
LIST OF TABLES

3.1	Data support for bigram and trigram estimation.	82
3.2	A summary of the automatic speech recognition system from a training corpus of 45 h with 2.2 million words from the 1997 CLSP Workshop [109] on a test corpus of 93 minutes with the dimension of the features 39.	92
5.1a	Showing the maximum-likelihood estimates of the parameters μ , variance σ^2 , and restoring force a in the isotropic model.	147
5.1b	Estimates of the texture parameters $a_{11}; a_{12}; a_{22}$ for the actin–myosin complex and the mitochondria for the orientation dependent model.	147
5.2	Maximum-likelihood estimates derived from the simulated images of the top row of Figure 5.9. Left 3 columns show the parameters used for the simulations; right 3 columns show the estimates.	149
5.3	Maximum-likelihood parameter estimates derived from the simulated images.	151
8.1	Maximum-likelihood estimates of the parameters μ , noise variance σ^2 , and restoring force as estimated from the mitochondria data using the algorithm from Example 5.26 of Chapter 5.	242
16.1	Matching distances and geodesic distance for Christensen algorithm and LDDMM.	480
19.1	Table showing the maximum-likelihood estimates of the parameters μ , noise variance σ^2 , and restoring force a estimated from the mitochondria data using the algorithm from Example 5.26 of Chapter 5.	556
19.2	Quantifying the performance of the pixel-based optimal Bayesian hypothesis testing (left column) compared with the segmentation via the global shape models. The true labels were determined by segmentation performed manually. Percentages indicate relative frequency of proper pixel labeling for each model type.	559

LIST OF FIGURES

2.1.	Shannon’s source channel model for communications systems.	5
2.2.	Left shows absolute value; middle shows squared error; right panel shows a thresholding function.	7
2.3.	Left column: The figure shows the results from the N-butyl alcohol experiment. The top panel shows the original 2D spectrum of the N-butyl alcohol data. The bottom panel shows the spectrum reconstructed from the estimates of the EM-algorithm parameters. The data are taken from [23]. Right column: The figures show the EM algorithm reconstruction of the COSM data taken from Dr. Keeling of Washington University; the data reconstructions are from [30]. The top row shows X–Z sections through the COSM amoeba data. The bottom row shows the 200th EM algorithm iteration for sections through the COSM amoeba data (see also Plate 1).	33
2.4.	Panel 1 shows the EM algorithm fit of G,W,CSF and partial volume compartments to brain tissue histograms to illustrate the segmentation calibration. The top solid curve superimposes the measured histogram data almost exactly. The lower dashed lines depict each of the compartment fits taken separately. Panel 2 shows an MRI section of the cingulate; panel 3 shows the Bayes segmentation into G,W, CSF compartments of coronal sections of the cingulate gyrus; Panels 4 and 5 show the same as above (row 1) for the medial prefrontal gyrus. Data taken from the laboratories of Dr. John Csernansky and Dr. Kelly Botteron of Washington University (see also Plate 2).	37
3.1.	Panel 1 shows a directed graph consisting of sites D and edges E ; panel 2 shows an undirected graph in which all edges have an orientation with no cycles.	49
3.2.	Left panel shows a directed graph LINEAR associated with a Markov chain; middle panel a graph type TREE associated with a random branching process; right panel most general DAG directed acyclic graph.	50
3.3.	Showing a sequence of graphs generated by sequentially peeling the leaves (those sites with children).	52
3.4.	Bayesian belief network for medical diagnosis. From Cowell et al. [70].	53
3.5.	Genetic regulatory network candidate for yeast; taken from Hartemink et al. [71].	54
3.6.	Panels showing two family generation trees from the 1-type process with particles A , evolving under different choices of birthing rules. Note that both families have the identical states $Z_0 = 1, Z_1 = 3, Z_2 = 2$.	57
3.7.	Left panel showing an example family tree from the 2-type <i>reverse polish</i> , with particles that birth according to the two rules $A_1 \rightarrow A_1 A_1 A_2, A_1 \rightarrow \mathbb{N}, A_2 \rightarrow \{+, *\}$. The derivation is $((3\ 4\ +)\ 2\ *)$ with state sequence $Z_0 = [1, 0], Z_1 = [2, 1], Z_2 = [2, 1]$. Right panel showing a 3-generation realization from a 2-type process, $p_{1,(1,1)} = 1, p_{2,(0,0)} = 1$, demonstrating the recurrence of the state $[1, 1]$ starting from $Z_0 = [1, 0]$.	58
3.8.	Showing two finite-state graphs corresponding to run-length (left panel) and parity (right panel) languages. For the 1-1 languages, both states are accepting; for even parity the even state is accepting.	76
3.9.	Figure shows super-exponential growth rates of the arithmetic expression language. Panel 1 shows $\log - \log$ growth rate with derivation depth n . Panel 2 shows the near line ratio of the logarithms of successive generations. Panel 3 shows branching rate upper and lower bounds for ρ . Results taken from O’Sullivan [79].	80
3.10.	Three graph structures for natural language: Panel 1 Markov chain graph, panel 2 tree graph, and panel 3 directed acyclic graph.	81
3.11.	Stochastic context-free grammar constructed by Steven Abney.	83
3.12.	Left panel shows two components of a derivation tree T : the tree t deriving the preterminal string “Art N V” and the word string $W_t =$ “The dog ran”. Right panel shows comparison of the model entropy of four language models, bigrams, trigram, context-free, and directed acyclic graph.	86

3.13.	The Hidden Markov model with hidden state sequence X_1, X_2, \dots and conditionally independent outputs Y_1, Y_2, \dots given the state-sequence.	88
3.14.	Representation of the Viterbi algorithm on a three-state state space. The arcs in bold show the path which obtains the state sequence with the highest output probability.	89
4.1.	Figure showing various undirected graphs.	95
4.2.	Left panel shows a 3×3 lattice. Middle panel shows the neighborhood system $N_5 = \{i_2, i_4, i_6, i_8\}$ for i_5 . Right panel shows the clique system of i_5 , $\{i_2, i_5\} \cup \{i_5, i_8\} \cup \{i_4, i_5\} \cup \{i_5, i_6\}$.	96
4.3.	Panels 1 and 2 show electron micrographs at 30,000 magnification containing mitochondria and cytoplasm. Panels 3 and 4 show the Bayes segmentation via an MRF model of nearest-neighbors with zero boundary and four gray levels. Results from [127, 128]; data from Jeffrey Saffitz of the Department of Pathology at Washington University.	98
4.4.	Left panel shows the directed graph structure corresponding to the joint X, Y process. Right panel shows the completely connected undirected graph corresponding to the marginal process Y only.	99
4.5.	Left panel shows the directed graph branching process structure for the joint $T = R, W$ process. The $\gamma_1, \gamma_2, \dots, \gamma_n$ denote the pre-terminal symbols. Right panel shows resulting completely connected undirected graph corresponding to the marginal process on the word W .	100
4.6.	Left panel shows a 2D configuration of dipoles; right panel shows the graph structure for the Ising model.	103
4.7.	Showing the interior and closure of the sets G^0, \bar{G} . Double bold line denotes the set G , with G^0 the interior square, and \bar{G} the exterior square.	105
4.8.	Figure taken from Mumford [133] showing the solutions from the Ising model following the cooling schedule.	108
4.9.	Left panel shows pixel sites marked with circles and edge sites marked with crosses; right panel shows the dual lattice with edges and pixels.	109
4.10.	Panel 1 shows animal fur texture, panel 2 shows white noise, panel 3 shows maximum-entropy with 5 pixel Laplacian of a Gaussian filter f_1 , panel 4 shows added filters $f_{6,120^\circ}, f_{2,30^\circ}$, then panel 5 adding the filter $f_{12,60^\circ}$, and panel 6 adds the filter $f_{10,120^\circ}$ and DC component.	114
4.11.	Panel 1: observed texture of mud. Panels 2–7 show synthesized versions from Zhu model.	115
4.12.	Panel 1: observed texture of fabric. Panels 2–7 show synthesized versions from Zhu model.	115
4.13.	Left panel shows a configuration with positive boundary condition. Right panel shows the bounding box for the circuit of length L .	118
5.1.	Left panel shows the graph induced by the first order difference operator; right panel shows the graph associated with the second order Laplacian difference operator.	130
5.2.	Top row: Panel 1 shows a noise realization driving the stochastic difference equation with $\sigma = 100, \mu = 128$. Panels 2–4 show the solution of the stochastic difference equation with $a = 1$ (panel 2), $a = 4$ (panel 3), and $a = 12$ (panel 4). Bottom row: Histograms of statistics of the images.	132
5.3.	Random fields and histograms resulting from different sets of parameters. Rows have restoring forces $a = 1, a = 4$ and $a = 12$ (top to bottom) for fixed noise mean $\mu = 128$, with noise standard deviations of $\sigma = 50$ (columns 1, 2) and $\sigma = 200$ (columns 3, 4).	132
5.4.	Panels show textures being synthesized with stationary Gaussian process. Top row: Panel 1 shows the original fur synthesis, panel 2 shows synthesis with the largest 8 eigenvalues, panel 3 shows synthesis with the largest 128 eigenvalues, panel 4 shows synthesis with the largest 512 eigenvalues, and panel 5 shows synthesis with all of the eigenvalues. Bottom row is similar but for fabric. Taken from Dimitri Bitouk.	136
5.5.	Figure presents the animal spectrum and the spectrum estimated from the synthesized textures. The two panels on the left show sparsity pattern of the animal fur texture eigenvalues and the eigenvalues estimated from the synthesized texture images; the two panels on the right show empirical versus synthesized spectra. Taken from Dimitri Bitouk.	136

