

Accommodating Geometric And Thermodynamic Variability For Forward-Looking Infrared Sensors*

Matthew Cooper[†], Ulf Grenander[‡], Michael Miller[†], Anuj Srivastava[‡]

[†] Center For Imaging Science
Department of Electrical Engineering
Washington University, St. Louis, Missouri 63130

[‡] Division of Applied Mathematics
Brown University
Providence, Rhode Island 02912

ABSTRACT

Our work has focused on deformable template representations of geometric variability in automatic target recognition (ATR). Within this framework we have proposed the generation of conditional mean estimates of pose of ground-based targets remotely sensed via forward-looking infrared radar (FLIR) systems. Using the rotation group parameterization of the orientation space and a Bayesian estimation framework, conditional mean estimators are defined on the rotation group with minimum mean squared error (MMSE) performance bounds calculated following [1]. This paper focuses on the accommodation of thermodynamic variation. Our new approach relaxes assumptions of the target's underlying thermodynamic state, expanding thermodynamic state as a scalar field. Estimation within the deformable template setting poses geometric and thermodynamic variation as a joint inference. MMSE pose estimators for geometric variation are derived, demonstrating the "cost" of accommodating thermodynamic variability. Performance is quantitatively examined, and simulations are presented.

KEYWORDS: Lie Groups, Automatic Target Recognition (ATR), Conditional Mean Estimation, Monte Carlo Random Sampling

1. INTRODUCTION

Forward-Looking Infrared Radar (FLIR) systems are high resolution imaging sensors with resolution that is suitable for object recognition in numerous military applications. To exploit this sensor's capabilities in the Automatic Target Recognition (ATR) problem, image understanding algorithms are required to interpret remote FLIR observations of complex scenes. In this context, the targets under consideration are rigid body objects and the parameters of primary interest are target number, class, location, and orientation. Our group has previously presented an approach based on deformable templates for the ATR problem for representing the geometric variations of orientation, position, and scale, whose aim is a joint solution to the object detection, tracking, and recognition problems for video imagery [2, 3, 4, 5].

A major obstacle to the development of such deformable template representations for FLIR imagery is that the geometric variation in the appearance of a target at a given pose is only one of the principal sources of variability of the target signature. The second principal source of variation in target appearance is variation in the target's

*This work was supported by ARO/ASSERT DAAH0494G0209, ONR N00014-94-1-0859, ONR/ASSERT N00014-94-1-1135, ARO DAAH04-95-1-0494, and ARO-MURI DAAH04-96-1-0445. This work was presented at SPIE Aerosense '97 (SPIE's International Symposium on Aerospace/Defense Sensing, Simulation, & Controls), 24 April, 1997.




Figure 1: In the left panel a CAD model is shown as an example template illustrating the lattice we use to describe the target’s thermodynamic state. On the right, a FLIR observation of a T62 tank generated by PRISM demonstrating the vector field on the surface of the tank.

thermodynamic state. For this, we propose the use of deformable templates extended to describe target thermodynamic state. In previous work, we have shown such representation to yield accurate low dimensional thermodynamic descriptions [6]. In this paper, the extended templates are applied to pose estimation problems, with minimum mean squared error (MMSE) performance computed demonstrating the accuracy with which the representations accommodate thermodynamic variability.

2. DEFORMABLE TEMPLATE REPRESENTATIONS FOR ATR

We have been developing probabilistic representations in terms of a pattern theoretic formalism [7] for the inference and recognition of rigid and deformable objects in image understanding tasks. Objects are represented using templates; their infinite variety of shape is represented via transformations which act on the templates. In the context of ATR involving rigid objects, the transformations form groups; a deformable template is then the orbit under a group action, generally the Euclidean group, its subgroups, and products. Estimation becomes identification of the group action and target type, requiring matching the observed remotely sensed image with the particular instance of the template.

2.1 Representing Geometric Variation Via Deformable Templates

Using the pattern theoretic framework for ATR (see [5], for example), scenes consist of targets of type $\alpha \in \mathcal{A}$, where \mathcal{A} denotes an alphabet of possible targets, each with an associated parameter vector $s \in \mathcal{S}$ the group of transformations describing position and orientation. In the context of rigid body ATR, a template is constructed corresponding to CAD representations of the 2-dimensional surface manifolds of the rigid objects. Denote such an ideal template as $I_{temp} \equiv \{I_{temp}(x), x \in \mathcal{X}\}$, \mathcal{X} the space indexing location in the target. Shown in the left panel of Figure 1 is a CAD model defining the rigid template.

