] C. Papadimitriou and J. Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.

[12] J. Pineau, G. J. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *J. of Artificial Intelligence Research*, 27:335–380, 2006.

[13] S. Prentice and N. Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *Intl. J. of Robotics Research*, 2009.

[14] A. Singh, F. Ramos, H. Durrant-Whyte, and W. J. Kaiser. Modeling and decision making in spatio-temporal processes for environmental surveillance. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 5490–5497, 2010.

[15] J. K. Uhlmann, S. J. Julier, and M. Csorba. Nondivergent simultaneous map building and localisation using covariance intersection. In *The Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Orlando, Florida*. SPIE, 1997. Navigation and Control Technologies for Unmanned Systems II.

[16] T. Vidal-Calleja, A. J. Davison, J. Andrade-Cetto, and D. W. Murray. Active control for single camera slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1930–1936, 2006.