5.6.	Panel 1 shows the micrograph data including mitochondria and cytoplasmic structures. Panel 2 shows a decomposition into two region types, mitochondria (white) and background (black). Panel 3 shows the pixel-by-pixel segmentation of the image based on the Bayesian hypothesis test under the two models. Data taken from Dr. Jeffrey Saffitz of the Department of Pathology at Washington University.	137
5.7.	Showing maximum-likelihood estimates of μ, σ, a from an experiment of 200 random fields generated with $\mu = 128, \sigma = 50$, and $a = 1$. Results courtesy of David Maffit, Washington University at St. Louis.	146
5.8.	Panels 1–3 show example micrographs containing mitochondria and actin–myosin complexes; panel 4 illustrates the hand segmentations of the mitochondria and actin–myosin shapes used to estimate the oriented texture parameters. Everything not labelled in gray is not mitochondria or actin–myosin. Data taken from Dr. Jeffrey Saffitz of the Department of Pathology at Washington University.	146
5.9.	Top row: Synthetic Gaussian texture images with directional components. Panels 1–4 show simulated random fields with varying parameter choices. Panel 1 shows fields based on parameters estimated from micrographs of the actin–myosin complexes; panel 2 shows similar parameters rotated to align with the Y -axis; panel 3 shows the thicker grain parameters; panel 4 shows these rotated by 60° . Middle row: Simulated mitochondria textures; panels 5 and 6 show the preferable first-order model and the less-likely second order model, with panels 7 and 8 showing the mixed and first-order models, respectively. Bottom row: Panel 9 shows cityscape observed via an infrared imager. Panel 10 shows a subimage of urban clutter extracted from the left portion of the cityscape. (Data courtesy Howard McCauley, Naval Air Warfare Center Weapons Division, NAWCWPNS, China Lake, CA. Approved for public release; distribution is unlimited.) Panels 11,12 show simulated infrared urban clutter; panel 11 shows the preferable second-order model; panel 12 shows the less-likely first order model.	149
6.1.	Panel 1 shows two generators with their bonds; panel 2 shows a configuration with the set of generators, bonds, and edges in the graph.	155
6.2.	Top and bottom rows show configurations $LINEAR(n, g_1, \dots, g_n)$ and $CYCLIC(n, g_1, \dots, g_n)$.	156
6.3.	Top row: Panel 1 shows the triangulated graph for the template representing amoeba corresponding to spherical closed surfaces. Generators are $g_i = (v_1, v_2, v_3)$, elements of \mathbb{R}^9 . Panel 2 shows the triangulated for the closed surface representing the bounding closed surface of the hippocampus in the human brain. Panel 3 shows the surface representation of a macaque brain from the laboratory of David Van Essen. Bottom row: Shows a generator shown in standard position, orientation under rotation and translation transformation. (see plate 3).	157
6.4.	Figure showing graph transformations forming and deleting tracks and track segments.	158
6.5.	Rewriting rules corresponding to the finite-state, context-free, and context-sensitive generators.	159
6.6.	Panel 1 shows the configuration diagram $TREE(6, r_1, r_2, r_3, r_6, r_4, r_7)$ for the sentence <i>The frogs jump</i> ; panel 2 shows the configuration diagram $TREE(6, r_1, r_2, r_3, r_6, r_4, r_5)$ for the sentence <i>The frogs eggs</i> , which is not locally regular. Panel 3 shows the configuration $PHRASE-FOREST(2, c_1, c_2)$ of two phrases $PHRASE(3, r_2, r_6, r_4)$, $PHRASE(2, r_3, r_7)$ each locally regular. Panel 4 shows the configuration $DIGRAPH(8, r_1, r_2, r_4, r_5, r_6, r_7, r_3, r_8)$ which is context-sensitive and is a $DIGRAPH$.	161
6.7.	Left panel depicts MON and $MULT$ extensions. Right panel shows figure tracking airplanes as $LINEAR$ graphs with generators.	163
6.8.	Top row shows graphs from $\Sigma = PHRASE-FOREST$ corresponding to a sequence of graph transformations generated during the parsing process of the sentence <i>The frogs jump</i> . Left panel shows the starting configuration; successive panels show parsing transformations from the phrase–structure context-free grammar. Bottom row shows a sequence of transformations through $\Sigma = PHRASE-FOREST$ to the root node S resulting in the synthesis of the sentence <i>The frogs jump</i> .	164
6.9.	Top panel shows the set of elementary graph transformations over a sample DAG . Lower panel depicts a sequence of transformations resulting in an illegal structure.	165

6.10.	Gene subnetwork for mating response in <i>Saccharomyces cerevisiae</i> , from [D. Pe'er: From Gene Expression to Molecular Pathways. Ph D dissertation Hebrew University]. The widths of the arcs correspond to feature confidence and they are oriented only when there is high confidence in their orientation.	165
6.11.	Top row shows $m = 1$ -memory Markov chain with generators $g = (x, y)$ the pairs in the boxes and bond values $\beta_{\text{in}}(g) = x, \beta_{\text{out}}(g) = y$. Bottom row shows for $m = 2$ -memory, with $g = (x, y, z)$ and $\beta_{\text{in}}(g) = (x, y), \beta_{\text{out}}(g) = (y, z)$.	168
6.12.	Left panel showing various cliques of size three and four transformed to the binary graph case (right panel).	170
6.13.	Top row: Panel 1 shows the grouping of lattice sites into generators; panel 2 shows the bond values for generator $g_{i,j} = (X_{i,j}, X_{i+1,j}, X_{i,j+1})$ for the derivative operator. Bottom row shows the same for the Laplacian operator.	172
7.1.	The panels 1–3 show various CAD models for objects. Panel 4 shows the human brain section depicting symmetry.	180
7.2.	Showing the circular closed contour template and its deformation via the scale-rotation groups.	184
7.3.	Top row: Panel 1 shows the basic leaf connector; panel 2 shows the forest of the maple leaf with panel 3 showing the corresponding independent loops. Bottom row: Panel 4 shows a second forest with panel 5 showing its associated independent loops.	187
7.4.	(i) Panel 1 shows the three orthonormal vector fields T, N, B of the curve x . Panel 2 shows a surface with the normal vector field and the corresponding shape operator $S(v)$.	189
7.5.	Panels 1–3 show the three macaque brain sulcal curves used for estimating the random sulcal model. Panel 4 shows the mean sulcal curves computed from the three brains displayed on the average brain generated from the 3 in the left panels; panel 5 shows samples of the sulcal curves from the distribution on curvature and torsion. Data courtesy of David Van Essen, Washington University.	193
7.6.	Shape operator on surfaces.	198
7.7.	Panel 1 shows the positive and negative normal curvatures. Panel 2 shows the case of $K(p) > 0$. Panel 3 shows the case of $K(p) < 0$; panel 4 shows the case $K(p) = 0$.	200
7.8.	Figure depicts the neighborhood of a point p and the three vectors E_{3p}, E_{1p}, E_{2p} used for fitting the local quadratic charts; taken from Joshi et al. [158].	201
7.9.	Panel 1 shows the resulting single triangle (shown as thick lines) from a cube where only 1 vertex (solid circle) is considered inside the surface. Panel 2 shows a slightly more complicated example where 3 vertices are considered inside the surface.	202
7.10.	Top row: Panel 1 shows the triangulated graph representing the hippocampus generated by isocontouring; panel 2 shows the mean curvature map superimposed on the hippocampus surface of the human. Data taken from Haller et al. [162]. Bottom row: Panels 3 and 4 show the same as above for macaque cortical surface reconstructed from cryosection data from the laboratory of Dr. David Van Essen. Bright areas represent areas of high positive mean curvature; dark areas represent areas of high negative mean curvature (see Plate 4).	203
7.11.	Top row: Panel 1 shows the atlas depiction of the occipital cortex. Panel 2 shows the reconstruction of the occipital cortex depicting the major sulcal and gyral principal curves including the inferior Calcarine sulcus and Lingual and Parietal gyri. Data taken from Dr. Steven Yantis of the Johns Hopkins University (see Plate 5).	204
7.12.	Left column shows the medial prefrontal cortex section from the Duvernoy atlas. Middle column top and bottom panels show the isosurface reconstruction of the prefrontal medial cortex. Sections through the two different MRI brains show the embedded surfaces. Right column, top and bottom panels, shows the medial prefrontal cortex reconstructions. Data taken from Dr. Kelly Botteron of Washington University (see Plate 6).	204
7.13.	Top row: Panel 1 shows eight geodesics generated on the neocortex by picking the start and end points manually. Panel 2 depicts geographical landmarks on the macaque cortex; labels 1, 2, 3, 4, 5, 6. Panel 3 shows a table of Riemannian distances in millimeters between the predefined points. Data taken from the laboratory of David Van Essen, Washington University. Bottom row: Figure shows optimality of dynamic programming. Panel 4 shows the Visible Human cortex extracted by David Van Essen; panel 5 shows choosing multiple terminal points for the DP solution; panels 6, 7 show the DP generation of the superior	

	temporal sulcus jumping across the break connecting the start and end points which were manually selected (see Plate 7).	209
7.14.	Top row panel 1 shows the external view of the Superior Temporal Gyrus (STG). Panel 2 shows Heschl's gyrus and the posterior boundary of the plenum temporale (PT) defined via dynamic programming. Panel 3 shows the delineation of the STG surface into two with the PT as the blue region which is extracted. Bottom row shows the application of dynamic programming to extract the PT from STG surface. Panel 4 tracks Heschl's gyrus; panel 5 tracks the STG as far as the posterior ascending (or descending) ramus; panel 6 tracks the geodesic from the end of the STG to the retro-insular end of the Heschl's gyrus. Data taken from the laboratory of Drs. Godfrey Pearlson and Patrick Barta, reconstructions from Dr. Tilak Ratnanather (see Plate 8).	209
7.15.	Panel 1 (Panel 3) show the radius of circle at vertex with positive (negative) curvature must be reduced (increased) to ensure that circles at all vertices are tangent to each other as shown in Panel 2. Figure from Ratnanather et al., (2003).	211
7.16.	Left column: Top two panels show automatically and hand generated PT surfaces from one MRI brain; bottom two panels show a second PT generated automatically and by hand contouring. Right column shows the surface of the left PT from the STG shown superimposed with the mean-curvature map drawn over the planar coordinates at the location defined by the bijection $\phi : S(\Delta) \rightarrow D$. The retro-insular end of the HG and the positive y-axis passes through the posterior STG where the ramus begins. Also superimposed is the Heschl's sulcus in blue generated by dynamic programming tracking on the original surface. Data taken from the laboratory of Drs. Godfrey Pearlson and Patrick Barta; reconstructions from Dr. Tilak Ratnanather (see Plate 9).	211
7.17.	Top row: Panel 1 shows the reconstruction of the cortical surface with the curvature map superimposed. Shown depicted are various sulcal principal curves generated via dynamic programming. Panel 2 shows the planar representation of the medial cortex. Bottom row: Panel 3 shows the reconstruction of the left and right MPFC from Dr. Kelly Botteron. Panels 4 and 5 show the planar maps of the MPFC reconstructions superimposed curvature profiles. Data taken from Dr. Kelly Botteron of Washington University (see Plate 10).	212
8.1.	Panel shows a curved submanifold $M \subset \mathbb{R}^3$ with tangent plane $T_p(M)$ at point $p \in M$.	215
8.2.	The basic model of transforming manifold M with local coordinates ϕ under smooth mapping $F : M \rightarrow N$ with local coordinates ψ .	219
8.3.	Left: shows the spherical template (top panel) and its deformed version (bottom panel) resulting from translation group applied to 4096 generators on the template. Right: Panels show deformations via the first four surface harmonics of the surface of the sphere (see also Plate 11).	228
8.4.	Panel 1 depicts an electron micrograph at 30,000 times magnification showing connected subregions of mitochondria. Panel 2 shows the disjoint partition of the connected subregion micrograph images into $\cup_j D(g_j)$ and the bounding contours $\partial D(g_j)$ representing the closed active contour models representing these regions. Data taken from the laboratory of Dr. Jeffrey Saffitz of Washington University.	230
8.5.	Panel shows the computation required for a small perturbation of the active model $g(\gamma) \rightarrow g(\gamma + \epsilon)$ depicts the regions over which the energy density in the integral is computed.	233
8.6.	Left panel shows the linear membrane model with midline curve f_s shown dashed, and two boundary components defined by the midline curve and normals $f_s^{+w} = f_s + wn_s$ and $f_s^{-w} = f_s - wn_s$. Right panel shows results from the linear active contour model.	236
8.7.	Results from level set evolution showing different iterations for a single face (see also Plate 12).	239
8.8.	Panels 1, 2, and 3 show the segmentation via the closed active contour models. Mitochondria interiors and cytoplasmic exteriors were represented using the Gaussian random field asymptotic partition function. Each circular template shape was manually placed; data were taken from the laboratory of Dr. Jeffrey Saffitz at Washington University.	243
9.1.	Top row: Panel 1 shows an electron-micrograph of rat-heart myocyte containing a linear organelle at 20,000 \times magnification taken from the laboratory of Jeffrey Saffitz of the department of Pathology at Washington University. Panel 2 shows the average mitochondria generated from 41 micrographs and hand tracing 497 mitochondria. Bottom row: Panel 3	