Geometric variation due to variability of pose of the targets are introduced via the rigid motions of translation and rotation. The set of transformations \mathcal{S} are Lie group actions on the templates. For ground-based scenes, we use the special orthogonal group of 2×2 rotation matrices, $\mathbf{SO}(2)$, and translations in the plane, \mathbb{R}^2 . The transformations are of the form, $s = (O, a)$, $O \in \mathbf{SO}(2)$, and $a \in \mathbb{R}^2$ a position vector. Then $s \in \mathcal{S} \equiv \mathbf{SE}(2)$, the Special Euclidean Group. $\mathbf{SE}(2) \equiv \mathbf{SO}(2) \ltimes \mathbb{R}^2$, where \ltimes here denotes the semi-direct product [8], with $\mathcal{S} : \mathcal{X} \leftrightarrow \mathcal{X}$ according to $sx \mapsto Ox + a$. The entire deformable template over which inference occurs is the orbit under $\mathbf{SE}(2)$; $\mathcal{I} \equiv \{I_{temp}(Ox + a) : (O, a) \in \mathbf{SE}(2)\}$, with elements $I(O, a) \equiv \{I_{temp}(Ox + a) : x \in \mathcal{X}\}$. Estimation of an element of the orbit $I(O, a) \in \mathcal{I}$ is estimation of its parametric representation $(O, a) \in \mathbf{SE}(2)$

2.2 Extending Deformable Templates To Represent Thermodynamic State

FLIR systems passively sense objects via their reflected or emitted infrared radiation. Unlike the video imaging systems described in [2, 3, 4], FLIR observations of a target vary dramatically with changes in the target’s thermodynamic state. Estimation of the thermodynamic state of a target is fundamental to the estimation of the target’s

type, position, and pose, given FLIR observations. Hence, understanding FLIR observations of a target entails the representation of thermodynamic variation, within our existing framework. Our group has previously presented jump-diffusion algorithms for the automated understanding of FLIR scenes [9, 10]. In that work, the thermodynamic state of the target was assumed known and constant. That approach requires the use of image libraries of templates at different thermodynamic states and target poses.

A complete description of the thermodynamic state of a target specifies the emitted radiance as a function of position on the surface of the rigid target. This surface can be decomposed into regions consisting of polygonal facets constituting a CAD model representation of the surface. This is the case in the models describing the geometry of the surface that are used by PRISM [11], the FLIR simulation software used by our group. The thermodynamic state is then described by a vector whose elements are the radiance at the lattice points defined by the template (CAD model). Define $T \equiv \{T(l) : l \in \mathcal{L}\}$ to be the thermodynamic state, defined as the radiant intensity indexed over the surface lattice \mathcal{L} given by the CAD model for the target with dimensionality $|\mathcal{L}| = L$. To extend the deformable template to accommodate thermodynamic variability, the template becomes \mathcal{I} , the set of all targets generated under the group action with superimposed scalar thermodynamic field T . We denote an element by $I(O, a, T) \in \mathcal{I}$, $\mathcal{I} \equiv \{I_{temp}(Ox + a, T) : (O, a, T) \in \mathcal{S}\}$. To illustrate, shown in the right panel of Figure 1 is the template rotated 90° with the scalar field representing thermodynamic state. Vectors are drawn normal to the surface at the various lattice sites with length proportional to the region’s radiance in (Watts/Steradian-meter²).

2.3 Expanding Thermodynamic Information as Scalar Fields Via Principal Components Analysis

Introducing thermodynamic state increases the dimensionality of the parameterization from 2 ($\mathbf{SE}(2)$) to $L + 2$. We want to reduce the dimensionality of the representation of thermodynamic state, L , to simplify the inference problem. For this we follow [6], where we defined the radiant intensity to be a scalar random field $T = \{T(l) : l \in \mathcal{L}\}$, where \mathcal{L} is the surface lattice indexing the CAD model for the target with dimensionality $|\mathcal{L}| = L$. Typically, L is several hundred, resulting in a considerable increase in dimensionality. We model T as a random field on the lattice with covariance K , where $K_{ij} = E\{(T(i) - \bar{T}(i))(T(j) - \bar{T}(j))\}$, \bar{T} is the mean thermodynamic profile of the database, and K is an $L \times L$ matrix. In the study of neuro-anatomical sub-manifolds [12], principal components analysis was established for random fields on surfaces. Following Joshi, *et al.*, [12], we apply Eigen-expansion techniques to the covariance to derive a basis representing thermodynamic variability. We represent the scalar field using a complete orthonormal basis $\{\Phi_i\}$. The Φ_i are chosen as the Eigenvectors of the empirical covariance K weighted by the surface measure Λ , satisfying $K\Lambda\Phi_i = \lambda_i\Phi_i$ (see Cooper, *et al.*, [6]). Then $T = \bar{T} + \sum \alpha_i\Phi_i$, where the coefficients $\{\alpha_i\}$ are given by $\alpha_i = \langle (T - \bar{T}), \Phi_i \rangle$.

We present principal thermodynamic Eigenprofiles for three databases. Databases were constructed to isolate two primary modes of thermodynamic variation: meteorological and operational variability. Meteorological variation, resulting from changes in the weather or atmospheric conditions, was isolated in the “static” database. Operational variation resulting from changes in the operational state of the target, such as the speed of the tank, whether the motor is on or off, how long the motor has been on, whether the main gun is firing or not, etc., was isolated in the “dynamic” database. A third, “composite” database was generated containing both of the previous two. Details regarding database construction may be found in [6].

Intuitive conclusions may be drawn from the plots and renderings of the first thermodynamic Eigen-profiles (first thermodynamic principal components) from the three databases, shown in Figure 2. The variation expressed in the first thermodynamic Eigen-profile of the static database demonstrates uniform heating of the tank as the sun moves across the sky during the twelve hour simulations. The variation represented by the first Eigen-profile of the dynamic database demonstrates the intense heating of the exhaust system of the tank as the motor runs at different speeds.

We use the empirical covariances of databases of radiance profiles over the target surface to characterize the thermodynamic variation, as observed by the FLIR sensor. Performing an Eigen-expansion of the empirical covariance matrix is equivalent to the classical statistical technique of principal components analysis [13]. Hence the Φ_i are the thermodynamic principal components. The plots of the power spectrum of the Eigenbasis representing thermodynamic state (left panel), and the mean squared error (right panel) in Figure 3 indicate that truncated Eigen-expansions yield accurate representations. We then expand the template using the coefficients of a truncated Eigen-expansion rather than the full L dimensional representation for thermodynamic state. In this way, the templates’ low-dimensionality is preserved, now accommodating the infinity of both geometric and thermodynamic variations of targets in FLIR observations. If N terms are retained in the truncated Eigen-expansion, the deformable template with reduced parameterization becomes $\mathcal{I} \equiv \{I_{temp}(Ox + a, \bar{\alpha}) : (s, \bar{\alpha}) \in \mathcal{S} \times \mathbb{R}^N\}$. Here, $\bar{\alpha} = [\alpha_1 \cdots \alpha_N]$, where $T = \bar{T} + \sum \alpha_i\Phi_i$. From Figure 3, $10 \leq N \leq 20$ appears to be sufficient for accurate representation of thermo-

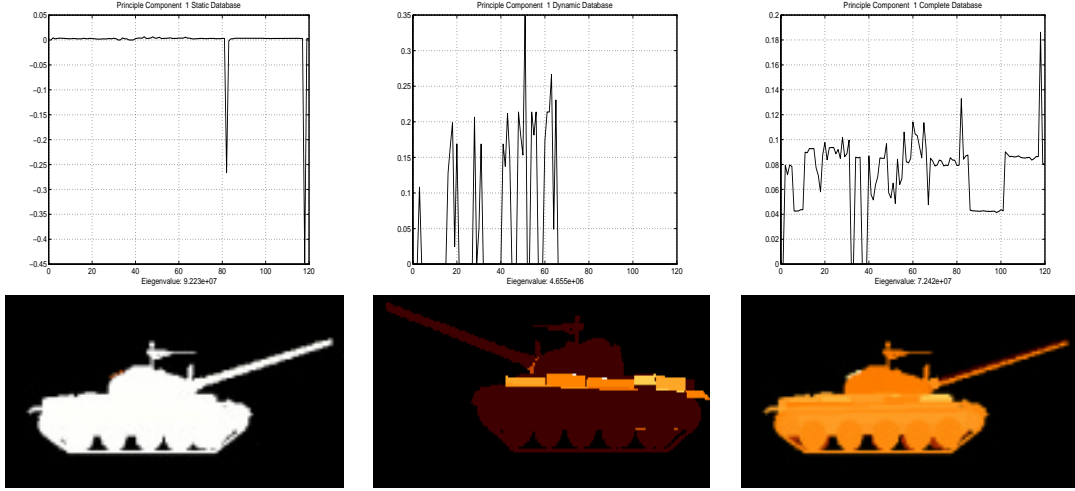


Figure 2: Visualizations of the first thermodynamic Eigen-profile given by the first principal component $\Phi_1(l), l \in \mathcal{L}$ from three databases: the top row shows the radiant intensities plotted versus lattice site index, and the bottom row shows the first principal thermodynamic profiles rendered by the SGI. From left to right, the first principal thermodynamic representing meteorological, operational, and composite modes of variation, respectively.

dynamic state. This represents a significant reduction in the dimensionality of the templates as CAD templates are routinely comprised of several thousand vertices and facets, and several hundred regions.

3. PERFORMANCE ANALYSIS FOR POSE ESTIMATION

Within the deformable template representation, performance bounds may be evaluated in the familiar sense of minimum mean squared error (MMSE) bounds for the expected variance of parameter estimators. We will review the derivation of these bounds, presented in [14], for the ground based target pose estimation problem given target class and position.