	shows the spectrum of the covariance of the stationary Gaussian prior on the deformations of the mitochondria. Panel 4 shows 8 mitochondria generated from the prior distribution placed uniformly in the scene.	260
9.2.	Figures shows deformations corresponding to the solution of the random equation, $LU = W$. Left figure: Top left panel (a) shows the original image with the area of dilatation depicted; the top right panel (b) shows a contracting field $d_1 < 0$ and $d_2 < 0$; the bottom left panel (c) shows an expanding field $d_1 > 0$ and $d_2 > 0$; the bottom right panel (d) shows a shearing field $d_1 < 0$ and $d_2 > 0$. Right figure shows analogous deformations to the ventricles (see also Plate 13).	281
9.3.	Three sections for each of three randomly generated macaque monkey brains.	283
9.4.	Panels 1–7 show the spherical harmonics 2–8 on the unit sphere computed numerically from the Laplacian operator on the sphere. Panel 8 shows the eigenvalues of the shift invariant covariance associated with the Laplacian operator (see also Plate 14).	287
9.5.	Top row shows 2D X-Z sections with pseudopod (panel 1) and high contrast data (panel 2) and low contrast data (panel 3). Bottom row (panel 4) shows 3D surface; (panel 5) shows high contrast, and (panel 6) low contrast surface reconstruction using active sphere.	292
9.6.	Left panel shows triangle element with vertices P_1, P_2, P_3 and arbitrary point within the triangle P . Right panel shows one ring centered at vertex j with triangle T_i having three vertices j, j_1, j_2 and angle θ_{ij_1} and θ_{ij_2} opposite to edges jj_1 and jj_2 , respectively.	294
9.7.	Rows 1 and 2 show surface harmonics 1–8 of the Laplace Beltrami operator on the planum temporale. Rows 3 and 4 are identical for the central sulcus. Central sulcus results are visualized via a bijection to the plane. Surface harmonics taken from Qiu and Bitouk (see also Plate 15).	297
9.8.	Figure shows planum temporal. Shown are the curvature profiles expanded in the complete orthonormal basis of the Laplace Beltrami operator (see also Plate 16).	299
9.9.	Panel 1 shows unconstrained likelihood estimator demonstrating dimensional instability. Panels 2 and 3 show the likelihood estimators with Gaussian prior. Results taken from Snyder and Miller [243].	307
9.10.	Panel 1 shows example point-spread functions from time-of-flight PET. Panel 2 (top row) show the Pie Phantom. Panel 3 (bottom row, right column) shows results of 1,000 EM iterations with no smoothing ($\alpha = 0$) of Eqn. 2.174 from Chapter 2. Results taken from Lanterman.	308
10.1.	Panel 1 shows perturbation of the flow $g_t \rightarrow g_t(\epsilon) = g_t + \epsilon \eta_t(g_t) + o(\epsilon)$ in exact matching with $g_1(\epsilon) = g_1$. For the matrix case $g_t \rightarrow g_t(\epsilon) = (\text{id} + \epsilon \eta_t)g_t + o(\epsilon) = g_t + \eta_t g_t + o(\epsilon)$. Panel 2 shows the variation of the flow g_t in inexact matching with $g_1(\epsilon)$ free, i.e. $g_1(\epsilon) \neq g_1$.	325
11.1.	Figure shows the Lagrangian description of the flow depicting the ODE $(dg_t/dt) = v_t(g_t), g_0 = \text{id}$.	332
11.2.	Top row: Panels show the vector fields $v_t(\cdot), t_1 = 0, t_2, t_3, t_4, t_5, t_6$ depicting the flow of a sequence of diffeomorphisms satisfying $\partial g_t / \partial t = v_t(g_t), g_0 = \text{id}$. Middle row: Panels depict the flow of diffeomorphisms g_t^{-1} associated with the vector fields in the top row. Bottom row: Panels depict the flow of particles which carry black and white labels corresponding to the circular movement of the patch under the flow of diffeomorphisms $I \circ g_t^{-1}$.	333
11.3.	Panels show the diffeomorphisms applied to the regular coordinates generated from the landmark metric matching; $x_n = \circ, x'_n = *$. Figure shows the scale (column 1), S-curve (column 2), and circular rotation (column 3) for $\sigma = 0.3$ (top row) and $\sigma = 1.0$ (bottom row). The line emanating from the landmarks are the trajectory onto the corresponding landmark.	343
11.4.	Results of landmark mapping of surface geometries. Top row shows the face mapping; bottom row shows the hippocampus mapping. Column 1 shows the starting template; column 6 shows the 3 target surfaces. Columns 2–5 show examples along the geodesic generated by solving the diffeomorphic mapping of the faces. Results taken from Vaillant [284].	344
11.5.	Histograms of accuracy showing distance between template and target surface geometries after landmark mapping. Results taken from Vaillant, PhD.	344

- 12.1. Row 1 shows objects for translation via diagonal translation (panels 1,2), scale (panel 3,4), and mitochondria (panels 5,6). Row 2 shows the density of momentum at the identity $Av_0 = (I'(g_1) - I)|Dg_1| \|\nabla I\|$ selected to match object 1 to object 2 constructed via the Faisal Beg large deformation diffeomorphic metric mapping algorithm (see Chapter 16) to satisfy Eqn. 12.38. Row 3 shows the image transported $I \circ g_1^{-1}$ by integrating the vector fields from the momentum at the identity along the flow $\dot{g}_t = v_t(g_t), g_0 = \text{id}$. 353
- 12.2. Experiments comparing momentum for translation (column 1), scale (column 2), and mitochondria (column 3). Row 1 shows comparison between momentum at the identity Av_0 generated by the Beg large deformation diffeomorphic metric mapping algorithm (see Chapter 16) and the gradient of the image ∇I ; arrows depict the direction vector of each. Row 2 shows the vector fields v_0 at the identity corresponding to the momentum Av_0 . 354
- 12.3. Top row: Panel 1 shows the grid test pattern matching $A \rightarrow B, C \rightarrow D$ with the corners fixed. Panel 2 shows the movement of the grid under the diffeomorphism g_1 ; panel 3 shows the determinant of the Jacobian $|Dg_1|$. Bottom row: Panel 4 shows the trajectories of the particles $g_t(x_i), i = 1, 2, t \in [0, 1]$ traced out by the landmark points A,C and the four corners of the image projected into the plane. Panel 5 shows the small deformation mapping applied to the grid; panel 6 shows the determinant of the Jacobian $|Dg_1|$. Black to white color scale means large negative to large positive Jacobian. The variances were $\sigma^2 = 0.01$; mappings from Joshi [282] (see also Plate 17). 364
- 12.4. Figure shows results of large deformation diffeomorphic landmark matching. Row 1: Panels 1 and 2 show brains 87A and target 90C, panel 3 shows 87A matched to 90C transformed via landmarks. Row 2: Panels 5, 6, 7, show sections through 87A, 90C, and 87A matched to 90C, respectively. Panel 4 shows the difference image between 87A and 90C; panel 8 shows difference image after landmark transformation. Mapping results taken from Johis [282]; data taken from David Van Essen of Washington University (see also Plate 18). 364
- 12.5. The source of images are the pair of photometric intensities and geometric transformation $(I, g) \in \mathcal{A} = \mathcal{H} \times \mathcal{G}$. The identification generate the images $I \circ g^{-1}, g \in \mathcal{G}$. 365
- 12.6. Shown is a comparison of objects I and J by looking for the smallest distance within the set $\mathcal{G}(\text{id}, I)$ and $\mathcal{G}(\text{id}, J)$. 366
- 12.7. Shows two rows from Younes [286] depicting exclusively photometric variation. Top row shows photometric change during matching. Bottom row shows grid undergoing no deformation. 371
- 12.8. Shows all grid deformation from Younes [286]. Top row shows photometric change required to generate matching. Middle row shows photometric change put through geometric variation. Bottom row shows deformation to the grid. 372
- 12.9. Rows show the time series of geometric and intensity transformation in brain tissue used for tumor growth. Row 1 shows the insertion of photometric change depicted by the small black dot. Row 2 shows the tumor developing under mostly geometry deformation. Row 3 shows the geometric deformation of the grid. 372
- 12.10. Top row shows photometric and high-dimensional geometric motion for the glasses experiment. Bottom rows show the eye opening video. Results taken from Younes showing image flow $J_{t_k}(x)$. 373
- 13.1. The source of images: An orbit under matrix groups of transformations. The observable I^D from multiple sensors. 378
- 13.2. Panels show the projective transformations in noise. Cubes depict objects in 3D $I_\alpha \circ g$; columns show projection mean $TI_\alpha(g)$ (bottom panel) and projection in Gaussian noise $I^D = TI_\alpha(g) + W$. 380
- 13.3. Top row: Panel 1 shows the true rendered object $TI(g)$. Panels 2–4 show different choices of the object at different poses and different object types. Bottom row: Panel 5 shows the observed Gaussian random field $I^D = TI(g) + W$ for the true object. Panels 6–8 show the difference between the optical data synthesized according to the Gaussian model from the target at its estimated pose, subtracted from the true measured data. Panel 6 shows the

- correct object at correct pose. Panels 7,8 show mismatched pose and object type, respectively. 381
- 13.4. Panel 1 shows the Hilbert Schmidt distance for $\mathbf{SO}(2)$; panel 2 shows the same for $\mathbf{SO}(3)$. Panel 3 shows tiling of $\mathbf{SO}(3)$ for discrete integration. Each point on S_+^2 is where a rotation vector passes. 383
- 13.5. Top row: Panels 1–5 show CAD models within one error ball of panel 3 depicted by the middle X at 1.0 unit HS MSE. Bottom row: Panel 6 shows $I^D = TI(g) + W$ for noise level $\sigma = 0.4$ corresponding to the middle X in panel 7 for HS performance of MSE=1. Panel 7 shows the mean-squared error bound as measured by the HSB on estimating the orientation using VIDEO projective imagery as a function of noise level. 386
- 13.6. Panel 1 shows the CAD airframe with the inertial coordinates with linear and angular velocities. Panel 2 shows MSE as measured by the HSB versus noise in the case of three different motion priors: (i) no prior (broken line), (ii) motion prior, and (iii) strong prior (.-.-)(see also Plate 19). 387
- 13.7. Panel 1 shows VIDEO tanks as the camera closes; panel 2 shows the HSB bound for mean-squared error performance of position; panel 3 shows the orientation bound. The different curves correspond to successive estimates of position and orientation as a function of scale as the camera closes. The different lines show the performance curves as a function of number of images used for the estimator. Results from Srivastava [307]. 388
- 13.8. Left Half: Publicly available SAR data from the MSTAR program showing vehicles (rows 1,3) and SAR datasets (row 2,4). Right Half: Shows estimated variance for each pixel for the 72 azimuth angles spaced from 5° to 360° of a T72. Variances were estimated from training data in the MSTAR database (see also Plate 20). 389
- 13.9. Panel 1 shows relative histogram of azimuth estimation errors converted from squared Hilbert–Schmidt norm to an equivalent error in units of degrees. Panel 2 shows average pose estimation error as a function of true vehicle azimuth. Panel 3 shows average classification error rate as a function of true vehicle azimuth. Results from the PhD thesis work of Devore and O’Sullivan [309] (see also Plate 21). 391
- 13.10. Confusion matrix for the 10-class recognition problem. Row headings name true vehicle classes and the value in any column is the number of images classified as the vehicle named in the column heading. Results from O’Sullivan et al. [308] 391
- 13.11. Top row: Panel 1 shows wireframes of targets viewed under perspective projection and obscuration. Panel 2 shows noise-free LADAR range image with range ambiguity. Panel 3 shows sample LADAR range image with range ambiguity, anomalous pixels, and range-dependent measurement errors. Bottom row: Panels show results from the LADAR experiments. Panel 4 and 5 show the mean LADAR signal and with noise. Panel 6 shows the HSB mean-squared error performance for pose as a function of CNR (see also Plate 22). 392
- 13.12. Left Half: Panel 1 shows M2 tanks through perspective projection; panel 2 shows the scene in Poisson noise. Panel 3 shows scene with Gaussian blur; panel 4 shows Poisson noise. Right Half: Panel 5 shows the log-likelihood of the data on the pixel array for the two matches $\log p(I^D(y)|I(g))$. Brightness means higher relative log-probability. 394
- 13.13. FLIR: Panels show results from the FLIR experiments. Panel 1 shows the mean, and panel 2 shows the signal in noise. Panel 3 shows the HSB mean-squared error performance as a function of SNR (see also Plate 23). 395
- 13.14. Azimuth-elevation power spectrum of the X-29 at several poses. Panels 1,3 show renderings of the object; panels 2,4 show high-resolution radar profiles generated from XPATCH for the X-29 targets at their respective orientations. 396
- 13.15. Panels 1 and 2 show the HRR range profile mean-fields and samples with noise for a T1 tank at 60° . Panel 3 shows the HSB mean-squared error performance as a function of SNR (see text for definition). 397
- 13.16. Top row: Panels 1, 2 show the template section I_α MRI image, and the image after rotation and translation with noise $I^D = I_\alpha \circ g + W$. Bottom row: Panels 3, 4 show the HSB for rotation, and rotation plus translation $\mathbf{SE}(2)$ error. 398

- 13.17. Panel 1 shows the T2 weighted MR image; panel 2 shows the image with a simulated tumor. Panel 3 shows the image with simulated inhomogeneity. Panel 4 shows the HSB bounds for rotation error for the template in panel 1 matched to the tumor and inhomogeneity data. 399
- 13.18. Panels 1,2 show tanks at various brightnesses and orientations. Panel 3 shows $1/2HSB(\theta, \sigma)$ for a set of θ (x -axis) versus $\sigma^2/\|\partial_\theta TI(\theta)\|^2$ (y -axis) (see also plate 24). 401
- 13.19. Top and middle rows show images from 1000 and 250 pixels on target as a function of $\sigma = 0.1, 0.3, 0.5$ (columns 1, 2, 3), respectively. Bottom row shows pose performance as measured by $HSB/2$ versus noise $\sigma^2/\|TI(\theta)\|$ showing the three curves for 100, 500, 1000 pixels on target. The solid line shows the $\frac{1}{2}HSB$; the dashed line shows the MSE as given by $E|\theta - \hat{\theta}|^2$. Note the near linear relation. 402
- 13.20. Columns 1, 2 and 3 show imagery of the vehicle (top row) and brain (bottom row) as a function of increasing SNR $SNR = -10, 0, 10$ dB. Column 4: panels show $\frac{1}{2}$ HSB versus inverse Fisher information F^{-1} ; dots show the MSE. 403
- 13.21. HRR,FLIR,LADAR FUSION: Panels show the performance-bound curves for 9° root-mean-square pose estimation error for LADAR, HRR (panel 1), LADAR,FLIR (panel 2), FLIR,HRR (panel 3). 404
- 13.22. Panels 1, 2 show two different object orientations depicting near symmetry visualized through the perspective projection of video imagery $TI(g_1) \approx TI(g_2)$. Panel 3 shows the posterior density of the Gaussian projection model with a uniform prior on the orientation. Note how they are virtually identical. 407
- 13.23. Top row: Panels 1,2 show VIDEO images of a tank at noise levels $\sigma/\|TI(g)\| = 0.01, 0.1$; panel 3 shows curves denoting the log-probability of misidentifying the tank as a truck for a fixed object pose. Solid curves show asymptotic analytical estimates (solid line); dashed curves show the likelihood ratio test calculation computed via numerical integration. Bottom row: Panels show probability of false alarm comparing Monte-Carlo simulation and asymptotic approximation from Shapiro and Yen [327]. Panel 4 shows FLIR identification between two targets comparing Monte-Carlo (circles) and asymptotic approximation (solid). Panel 5 shows similar results for M-ary recognition. 412
- 14.1. The source of images: an orbit under groups of transformations. The source are images and group elements $(I, g), I \in \mathcal{H}, g \in \mathcal{G}$, \mathcal{H} the Hilbert space of photometric intensities, and \mathcal{G} the group of geometric transformations. The observations are $I^D \in \mathcal{I}^D$. 414
- 14.2. Shows photometric variability of FLIR imagery. 415
- 14.3. Examples of faceted models used for studying photometric signature variation; shown depicted are triangulated graphs $X = S(\Delta) \subset \mathbb{R}^3$ which represent the surfaces on which we will want to understand photometric variation along with regular subvolumes. 416
- 14.4. Column 1: Panels show eigen-images 2 and 3. Columns 2–6 show images from the illumination variability $\mathcal{I} = \sum_n I_n \phi_n(\cdot)$ corresponding to one (top row) and two (bottom row) light sources from [329]. 419
- 14.5. Summaries of the contents of the empirical databases used for the KL expansion of the signature random field in the video context. The number of signatures used in the PCA is 2592 for the teacup, and 7776 for the teapot. Right panel shows the normalized power spectrum of the eigenbasis. 420
- 14.6. Top row shows visualizations of the first five eigenfunctions of the teacup database. Bottom row panels show the first five eigenfunctions of the teapot database. 420
- 14.7. Summaries of the contents of the empirical databases used. The static database isolates meteorological variability while the dynamic database isolates operational variability. The composite database is composed of the two modes of variability. The Right panel shows the normalized power spectrum of the eigenbasis. 421
- 14.8. Top row shows visualizations of the three eigenfunctions of the static database of the tank for FLIR photometric variation. Left column: Top 3 panels show first 3 static database eigenfunctions; bottom three panels show the dynamic database eigenfunctions. Taken from Cooper [332–334]. Right column: Figure shows *Eigentanks* from Lanterman. The rows show