3.1 Minimum Mean Squared Error Estimators

Parameterize the set of possible orientations of the target(s) being considered by rotation matrices in $\mathbf{SO}(2)$. Select as an error metric the Hilbert-Schmidt (HS) norm: for $\mathbf{SO}(N)$, $\|X\|_{HS} = \sum_{i,j=1}^N x_{ij}^2$. In $\mathbf{SO}(N)$ then, $\|X_2 - X_1\|_{HS}^2 = 2(N - \text{Tr}(X_2 X_1^t))$. Here the superscript t denotes the matrix transpose and the operation $\text{Tr}(\cdot)$ is the matrix trace. For $N = 2$, the matrices of $\mathbf{SO}(2)$ take the form:

$$O_\theta = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}.$$

The HS norm squared of the difference between rotation matrices corresponding to the angles θ_2 and θ_1 , O_2 and O_1 , respectively is $\|O_2 - O_1\|_{HS}^2 = 4 - 4\text{Cos}(\theta_2 - \theta_1)$.

The principal difficulty addressed in [14] is the evaluation of estimators taking values on non-flat spaces, such as matrix Lie groups, including the special orthogonal groups. Summarizing that work, the MMSE estimator is defined as follows:

$$O_{HS}^*(I^D) = \underset{O \in \mathbf{SO}(2)}{\text{ArgMin}} E\{\|O - \hat{O}\|_{HS}^2 | I^D\}. \quad (1)$$

The expected squared error of this estimator then represents a lower bound, where error is measured via the HS norm. Given an observation of our target, I^D , we may evaluate the posterior distribution of the orientation conditioned on the observation. Using Bayes' rule, the posterior density $\pi(O|I^D)$ is given by the product of the data likelihood, and the prior density on target orientation:

$$\pi(O|I^D) \propto P(I^D|O)P(O). \quad (2)$$

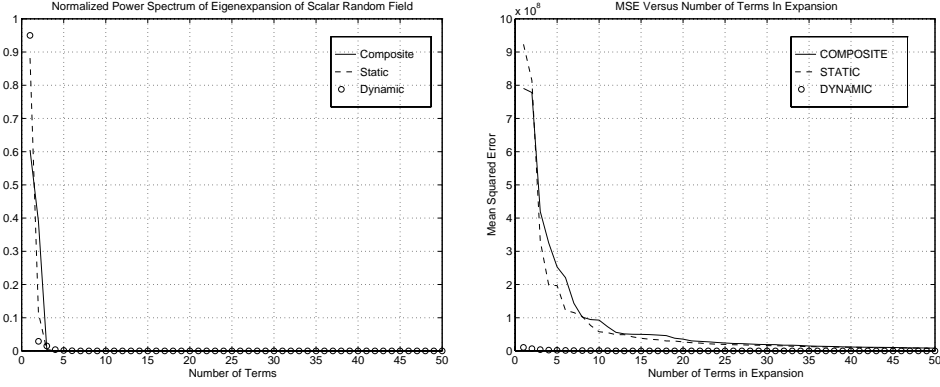


Figure 3: The left panel shows the normalized power spectrum ($\frac{\lambda_k^2}{\sum_{k=1}^L \lambda_k^2}$) of the Eigen-basis representing FLIR variability is shown. On the right appears the squared error versus the number of terms used in the Eigen-expansion, averaged over three databases.

Assuming a uniform prior, as proved in [14], the MMSE estimate is:

$$O_{HS}^*(I^D) = \underset{O \in \mathbf{SO}(2)}{\text{ArgMin}} \left\{ \int_{\mathbf{SO}(2)} \|O - \hat{O}\|_{HS}^2 \pi(\hat{O}|I^D) \gamma(d\hat{O}) \right\}, \quad (3)$$

$$= \underset{O \in \mathbf{SO}(2)}{\text{ArgMin}} \left\{ 4 - 2 \int_{\mathbf{SO}(2)} \text{Tr}(O\hat{O}^t) \pi(\hat{O}|I^D) \gamma(d\hat{O}) \right\}, \quad (4)$$

$$= \underset{O \in \mathbf{SO}(2)}{\text{ArgMax}} \text{Tr}(OA^t), \quad \text{where } A = \int_{\mathbf{SO}(2)} \hat{O} \pi(\hat{O}|I^D) \gamma(d\hat{O}), \quad (5)$$

$$= \frac{1}{|A|} A. \quad (6)$$

Here, $|\cdot|$ denotes the matrix determinant. Hence, estimation is reduced to computing $O_{HS}^*(I^D)$ as in Equation 6. The HS error (HSE) is then given by

$$HSE(O_T) = 4 - 2\text{Tr}(O_T(O_{HS}^*(I^D))^t) = 4 - \frac{2}{|A|} \text{Tr}(O_T A^t). \quad (7)$$

Here O_T is the true underlying target orientation. The HSE is a function of O_T , thus it is the conditional error with respect to the HS norm for the given estimation experiment.