- varying orientation; the columns show increasing detail as generated by increasing numbers of eigenfunctions. Shown are estimated eigensignatures for an M60 at different hypothesized orientations and with different numbers of eigentanks. From left to right, columns display signatures with the first 1, 3, 9, and 50 eigentanks used in the expansion (see also Plate 25). 421
- 14.9. Left column: Panel 1 shows a FLIR image of an M60 facing away from the detector provided by NVESD. Middle column: Panel 2 shows the log-likelihood varying with respect to orientation for the NVESD image with an M60 at 270° . The correct target type, an M60, is assumed. Lines correspond to log-likelihoods computed using 1, 3, 9, and 50 eigentanks. Panel 3 shows the log-determinant of the Fisher information penalty, with respect to the number of eigentanks for the M60 oriented at 270° (top line), and 0° (bottom line). Right column: Panel 4 shows the log-likelihood (top line) and penalized log-likelihood (bottom line) with respect to the number of eigentanks used for a tank hypothesized at the correct orientation. Panel 5 shows the MAP estimator of the M60 data (panel 1). 425
- 14.10. Left panels show synthetic T62 superimposed over a real infrared background (courtesy Night Vision and electronic Sensors Directorate). Right column shows log-likelihood varying with respect to orientation for data generated with a true orientation of 45 degrees. From bottom to top, lines correspond to likelihoods computed using 1, 3, 5 and 17 eigentanks. Bottom panel shows the log-likelihood (top line) and penalized log-likelihood (bottom line) with respect to the number of eigentanks employed. 426
- 14.11. Panels 1–4 show four signatures for the teapot $I = 1044, 336, 2000, 1$. Panel 5 shows performance degradation due to lighting signature mismatch. The true signature was $I = 1044$, other signatures studied are $I = 1, 2000$, along with Bayes integration over all signatures denoted with X's and labelled "Random". 428
- 14.12. Panels 1–4 show T62 tanks with four temperature signatures $I = 8, 45, 75, 140$; panel 5 shows performance degradation due to signature mismatch. The true signature was $I = 140$, other signatures studied are $I = 45, 75$ along with Bayes integration over all temperature signatures depicted via X's and denoted "Random" (see also Plate 26). 428
- 14.13. Panel 1 shows cryosection imagery depicting clearly the regions WM=white matter, GM=gray matter, CSF=cerebrospinal fluid, and bone. Panel 2 depicts the regions of the compartments, a disjoint partition of the full brain. Data taken from the Visible Human Project of the National Library of Medicine. Panel 3 of tissue voxel intensity values for MRI T1, MRI T2, MRI Proton Density (PD), and CT imaging modalities. Values correspond to four disjoint compartments White Matter, Gray Matter, CSF, and skull. 429
- 14.14. Column 1 shows examples of segmentations of the T1 Weighted MR Image. Columns 2 and 3 show T2, and proton density images (top row) with multimodality registration bounds for mutual information and the Hilbert–Schmidt estimator. Panel 5 shows the comparison between mutual information and the HSE for T1 and T2 weighted MR images. Panel 6 shows the T2 and proton density registration. Data generated at the Kennedy Krieger institute. 430
- 14.15. Panel 1 shows a scene constructed from CAD models of trees placed via the transported generator clutter model. Panel 2 shows the same for 3D CAD models which were ray-traced; taken from Bitouk (see also Plate 27). 432
- 14.16. Left column: Figure shows the various tree generators. Right column: Top panel 2 shows kurtosis measured for different trees. Bottom panel shows average kurtosis plot as a function of density parameter λ . 436
- 14.17. Panels 1 and 2 show examples from Bitouk of synthesized target chips generated via the ray tracing algorithm. Panel 3 shows the derivative statistic $\delta \ln I$ for synthetic clutter images; solid curve is the observed, dashed curve is the best fit with the generalized Laplace distribution $(\alpha, s) = (0.637, 0.254)$ (see also Plate 28). 437
- 14.18. Orbits under actions of the groups \mathcal{G} , \mathcal{H} and $\mathcal{A} = \mathcal{G} \otimes \mathcal{H}$. 438
- 14.19. Figure shows multiple CAD models and metric distances between within class and across class models in clutter. The diagonal entries show the within class distance in clutter; smaller numbers mean increased similarity providing a method for clutter independent identification. These distances are computed by solving the minimum over all rotations. 440

- 14.20. Panel 1 shows the tanks in dense foliage corresponding to obscuring clutter. Panels show image flow along the geodesic under the Euclidean norm $\| \cdot \|^2$ for $J(x, t_k)$ at $t_k = k/5, k = 0, 1, \dots, 5$. Results taken from Bitouk. 443
- 14.21. Top row : observed images. Bottom row : images synthesized using the second-order model. Images taken from Bitouk. 444
- 14.22. Top row: ROC curves for detection of a T72 tank for different densities of clutter. Solid line—covariance norm; dashed line—Euclidean norm. Bottom row: ROC curves for classification of a T72 tank versus a Jeep for different densities of clutter. Solid line—covariance norm; dashed line—Euclidean norm. Results from Bitouk. 446
- 15.1. Panel 1 shows the conditional entropies $H(O|I_{FLIR}^D)$ (—), $H(O|I_{VIDEO}^D)$ (- - -), $H(O|I_{FLIR}^D, I_{VIDEO}^D)$ (x) are plotted versus increasing SNR. Panel 2 the mutual informations $\Delta H(O; I_{FLIR}^D)$ (—), $\Delta H(O; I_{VIDEO}^D)$ (- - -), $\Delta H(O; I_{FLIR}^D, I_{VIDEO}^D)$ (x) are plotted versus increasing SNR. Panel 3 shows Bits expressed in log base 2. 449
- 15.2. Top row: Panel 1 shows the mutual information for the FLIR sensor $\Delta H(O; I^D, I_{temp})$ (—) and $\Delta H(O; I^D)$ (- - -), showing the information loss due to the absence of knowledge of the object signature. Panel 2 shows the average information loss due to signature variation, as measured by the Kullback–Leibler distance, $D(p(O|I^D, I_{temp}) \| p(O|I^D))$, versus increasing signal to noise ratio. Bottom row: Panels 3, 4 measure information loss between the marginalized model and the true signature model $I_{temp} = 8$ (panel 3) and $I_{temp} = 75$ (panel 4). Each panel shows the information loss taken by choosing models with the incorrect signature as indicated by the dashed lines. The marginal model information loss is shown by the x's. In both panels, information loss is minimized by the Bayes signature model without *a priori* assumption of the object signature. Bits expressed in log base 2. 452
- 15.3. Top row shows CAD models. Bottom row shows empirical entropies for pose versus analytic entropy predicted by the variance of the Fisher information (see Plate 29). 453
- 15.4. Column 1: Video imagery of the T62 tank at SNR = 5 dB (top) and 30 dB (bottom). Column 2: Numerical (—) and analytic (- - -) approximations to $H(\Theta|i^D)$ for $\theta_T = 30^\circ$. 454
- 15.5. Top row shows FLIR observations courtesy of Dr. J. Ratches U.S. Army NVESD. Second row shows renderings of the MAP estimates of orientation and signature from the corresponding observation of the top row. Third row: Panel 7 shows the mean squared error for orientation estimation via FLIR with and without signature information. Panel 8 shows the mutual information curves for FLIR with and without signature information. In both cases, the solid lines correspond to performance with signature information, and the dashed curves without signature information. The operating SNR of the NVESD Commanche FLIR data is indicated on the curves by the solid dots. Bits expressed in log base 2 (see Plate 30). 455
- 15.6. CAD models JEEP, T62, and CUP—4 orientations of CAD models without noise through camera projection (image size 32×32 , 8 bit/pixel). 463
- 15.7. Orientations of each model at SNR = -6 dB. 463
- 15.8. Left column shows rate-distortion functions for MSE distortion for objects “JEEP”(top row) and T62”(bottom row). Right column shows the same for HS distortion. 464
- 15.9. Rate-Distortion function for the model “CUP”. Left panel MSE; right panel HSE. 464
- 15.10. Output distribution densities for HSE distortion for object “JEEP” at rate 0.5 (columns 1,2), 1.0 (columns 3,4), 2.0 (columns 5,6) bits. 465
- 15.11. Output distribution density for object “T62” at rates 0.5 (columns 1,2), 1.0 (columns 3,4), 2.0 (columns 5,6) bits/orientation. 466
- 15.12. Output distribution densities for the HSE distortion measure for the CAD object “CUP” at rates 0.5 (columns 1,2), 1.0 (columns 3,4), 2.0 (columns 5,6) bits/orientation. 466
- 16.1. Figure shows the diffeomorphisms between brain volumes and submanifolds within the brain; figure taken from Van Essen et al. [167]. 470
- 16.2. The source for the **anatomical model** is the set of all images \mathcal{I} of the exemplars $I_\alpha(g_1^{-1})$. The observed data I^D are observations seen via the sensors (MRI, CT, CRYOSECTION pictures) and are conditional Gaussian random fields with mean fields $g_1 \cdot I_\alpha$. For the **static metric mapping model** there is only one observation I^D the mean field of the exemplar $I_\alpha(g_1^{-1})$. 471

- 16.3. The source for the **dynamic growth model** is the set of all growth flows $\mathcal{I}[0, 1]$ of the exemplars $g_t \cdot I_\alpha, t \in [0, 1]$. The observed data $I_t^D, t \in [0, 1]$ are flows of observations seen via the sensors (MRI, CT, CRYOSECTION pictures) and are conditional Gaussian random fields with mean fields $g_t \cdot I_\alpha, t \in [0, 1]$. 471
- 16.4. Rows 1,2, and 3 show the sequence of metric maps $I(g_t^{-1})$ for t_0, t_6, t_{12}, t_{19} , target for the TRANSLATE, SCALE, and C experiments with metric $\int_0^t \|v_s\|_V ds$. Column 6 shows the target through the forward map $I^D(g_1)$. Rows 4 and 5 show the HEART and MACAQUE. Mappings taken from Beg [285]; heart data taken from Dr. Raimond Winslow of Johns Hopkins University and macaque taken from Dr. David Van Essen of Washington University. 477
- 16.5. Columns 1,2, and 3 show the vector fields v_0 (row 1), v_{10} (row 2), and v_{19} (row 3) for TRANSLATION, SCALE, and C experiments. Columns 4 and 5 show the vector fields for the HEART experiment and MACAQUE for v_0, v_{10}, v_{19} . The mapping results are from Beg [285]; the heart data is courtesy of Dr. Raimond Winslow, The Johns Hopkins University. The macaque data is taken from David Van Essen. 477
- 16.6. Figure shows vector fields for the $\frac{1}{2}C$ (panels 1,2, and 3) and macaque (panels 4,5, and 6) experiments. 477
- 16.7. Rows 1 and 2 show results from the hippocampus mapping experiment for the YOUNG to SCHIZOPHRENIC (row 1) and ALZHEIMER'S (row 2). Shown are the sequence of geodesic mappings $I(g_t^{-1})$ connecting the Young to the targets for $t_0, t_3, t_6, t_9, t_{12}, t_{15}, t_{19}$. Plotted below each is the metric distance. Column 6 shows the vector fields v_0 at the identity. Mapping results from Beg [285]. Data taken from the laboratory of Dr. John Csernansky of Washington University. 478
- 16.8. Figure shows 3D hippocampus mapping results from Schizophrenia (row 1) and Alzheimer's (row 2). Top row shows template I_α , panel 2 shows the Schizophrenic hippocampus hand labeled I_{SCHIZ} , panel 3 shows $I_\alpha(g_1^{-1}(\cdot))$, and panel 4 shows $I_{SCHIZ}(g_1(\cdot))$. Row 2 shows similar results for the Alzheimer's. Data taken from the laboratory of Dr. John Csernansky of Washington University. 478
- 16.9. Column 1 shows electron microscopy images. Columns 2–5 show computed geodesic distances between the template shapes (column 2) and other mitochondrial shapes (columns 3,4, and 5). 479
- 16.10. Top row shows vector fields generated from the metric correspondence algorithm; bottom row shows vector fields from the Christensen algorithm. Column 1 shows translation, column 2 shows the heart, and column 3 shows the macaque sections. Shown in each panel are the superposition of 20 vector fields along the flow. 480
- 16.11. Figure shows 3D MRI-MPRAGE head matching. Left three panels show template and target, and the 3D transformation of the template matched to the target. Right three panels show the same for the second individual. Figure reproduced from Christensen et al., Volumetric Transformation of Brain Anatomy, 1997. 481
- 16.12. Column 1: Panel shows the template hippocampus (green) embedded in the MRI volume. Columns 2,3,4: Panels show the template surface (top), the target surface (middle), and the mapped template surface (bottom). Top row shows the hippocampus surface through the mapping; bottom row shows the surface embedded in the mapped volumes. Mapping results from Christensen et al. [220] and Haller et al. [162] (see Plate 31). 482
- 16.13. Parasagittal cryosection cortex sections from the macaque occipital lobe. Left column: Panels show photographs of the right hemisphere of macaque 92K (top) and 92L (bottom). Arrows show the cuts that were made to remove part of the visual cortex. Bottom rows: Right columns show automated segmentations of sections 23, 37, 52, 61 showing 92K, 92L, and 92K \rightarrow 92L. Columns 2 and 3 show hand segmentations; column 4 shows automated segmentation from mapping 92K \rightarrow 92L. Mapping results taken from Christensen et al. [220]; data taken from the laboratory of David Van Essen of Washington University. 482
- 16.14. Van Essen atlas of macaque visual cortex. Left column shows the macaque atlas (top row) and flattened version (bottom row) generated by David Van Essen. Right column shows flattened surface-based mapping from an individual flat map via landmark matching in the plane to the macaque atlas. **A** shows flat map of Case 1, **B** pattern of grid lines after deformation of Case 1, **C** shows the displacement vector field, **D** shows the map of