The mean squared error is the expected squared error over the joint parameter space, $\mathcal{I}^D \times \mathbf{SO}(2)$, where \mathcal{I}^D is the space of possible observations. The MSE is computed as in Equation 8, in which the MSE is defined for an arbitrary estimator O^* :

$$MSE = \int_{\mathbf{SO}(2)} \int_{\mathcal{I}^D} \|O^*(I^D) - O\|_{HS}^2 P(O, I^D) \gamma_{\mathcal{I}^D}(dI^D) \gamma(dO). \quad (8)$$

Equivalently,

$$MSE = \int_{\mathbf{SO}(2)} \int_{\mathcal{I}^D} \|O^*(I^D) - O\|_{HS}^2 P(I^D|O) \gamma_{\mathcal{I}^D}(dI^D) P(O) \gamma(dO). \quad (9)$$

The MSE conditioned on the underlying target orientation, O_T , is then given by the inner integral of Equation 9:

$$MSE(O_T) = \int_{\mathcal{I}^D} \|O^*(I^D) - O\|_{HS}^2 P(I^D|O_T) \gamma_{\mathcal{I}^D}(dI^D) = \int_{\mathcal{I}^D} \left[4 - \frac{2}{|A|} \text{Tr}(O_T A^t) \right] P(I^D|O_T) \gamma_{\mathcal{I}^D}(dI^D). \quad (10)$$




Figure 4: Here the T62 tank is displayed in four thermodynamic states. In the top row: profile 140 on the left and profile 75 on the right. In the bottom row: profile 45 on the left and profile 8 on the right. These differing thermodynamic states are used to quantify performance variation due to inaccurate thermodynamic information.

The second equality is given by substituting the HS estimator for O^* . The MSE is evaluated using numerical approximations to the integral over $\mathbf{SO}(2)$ (concealed in A , Equation 5) and Monte Carlo random sampling to compute the expectation over the space of possible observations, \mathcal{I}^D . Thus, the MSE is computed in a specific noise environment specified by the noise standard deviation, σ in the case of Gaussian noise, or more generally the signal to noise ratio, according to

$$MSE_{\sigma}(O_T) = \frac{1}{N} \sum_{i=1}^N \|O_{HS}^*(I_i^D(\sigma)) - O_T\|_{HS}^2 = \frac{1}{N} \sum_{i=1}^N \left[4 - \frac{2}{|A_i(\sigma)|} \text{Tr}(O_T A_i^t(\sigma)) \right]. \quad (11)$$

Here, $A_i(\sigma)$ represents the average matrix computed using Equation 6 given the i^{th} observation of a set of N observations, $\{I_i^D(\sigma)\} \subset \mathcal{I}^D$, in which the noise environment is parameterized by σ .

To evaluate the posterior density $\pi(O|I^D)$, images of the target are simulated with known thermodynamic state and a sampling of orientations $\mathcal{O} = \{O_i : i = 1 \dots M\} \subset \mathbf{SO}(2)$, These are stored in a database. The likelihood of the data conditioned on true orientation O_i is evaluated using the observation and the corresponding database image. Equation 2 is then applied to give $\pi(O_i|I^D)$. $\pi(O_i|I^D)$ is computed for each $O_i \in \mathcal{O}$ and normalized appropriately. The A matrix is then empirically computed per Equation 5. The HS estimate and the HSE are evaluated. This procedure is repeated for a number of realizations of the noise process for a specific value of σ to empirically approximate the conditional MSE, as in Equation 11. Results using FLIR sensors in Gaussian noise environments are shown in Figure 5. In the simulations presented throughout this paper, only orientations in the interval $[0, \frac{\pi}{2})$ are considered, so that target symmetry effects may be ignored (see [1]). In the left panel, the conditional mean squared error (MSE(30)) is plotted versus noise standard deviation for orientation estimation of the T62 tank, conditioned on an underlying orientation of 30° . The curves were generated using observations of the tank at the same pose with the underlying thermodynamic profiles of Figure 4. The databases corresponding to the thermodynamic state of the observation are then used to calculate performance curves. Notice the performance variation due to the thermodynamic variability.

To this point we have assumed that the target's thermodynamic profile is known and constant. Relaxing this assumption, the performance loss incurred by assuming incorrect thermodynamic information may be quantified. Working with the set of thermodynamic profiles shown in Figure 4, noisy FLIR observations are generated using thermodynamic Profile 140. Databases are generated for the set of orientations \mathcal{O} assuming the true thermodynamic state and the two additional target thermodynamic profiles shown in Figure 4: Profile 75 and Profile 45. Performance curves are computed using each database, conditioned on a true target orientation of 30° . In the latter two

cases, the observation and the database images compared via the data likelihood assume different underlying target thermodynamic profiles. Resulting performance losses are quantified by the curves shown in the right panel of Figure 5. Performance given correct thermodynamic information corresponds to the curve labeled “Database 140”. The mismatched curves demonstrate performance dependence on accurate thermodynamic information.

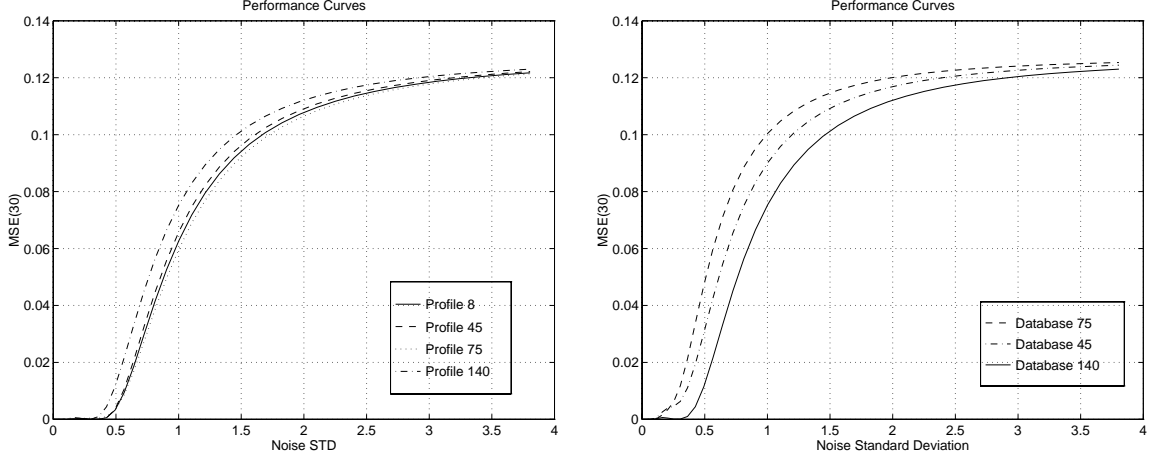


Figure 5: In the left panel, the plot of MSE versus noise standard deviation for orientation estimation of a T62 tank is shown. The plot shows performance variations due to changes in the underlying thermodynamic profile of the target. In the right panel, performance curves appear in which the underlying thermodynamic profile of the observed target is given by Profile 140, while the estimation algorithms assume Profiles 140, 45, and 75. Performance losses are observed in the mismatched cases.

The unconditional MSE may be evaluated by evaluating the outer expectation of Equation 9 over $\mathbf{SO}(2)$. The unconditional MSE is evaluated as:

$$MSE = E\{E\{\|O_T - O_{HS}^*\|_{HS}^2 | O_T\}\}. \quad (12)$$

Assuming a uniform prior on the underlying target orientation O_T , the MSE maybe calculated empirically by averaging the conditional MSE over \mathcal{O} :

$$MSE = \frac{1}{M} \sum_{i=1}^M MSE(O_i). \quad (13)$$

A plot of the MSE appears in Figure 6.

4. PERFORMANCE ANALYSIS & POSE ESTIMATION ACCOMMODATING THERMODYNAMIC VARIABILITY

The pose estimation experiment has been performed employing the representation for thermodynamic variation expanding thermodynamic state via the thermodynamic Eigenbasis, relaxing assumptions regarding the target’s thermodynamic state. In this case, the joint posterior is evaluated over the parameter space of Section 2.3, $\mathbf{SO}(2) \times \mathbb{R}^N$.

$$\pi(O|I^D) = \int_{\mathbb{R}^{20}} \pi(O, \bar{\alpha}|I^D) \gamma_{\bar{\alpha}}(d\bar{\alpha}) \quad (14)$$

The target’s thermodynamic state is represented by a vector of coefficients, which are the projections of the thermodynamic profile on to the basis of thermodynamic Eigen profiles. This vector is denoted $\bar{\alpha}$, where $T = \bar{T} + \sum \alpha_i \Phi_i$. Here, the first twenty terms of the generalized Karhunen-Loeve expansion are retained, i.e. $N = 20$, $\bar{\alpha} \in \mathbb{R}^{20}$. The




Figure 6: The unconditional MSE with respect to the HS norm for the ground-based pose estimation experiment is plotted versus noise standard deviation. It is evaluated by empirically averaging the conditional MSE given the underlying target orientation over the space of target orientations.

joint posterior $\pi(O, \bar{\alpha}|I^D)$ is given by Bayes' rule as the product of the joint data likelihood, $P(I^D|O, \bar{\alpha})$ and the joint prior on the target orientation and thermodynamic state, $P(O, \bar{\alpha})$. We make two remarks on the joint prior. First, the target's pose and thermodynamic state are assumed independent, considering a single observation. Second, the generalized Karhunen-Loeve expansion used to generate the Eigenbasis of the target's possible thermodynamic profiles implies a prior on the vector $\bar{\alpha}$. Modeling the thermodynamic state as a Gaussian random field, the coefficients of the target's Eigenbasis representation, α_i , are independent Gaussian random variables with zero mean and variance equal to λ_i^2 , where λ_i is the Eigenvalue of the i^{th} thermodynamic Eigenprofile. Assuming a uniform prior on target orientation:

$$\pi(O|I^D) = \frac{1}{Z} \int_{\mathbb{R}^{20}} P(I^D|O, \bar{\alpha})P(\bar{\alpha})\gamma_{\bar{\alpha}}(d\bar{\alpha}). \quad (15)$$

Z is the appropriate normalizer for the posterior. Monte Carlo random sampling in the space of possible thermodynamic states is used to evaluate this integral. Thus the estimation algorithm uses randomly sampled thermodynamic information, weighted by the prior $P(\bar{\alpha})$, in place of assuming a target thermodynamic profile *a priori*.