- geography (shading) and target registration contours (black lines) on the atlas map, and E shows the deformed borders of Case 1 on the atlas map. Results taken from Van Essen et al. [362] (See Plate 32). 484
- 16.15. Panel 1 shows 3D Visible Human Male. Panel 2 shows the boundaries of deformed macaque visual areas (black lines) superimposed on the fMRI activation pattern from an attentional task from the study of Corbetta et al. [447] after deformation to the Visible Man atlas by Drury et al. [448] and mapping results from [362]. Panels 3 and 4 show landmarks on macaque flat map and human flat map, respectively, used for performing the mappings shown in panel 2 (see Plate 33). 485
- 16.16. Columns 1 and 2 show hearts in grayscale MRI in normal (column 1) and failing (column 2) hearts; columns 3 and 4 show DTI of those sections depicting orientation of the principal eigenvectors by color. The top row shows coronal sections, the bottom row shows axial sections (see Plate 34). 485
- 16.17. Figure shows the template heart and the target heart images after transformation into the coordinate system of the template. The determinant of the Jacobian of the transformation is superimposed as a colormap on the surface rendering. Blue colors indicate regions where the determinant is less than unity whereas red regions are where the determinant is greater than unity. Data courtesy of Dr. Raimond Winslow; mappings courtesy of Faisal Beg and Pat Helm (see Plate 35). 486
- 16.18. Panels show points along the geodesic for the LDDMM of the DT vector data, with the red vectors showing the template, and the blue vectors showing the target. Taken from [449]. 487
- 16.19. Column 1 shows the geometries of the two hearts. Column 2, top panel shows the two geometries superimposed after rigid motion; blue is template, red is target. Column 2, bottom panel shows the LDDMM solution color representing the dot product of the corresponding vectors after alignment. Red color means total alignment; taken from [449] (see Plate 36). 488
- 16.20. MR images of postnatal mouse brains. Coronal MR T_2 and DT images were shown at the level of anterior commissure. Images have been aligned to ensure proper orientation and position. Intensity in MR T_2 images is related to tissue properties, such as the content of myelin protein. MR diffusion tensor images reveal local tissue orientation. Both types of images were utilized in our mouse brain developmental study. The color scheme for diffusion tensor images was illustrated by color arrows in the figure, with red for local tissue whose orientation is anterior-to-posterior (perpendicular to the current plane), green for horizontal orientation, and blue for vertical orientation. The scalar bar represents 1 mm. Structure abbreviations are—ac: anterior commissure; cc: corpus callosum; CPu: caudate putamen; CX: cortex (see Plate 37). 490
- 16.21. Panel 1 shows white matter tracts, anterior part of anterior commissure (aca), posterior (acp), cerebral peduncle (cb), hippocampus (H), and hippocampal commissure (hc). Panels 2,3, and 4 show visualization of landmarks on white matter tracts (green dots) and landmarks inside hippocampus (yellow structures) (blue dots), outer brain surface (orange) with boundary of white matter and cortical region shown in green, and landmarks depicted as intersections of white lines on the green structure (see Plate 38). 490
- 16.22. LDDMM of P7 mouse brain images to P30 images. Coronal MR images of P7 (A) and P30 (C) were shown with brain boundaries marked by orange curves. Deformed P7 (B) was overlaid with the brain boundary of P30. White arrows in A, B, and C are pointing to hippocampus (H). The transformation was visualized as color-coded vector plot (D) and Jacobian plot (F), with enlarged local areas shown in E and G, respectively. For vector plots, color of vectors represents the magnitude of local displacement. For Jacobian plot, color on deformed grids represents changes in local volume. Surface of P7, deformed P7, and P30 mouse brains were visualized in H, from top to bottom. Vector maps show how transformations deform a local voxel from its original location by taking the original location as the start of the vector and the destination as the end; length and color of vectors reflect the magnitude of tissue displacement. The Jacobian map shows local volume changes; a value greater than unity corresponds to volume increase, a value less than unity corresponds to volume loss, and unity for no volume change (see Plate 39). 491

- 16.23. Outer-shape changes of mouse hippocampus measured using LDDMM. The displacement due to growth of hippocampus was color coded and the direction of growth was visualized by 3D glyphs (see plate 40). 492
- 16.24. Series of volumes depicting aging in Huntington’s disease patient over time as manifest in the Caudate. Data taken from the laboratory of Dr. Elizabeth Aylward of the University of Washington (see Plate 41). 492
- 17.1. 2D Face Cartoon landmark LDDMM (see also Plate 42). 500
- 17.2. Non-linearity artifact in panel 4 overcome by initial momentum approach shown in panel 5. Results taken from Vaillant [284] (see also Plate 43). 500
- 17.3. First two eigenmodes of 3D PCA applied to the Morphable Faces database of 100 faces. Row 1 shows front views; row 2 shows side views. Column 1 shows template. Columns 2 and 3 show deformation via eigenfunction 1. Columns 4 and 5 show deformation via eigenfunction 2 (see also Plate 44). 501
- 17.4. First two eigenmodes of 3D PCA applied to hippocampi from Randy Buckner of Washington University. Column 1 shows front and side view. Top row shows deformation via eigenfunction 1. Bottom row shows deformation via eigenfunction 2 (see also Plate 45). 501
- 17.5. Left column panel 1 shows a whole MRI-MPRAGE image volume with a section through the brain delineating the surface of the hippocampus. Left column panel 2 shows the triangulated graph representing the mean state of the hippocampus. The data are courtesy of Dr. John Csernansky of the Department of Psychiatry at Washington University. Right column four panels show the first four surface harmonics visualized through deformation of the template (panel 2). Courtesy Sarang Joshi, Ph.D. thesis [242] (see also Plate 46). 503
- 17.6. Top row: Panels 1–3 show maps of the initial hippocampus to three in a population of 30 patients, $M_0 \circ g_1, M_0 \circ g_2, M_0 \circ g_3$ studied in Csernansky et al. [162]. Panel 4 shows the composite template generated from the average of 30 maps, $M_0 \circ \bar{g}$. Bottom left column shows the first two eigenshapes of the left and right hippocampus generated from a population of maps of normals and schizophrenics. Lower right panels show random instances of left and right hippocampi generated via the empirically estimated covariance function for the Gaussian random field on the hippocampus surface. Taken from Joshi [242] (see also Plate 47). 504
- 17.7. Top row: Schizophrenia. Panel 1 shows difference of hippocampal surface patterns between the control and schizophrenia groups visualized as z-scores on the mean surface of the control group. Inward surface deformations due to schizophrenia are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 2 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,3,4,6,10,15}. Multivariate ANOVA indicates significant between-group difference: $p = 0.0028$ ($F = 4.73$, $df = 1, 28$). Bottom row: Asymmetry in Schizophrenia. Panel 3 shows difference of hippocampal surface asymmetry patterns between the control and schizophrenia groups visualized as z-scores on the mean flipped right-side surface of the control group. Inward surface deformations due to differences in asymmetry are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 4 shows statistical testing of asymmetry pattern differences: log-likelihood ratios for a linear combination of basis functions {2,12,17}. Multivariate ANOVA indicates significant between-group difference in asymmetry: $p = 0.0029$ ($F = 6.03$, $df = 3, 26$) (see also Plate 48). 507
- 17.8. Top row: Alzheimer’s Disease. Panel 1 shows difference of hippocampal surface patterns between the elderly control (CDR 0) and AD (CDR 0.5) groups visualized as z-scores on the mean surface of the elderly group. Inward surface deformations due to AD are visualized in colder colors, outwards in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 2 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,5}. Multivariate ANOVA indicates significant between-group difference: $p = 0.0002$ ($F = 11.4$, $df = 2, 33$). Bottom row: Normal Aging. Panel 3 shows difference of hippocampal surface patterns between the younger and the elderly control (CDR 0) groups visualized as z-scores on the mean surface of the younger