We have arrived at a means of evaluating the posterior density of our parameter of interest, $\pi(O|I^D)$, without assumptions regarding the target's thermodynamic state. This new expression for $\pi(O|I^D)$ may be inserted into Equation 5 for the HS estimator, and the HSB and performance curves are calculated as before in Equations 10 and 11:

$$O_{HS}^*(I^D) = \frac{1}{Z} \int_{\text{SO}(2)} \hat{O} \int_{\mathbb{R}^{20}} P(I^D|O, \bar{\alpha})P(\bar{\alpha})\gamma_{\bar{\alpha}}(d\bar{\alpha})\gamma(d\hat{O}). \quad (16)$$

In Figure 7, the performance of the pose estimators with correct and randomly sampled thermodynamic information is compared. In these simulations, the underlying thermodynamic state of the tank is given by Profile 8 of Figure 4. The Y-Axis represents the conditional MSE given an underlying target orientation of 30° . The curves labeled "Ideal Performance" correspond to ideal performance with correct thermodynamic information, and the curves labeled "Randomized Performance" correspond to performance using the random sampling approach, without assumption of the target thermodynamic profile. The difference between these two curves represents the "cost" of knowledge of the target's thermodynamic profile, in the context of the estimation problem.

In Figure 8, the performance of the new estimators, accommodating thermodynamic variation, is compared with the performance of estimators with incorrect thermodynamic information, similar to those in the right panel of Figure 5. In these simulations, the observations were created using Profile 8 of Figure 4. In the left panel, the performance of




Figure 7: The performance of pose estimators with correct and randomly sampled thermodynamic information. The curve labeled “Ideal Performance” corresponds to ideal performance, and the curve labeled “Randomized Performance” corresponds to the performance of the estimator using randomly sampled thermodynamic information. In the left panel, MSE given by the HS norm squared is plotted versus increasing noise standard deviation. In the right panel, estimation error is plotted in degrees versus the noise standard deviation.

our new estimator, labeled “Random” is compared to ideal performance, “Database 8”, and two curves corresponding to performance assuming two incorrect thermodynamic profiles shown in Figure 5, “Database 45” and “Database 75”. The random sampling approach outperforms the non-randomized estimators with incorrect thermodynamic information. The estimator with correct known thermodynamic information performs most optimally. Plots for a second set of these simulations is shown in the right panel of Figure 8. In this case, the true thermodynamic profile is Profile 140 of Figure 4, and performance of the new conditional mean estimator, without assumption of thermodynamic state, is almost identical to that of the estimator with correct thermodynamic information.

5. ORIENTATION ESTIMATION ACCOMMODATING THERMODYNAMIC VARIATION

To apply the expanded template representations, a ground-based scene is considered and the goal is the estimation of the orientation of a target of known class at a known position using a noisy FLIR observation. The target’s orientation is described by a single rotation $O \in \mathbf{SO}(2)$. To illustrate the application of the extended template $I(a, O, \bar{\alpha})$, we examined the maximum *a posteriori* (MAP) estimator on sample data. The first 20 terms of the Eigen-expansion are retained, thus the thermodynamic state will be given by a vector of coefficients, $\bar{\alpha} \in \mathbb{R}^{20}$. The approach is to maximize the posterior probability over the *joint* parameter space, $\mathbf{SO}(2) \times \mathbb{R}^{20}$. We perform the maximization via gradient descent using numerical approximations to the derivative of the joint posterior [15].

The observations are generated by PRISM. In addition, the effects of both charge-coupled device (CCD) camera photocounting noise, modeled as Poisson noise, and a Gaussian camera pointspread function are also incorporated [16]. Furthermore, dead and saturated pixels are added to the observation. Two observations appear in the left column of Figure 9. The algorithm is initialized using an exemplar thermodynamic state and an arbitrary orientation, as seen in the middle column of Figure 9 for the experiments. The final MAP estimates produced by the algorithm appear in the right column of Figure 9, and in both trials, the target orientation is correctly estimated.

6. CONCLUSION

In this paper, a low dimensional representation accommodating the infinite variations of pose and thermodynamic state of targets observed via FLIR sensors has been presented. Performance analysis, facilitated by the model-based Bayesian formulation of the ATR problem, has been presented. Conditional mean estimators for ground-based target pose estimation problems have been reviewed using the Lie group pose parameterization. Thermodynamic




Figure 8: The performance of pose estimators with with correct, randomly sampled, and incorrect thermodynamic information. In both panels, the performance of the estimator with randomly sampled thermodynamic information, labeled “RANDOM”, is compared to the ideal performance, given by “Database 8” in the left panel, and by “Database 140” in the right panel. Also shown are curves corresponding to performance given incorrect assumptions of the target’s thermodynamic state: “Database 45” and “Database 75”.

information has been incorporated into this development, and estimators using the extended template have been derived. Simulations using Monte Carlo random sampling have been presented to quantitatively assess the impact of both thermodynamic variability and the extended template representation of thermodynamic state in the pose estimation scenario. Finally, a MAP estimation example was presented in which the target pose and thermodynamic state were jointly estimated.

Results presented throughout this paper have repeatedly demonstrated the intrinsic association of accurate thermodynamic representations to image understanding with FLIR systems in the ATR context. Performance in the pose estimation problem is seen to depend on accurate thermodynamic representation, and the random sampling approach in this scenario has been shown to yield high performance without *a priori* assumptions of the underlying target thermodynamic state. In the MAP estimation examples presented, the joint estimation approach again exploits this relationship effectively.

9. REFERENCES

- [1] U. Grenander, M.I. Miller, and A. Srivastava. Hilbert-Schmidt Lower Bounds For Estimators on Matrix Lie Groups. Submitted to *International Journal of Pattern Analysis & Artificial Intelligence*.
- [2] M.I. Miller, R. S. Teichman, A. Srivastava, J.A. O’Sullivan, and D. L. Snyder. Jump-diffusion processes for automated tracking-target recognition. In *Proceedings of the Twenty-Seventh Annual Conference Conference on Information Sciences and Systems*, pages 617–622, Baltimore, Maryland, March 24-26 1993. Johns Hopkins University.
- [3] A. Srivastava, M.I. Miller, and U. Grenander. Multiple target direction of arrival tracking. *IEEE Transactions on Signal Processing*, 43(5):1282–1285, May 1995.
- [4] M.I. Miller, A. Srivastava, and U. Grenander. Conditional-mean estimation via jump-diffusion processes in multiple target tracking/recognition. *IEEE Transactions on Signal Processing*, 43(11):2678–2690, November 1995.
- [5] M. Miller, U. Grenander, J. O’Sullivan, and D. Snyder. Automatic Target Recognition Organized Via Jump-Diffusion Algorithms. *IEEE Transactions on Image Processing*, 6(1):157–174, January, 1997.




Figure 9: In each row, simulation results for joint MAP estimation of target orientation and thermodynamic state. The left column shows two simulated noisy FLIR observations of a T62 tank. In the middle column initial estimates for the two trials generated by exemplar thermodynamic data and an arbitrarily selected orientation. The right column shows the final estimates for each trial, in which target orientation is correctly estimated.

[6] M.L. Cooper, A.D. Lanterman, S.C. Joshi, and M.I. Miller. Representing The Variation of Thermodynamic State Via Principal Components Analysis, Proceedings of the Third Workshop On Conventional Weapon ATR. Redstone Arsenal, AL, November 1996.

[7] U. Grenander. *General Pattern Theory*. Oxford Univ. Press, 1994.

[8] W. M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry*. Academic Press, 1986.

[9] A.D. Lanterman, M.I. Miller, D.L. Snyder, and W.J. Miceli. The unification of detection, tracking, and recognition for Radar/Ladar Processing, Proc. SPIE, Vol. 2562, Ed. William J. Miceli, July 1995, pp. 150-16

[10] A.D. Lanterman, M.I. Miller, and D.L. Snyder. Implementation of jump-diffusion processes for understanding flir scenes. In F.A. Sadjadi, editor, *Automatic Object Recognition V*, volume 2485, pages 309–320, Orlando, FL, April 1995. SPIE.

[11] Prism 3.1 Users' Manual. Keweenaw Research Center, Michigan Technological University, 1987.

[12] S. Joshi, M. Miller, and U. Grenander. On the Geometry and Shape of Brain Sub-Manifolds. To Appear in *International Journal of Pattern Analysis & Artificial Intelligence*.

[13] R. Muirhead. *Aspects of Multivariate Statistical Theory*. J. Wiley & Sons, 1982.

[14] A. Srivastava, Inferences on Transformation Groups Generating Patterns On Rigid Motions. Doctoral Thesis, Washington University, August 1996.

[15] A. D. Lanterman, Jump-diffusion algorithms for the automated understanding of forward-looking infrared scenes. Master's thesis, Washington University, May 1995.

[16] D. L. Snyder, A. M. Hammoud, and R. L. White, Image Recovery from Data Acquired with a Charge-Coupled-Device Camera, *Journal of the Optical Society of America*, vol. 10, pp. 1014-1023, May 1993.