- group. Inward surface deformations due to aging are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 4 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,2}. Multivariate ANOVA indicates significant between-group difference: $p < 0.0001$ ($F = 348$, $df = 2, 30$) (see also Plate 49). 508
- 17.9. AD Progression. Panel 1 shows deformations from baseline to follow-up for the CDR 0.5 AD group. Inward surface deformations due to AD are visualized in colder colors, outwards in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 2 shows deformations from baseline to follow-up for the CDR 0 control group. Inward surface deformations due to AD are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 3 shows follow-up-versus-baseline “spread” of the between-group inward surface deformation patterns, shown as Wilcoxon’s sign rank test map on the CDR 0 mean surface. Areas of significant ($p < 0.05$) inward deformation at baseline of CDR 0.5 group are shown in turquoise color, representing 38% of total hippocampal surface area. By follow-up, areas of significant inward deformation have increased to 47% of total hippocampal surface area. The increased affected areas are shown in purple color. Areas of non-significant surface deformation are shown in green color. Panel 4 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,2,4,11}. Multivariate ANOVA of the first 12 basis functions indicates significant between-group difference: $p = 0.014$ ($F = 2.66$, $df = 12, 31$) (See also Plate 50). 509
- 17.10. Depression. Panel shows differences of hippocampal surface patterns between the control and depression groups visualized as perpendicular displacements on the mean surface of the control group. Inward surface deformations due to depression are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Multivariate ANOVA of the first 10 basis functions indicates significant between-group difference: $p < 0.0001$ ($F = 34.1$, $df = 10, 58$) (see also Plate 51). 510
- 17.11. Empirical distribution \hat{F} from randomized Fisher’s test with 10,000 group permutations, between the group of control and schizophrenia subjects. Basis vectors {2,12,17} are selected. The $p = 0.0031$ value shown is calculated from Eqn. (17.25). Also shown are (i) $\hat{F}(T^2)$ value (solid blue line) of the control-versus-schizophrenia group comparison; (ii) theoretical F -distribution (solid red curve) with (3,26) degrees of freedom superimposed on the empirical distribution; and (iii) $p = 0.05$ (red dotted line) and $p = 0.01$ (red dot-dash line) for reference (see also Plate 52). 511
- 17.12. Figure showing heteromorphic deformations. The left panel shows an expansion or “push out”; the right panel shows a “push in” (see also Plate 53). 513
- 19.1. Top row: Panels 1 and 2 show the CAD models used for the simulations. Panel 3 shows the top-down and perspective view of the M60, M2, and T62 tank scene. Bottom row: Panels 4,5, and 6 show the radiant intensities on the CAD models generated via the PRISM simulation package. 533
- 19.2. Panel 1 shows a single track of $n(1)$ rotations and translations; panel 2 displays the projection system for optical imaging. 534
- 19.3. Top row: Panels 1 and 2 show 3–5 and 8–12 microns infra-red image of a natural scene obtained during the Grayling 1 Field Experiment, Smart Weapons Observability Enhancement Joint Test & Evaluation Program. Courtesy Robert Guenther of the Physics Division, Army Research Office. Bottom row: Panel 3 shows an ideal infra-red scene TI through the projective geometry and with Gaussian blur; panel 4 shows the corresponding ranging mean-field for the LADAR. 535
- 19.4. Columns 1 and 2 Death Process: Panel 1 shows an initial condition with extraneous target hypotheses; 2, 3, and 4 show the state of the jump diffusion process at iterations 2, 41, and 65, respectively. Columns 3 and 4 Identification Process: Column 3 shows the match of the M60 to a data from a scene containing an M60 (top) and the match of the T62 to an M60 scene (bottom). Column 4 shows the log-likelihood of the data on the pixel array for the two matches. Brightness means higher relative log-probability. Results taken from Lanterman. 544

- 19.5. Top row shows examples of optical imagery for the airplane. Middle row shows two poses of the target space; panels 5 and 7 show renderings of the target with panels 6 and 8 showing the difference between the optical data synthesized according to the Gaussian model from the target at its estimated pose, subtracted from the true measured data. Bottom row shows identification showing two different identifications. Panels 9 and 11 show renderings of the two different target types; panels 10 and 12 show the difference between the optical data generated from the proposed target at its estimated pose, subtracted from the true measured data. Right panels shows the correctly estimated identity and difference data. 545
- 19.6. Top row shows the target rendered at a sequence of four flight times, time increasing from left to right. Middle row shows the HRR data associated with the target imaged at these orientations. Bottom row shows the estimates of the pose generated at the four flight times. Results taken from O’Sullivan. 546
- 19.7. FLIR: Panel 1 shows the initial data from a configuration of M2 APC’s observed by the FLIR imager. Successive panels 2–9 show a sequence of states of the jump–diffusion process at iterations 1, 10, 12, 25, 27, 34, 55, and 75. Results taken from Lanterman. 547
- 19.8. FLIR: Panel 1 shows the measured scene for the FLIR imager. Panels 2–10 show iterations 1,3,24,32,34,68,87,88, and 117 of the jump–diffusion process. Results taken from Lanterman. 547
- 19.9. Results from LADAR: Panel 1 shows the initial data from a configuration of vehicles observed by the LADAR imager. Panels 2–9 show iterations 1, 3, 11, 12, 24, 32, 38, and 130 of the jump–diffusion process for the LADAR range data. Results taken from Lanterman. 548
- 19.10. Top row: Panel 1 shows the actual track drawn in gray with the mesh representing ground supporting the observation system in the inertial frame of reference. Panels 2 and 3 display intermediate results from the single track estimation with the estimates drawn in black. Rows 2 and 3: Sequence of jump moves adding segments to the estimated state from left to right; no diffusion. Bottom two rows: Sequence of panels showing diffusion towards the true track mean. 549
- 19.11. Top row: Panel 1 showing the track estimates drawn overlapping in black at several stages of the algorithm. Panel 2 depicts the posterior probability of a target at that position proportional to the size of the sphere drawn. Panels 3 and 4 show candidates from the prior distribution for target path estimation with the high prior probability candidates forming a cone at track end for the algorithm to sample from. 552
- 19.12. Panel 1 is the micrograph data at 30,000 times magnification with panel 2 showing a hand segmentation into the disjoint partition $U_j \mathcal{I}_{g_j}$ and with bounding contours g_j of the two region types, mitochondria (white) and background (black). Data taken from the laboratory of Dr. Jeffrey Saffitz of Washington University. 553
- 19.13. Panel 1 shows the jump detections scored via brightness indicating the amount of posterior probability. Panel 2 depicts the representation of mitochondria resulting from the jump–diffusion algorithm. Panel 3 shows a pixel-by-pixel segmentation of the image based on the optimal Bayes hypothesis test for each pixel under the two models. 559
- 19.14. Figure showing a birth of a mitochondria (left column), a death of a mitochondria (middle column), and a merge of two mitochondria (right column). 559
- 19.15. Comparison of segmentations based on 4-gray level MRF pseudolikelihood models (left two columns) and the Gaussian asymptotic partition function global Bayes model (right two columns) for two sets of mitochondria (top and bottom rows). Rows 1 and 2 show manual placement of mitochondria templates; rows 3 and 4 show automatic placement by the jump process. Columns 1 and 3 show image data; columns 2 and 4 show pixel-by-pixel segmentation generated by the likelihood ratio test (gray mitochondria, black cytoplasm). 560

LIST OF PLATES

1. Left column: The figure shows the results from the N-butyl alcohol experiment. The top panel shows the original 2D spectrum of the N-butyl alcohol data. The bottom panel shows the spectrum reconstructed from the estimates of the EM-algorithm parameters. The data are taken from [23]. Right column: The figures show the EM algorithm reconstruction of the COSM data taken from Dr. Keeling of Washington University; the data reconstructions are from [30]. The top row shows X-Z sections through the COSM amoeba data. The bottom row shows the 200th EM algorithm iteration for sections through the COSM amoeba data.
2. Panel 1 shows the EM algorithm fit of G,W,CSF and partial volume compartments to brain tissue histograms to illustrate the segmentation calibration. The top solid curve superimposes the measured histogram data almost exactly. The lower dashed lines depict each of the compartment fits taken separately. Panel 2 shows an MRI section of the cingulate; panel 3 shows the Bayes segmentation into G,W,CSF compartments of coronal sections of the cingulate gyrus; Panels 4 and 5 show the same as above (row 1) for the medial prefrontal gyrus. Data taken from the laboratories of Dr. John Csernansky and Dr. Kelly Botteron of Washington University
3. Top row: Panel 1 shows the triangulated graph for the template representing amoeba corresponding to spherical closed surfaces. Generators are $g_i = (v_1, v_2, v_3)$, elements of \mathbb{R}^9 . Panel 2 shows the triangulated for the closed surface representing the bounding closed surface of the hippocampus in the human brain. Panel 3 shows the surface representation of a macaque brain from the laboratory of David Van Essen. Bottom row: Shows a generator shown in standard position, orientation under rotation and translation transformation.
4. Top row: Panel 1 shows the triangulated graph representing the hippocampus generated by isocontouring; panel 2 shows the mean curvature map superimposed on the hippocampus surface of the human. Data taken from Haller et al. [162]. Bottom row: Panels 3 and 4 shows the same as above for macaque cortical surface reconstructed from cryosection data from the laboratory of Dr. David Van Essen. Bright areas represent areas of high positive mean curvature; dark areas represent areas of high negative mean curvature.
5. Top row: Panel 1 shows the atlas depiction of the occipital cortex. Panel 2 shows the reconstruction of the occipital cortex depicting the major sulcal and gyral principal curves including the inferior Calcarine sulcus and Lingual and Parietal gyri. Data taken from Dr. Steven Yantis of the Johns Hopkins University.
6. Left column shows the medial prefrontal cortex section from the Duvernoy atlas. Middle column top and bottom panels show the isosurface reconstruction of the prefrontal medial cortex. Sections through the two different MRI brains show the embedded surfaces. Right column, top and bottom panels, shows the medial prefrontal cortex reconstructions. Data taken from Dr. Kelly Botteron of Washington University.
7. Top row: Panel 1 shows eight geodesics generated on the neocortex by picking the start and end points manually. Panel 2 depicts geographical landmarks on the macaque cortex; labels 1, 2, 3, 4, 5, 6. Panel 3 shows a table of Riemannian distances in millimeters between the predefined points. Data taken from the laboratory of David Van Essen, Washington University. Bottom row: Figure shows optimality of dynamic programming. Panel 4 shows the Visible Human cortex extracted by David Van Essen; panel 5 shows choosing multiple terminal points for the DP solution; panels 6, 7 show the DP generation of the superior temporal sulcus jumping across the break connecting the start and end points which were manually selected.
8. Top row panel 1 shows the external view of the Superior Temporal Gyrus (STG). Panel 2 shows Heschl's gyrus and the posterior boundary of the plenum temporale (PT) defined via dynamic programming. Panel 3 shows the delineation of the STG surface into two with the PT as the blue region which is extracted. Bottom row shows the application of dynamic programming to extract the PT from STG surface. Panel 4 tracks Heschl's gyrus; panel 5 tracks the STG as far as the posterior ascending (or descending) ramus; panel 6 tracks the geodesic from the end of the STG to the retro-insular end of the Heschl's gyrus. Data taken

XXX

- from the laboratory of Drs. Godfrey Pearlson and Patrick Barta, reconstructions from Dr. Tilak Ratnanather.
9. Left column: Top two panels show automatically and hand generated PT surfaces from one MRI brain; bottom two panels show a second PT generated automatically and by hand contouring. Right column shows the surface of the left PT from the STG shown superimposed with the mean-curvature map drawn over the planar coordinates at the location defined by the bijection $\phi : S(\Delta) \rightarrow D$. The retro-insular end of the HG and the positive y-axis passes through the posterior STG where the ramus begins. Also superimposed is the Heschl's sulcus in blue generated by dynamic programming tracking on the original surface. Data taken from the laboratory of Drs. Godfrey Pearlson and Patrick Barta; reconstructions from Dr. Tilak Ratnanather.
 10. Top row: Panel 1 shows the reconstruction of the cortical surface with the curvature map superimposed. Shown depicted are various sulcal principal curves generated via dynamic programming. Panel 2 shows the planar representation of the medial cortex. Bottom row: Panel 3 shows the reconstruction of the left and right MPFC from Dr. Kelly Botteron. Panels 4 and 5 show the planar maps of the MPFC reconstructions superimposed curvature profiles. Data taken from Dr. Kelly Botteron of Washington University.
 11. Left: shows the spherical template (top panel) and its deformed version (bottom panel) resulting from translation group applied to 4096 generators on the template. Right: Panels show deformations via the first four surface harmonics of the surface of the sphere.
 12. Results from level set evolution showing different iterations 0, 1, 3, 5, 7, 10 for a single face.
 13. Figures show deformations corresponding to the solution of the random equation, $LU = W$. Left figure: Top left panel (a) shows the original image with the area of dilatation depicted; the top right panel (b) shows a contracting field $d_1 < 0$ and $d_2 < 0$; the bottom left panel (c) shows an expanding field $d_1 > 0$ and $d_2 > 0$; the bottom right panel (d) shows a shearing field $d_1 < 0$ and $d_2 > 0$. Right figure shows analogous deformations to the ventricles.
 14. Panels 1–7 show the spherical harmonics 2–8 on the unit sphere computed numerically from the Laplacian operator on the sphere. Panel 8 shows the eigenvalues of the shift invariant covariance associated with the Laplacian operator.
 15. Rows 1 and 2 show surface harmonics 1–8 of the Laplace Beltrami operator on the planum temporale. Rows 3 and 4 are identical for the central sulcus. Central sulcus results are visualized via a bijection to the plane. Surface harmonics taken from Qiu and Bitouk.
 16. Figure shows planum temporal. Shown are the curvature profiles expanded in the complete orthonormal basis of the Laplace Beltrami operator.
 17. Top row: Panel 1 shows the grid test pattern matching $A \rightarrow B, C \rightarrow D$ with the corners fixed. Panel 2 shows the movement of the grid under the diffeomorphism g_1 ; panel 3 shows the determinant of the Jacobian $|Dg_1|$. Bottom row: Panel 4 shows the trajectories of the particles $g_t(x_i), i = 1, 2, t \in [0, 1]$ traced out by the landmark points A,C and the four corners of the image projected into the plane. Panel 5 shows the small deformation mapping applied to the grid; panel 6 shows the determinant of the Jacobian $|Dg_1|$. Black to white color scale means large negative to large positive Jacobian. The variances were $\sigma^2 = 0.01$; mappings from Joshi [282].
 18. Figure shows results of large deformation diffeomorphic landmark matching. Row 1: Panels 1 and 2 show brains 87A and target 90C, panel 3 shows 87A matched to 90C transformed via landmarks. Row 2: Panels 5, 6, 7, show sections through 87A, 90C, and 87A matched to 90C, respectively. Panel 4 shows the difference image between 87A and 90C; panel 8 shows difference image after landmark transformation. Mapping results taken from Johis [282]; data taken from David Van Essen of Washington University.
 19. Panel 1 shows the CAD airframe with the inertial coordinates with linear and angular velocities. Panel 2 shows MSE as measured by the HSB versus noise in the case of three different motion priors: (i) no prior (broken line), (ii) motion prior, and (iii) strong prior (-.-).
 20. Left Half: Publicly available SAR data from the MSTAR program showing vehicles (rows 1,3) and SAR datasets (row 2,4). Right Half: Shows estimated variance for each pixel for the 72 azimuth angles spaced from 5° to 360° of a T72. Variances were estimated from training data in the MSTAR database.

21. Panel 1 shows relative histogram of azimuth estimation errors converted from squared Hilbert–Schmidt norm to an equivalent error in units of degrees. Panel 2 shows average pose estimation error as a function of true vehicle azimuth. Panel 3 shows average classification error rate as a function of true vehicle azimuth. Results from the PhD thesis work of Devore and O’Sullivan [309].
22. Top row: Panel 1 shows wireframes of targets viewed under perspective projection and obscuration. Panel 2 shows noise-free LADAR range image with range ambiguity. Panel 3 shows sample LADAR range image with range ambiguity, anomalous pixels, and range-dependent measurement errors. Bottom row: Panels show results from the LADAR experiments. Panel 4 and 5 show the mean LADAR signal and with noise. Panel 6 shows the HSB mean-squared error performance for pose as a function of CNR.
23. FLIR: Panels show results from the FLIR experiments. Panel 1 shows the mean, and panel 2 shows the signal in noise. Panel 3 shows the HSB mean-squared error performance as a function of SNR.
24. Panels 1,2 show tanks at various brightnesses and orientations. Panel 3 shows $1/2HSB(\theta, \sigma)$ for a set of θ (x -axis) versus $\sigma^2 / \|\partial_\theta TI(\theta)\|^2$ (y -axis).
25. Top row shows visualizations of the three eigenfunctions of the static database of the tank for FLIR photometric variation. Left column: Top 3 panels show first 3 static database eigenfunctions; bottom three panels show the dynamic database eigenfunctions. Taken from Cooper [332–334]. Right column: Figure shows *Eigentanks* from Lanterman. The rows show varying orientation; the columns show increasing detail as generated by increasing numbers of eigenfunctions. Shown are estimated eigensignatures for an M60 at different hypothesized orientations and with different numbers of eigentanks. From left to right, columns display signatures with the first 1, 3, 9, and 50 eigentanks used in the expansion.
26. Panels 1–4 show T62 tanks with four temperature signatures $I = 8, 45, 75, 140$; panel 5 shows performance degradation due to signature mismatch. The true signature was $I = 140$, other signatures studied are $I = 45, 75$ along with Bayes integration over all temperature signatures depicted via X ’s and denoted “Random”.
27. Panel 1 shows a scene constructed from CAD models of trees placed via the transported generator clutter model. Panel 2 shows the same for 3D CAD models which were ray-traced; taken from Bitouk.
28. Panels 1 and 2 show examples from Bitouk of synthesized target chips generated via the ray tracing algorithm. Panel 3 shows the derivative statistic $\delta \ln I$ for synthetic clutter images; solid curve is the observed, dashed curves is the best fit with the generalized Laplace distribution $(\alpha, s) = (0.637, 0.254)$.
29. Top row shows CAD models. Bottom row shows empirical entropies for pose versus analytic entropy predicted by the variance of the Fisher information.
30. Top row shows FLIR observations courtesy of Dr. J. Ratches U.S. Army NVESD. Second row shows renderings of the MAP estimates of orientation and signature from the corresponding observation of the top row. Third row: Panel 7 shows the mean squared error for orientation estimation via FLIR with and without signature information. Panel 8 shows the mutual information curves for FLIR with and without signature information. In both cases, the solid lines correspond to performance with signature information, and the dashed curves without signature information. The operating SNR of the NVESD Commanche FLIR data is indicated on the curves by the solid dots. Bits expressed in log base 2.
31. Column 1: Panel shows the template hippocampus (green) embedded in the MRI volume. Columns 2,3,4: Panels show the template surface (top), the target surface (middle), and the mapped template surface (bottom). Top row shows the hippocampus surface through the mapping; bottom row shows the surface embedded in the mapped volumes. Mapping results from Christensen et al. [220] and Haller et al. [162].

32. Van Essen atlas of macaque visual cortex. Left column shows the macaque atlas (top row) and flattened version (bottom row) generated by David Van Essen. Right column shows flattened surface-based mapping from an individual flat map via landmark matching in the plane to the macaque atlas. **A** shows flat map of Case 1, **B** pattern of grid lines after deformation of Case 1, **C** shows the displacement vector field, **D** shows the map of geography (shading) and target registration contours (black lines) on the atlas map, and **E** shows the deformed borders of Case 1 on the atlas map. Results taken from Van Essen et al. [362].
33. Panel 1 shows 3D Visible Human Male. Panel 2 shows the boundaries of deformed macaque visual areas (black lines) superimposed on the fMRI activation pattern from an attentional task from the study of Corbetta et al. [447] after deformation to the Visible Man atlas by Drury et al. [448] and mapping results from [362]. Panels 3 and 4 show landmarks on macaque flat map and human flat map, respectively, used for performing the mappings shown in panel 2.
34. Columns 1 and 2 show hearts in grayscale MRI in normal (column 1) and failing (column 2) hearts; columns 3 and 4 show DTI of those sections depicting orientation of the principal eigenvectors by color. The top row shows coronal sections, the bottom row shows axial sections.
35. Figure shows the template heart and the target heart images after transformation into the coordinate system of the template. The determinant of the Jacobian of the transformation is superimposed as a colormap on the surface rendering. Blue colors indicate regions where the determinant is less than unity whereas red regions are where the determinant is greater than unity. Data courtesy of Dr. Raimond Winslow; mappings courtesy of Faisal Beg and Pat Helm.
36. Column 1 shows the geometries of the two hearts. Column 2, top panel shows the two geometries superimposed after rigid motion; blue is template, red is target. Column 2, bottom panel shows the LDDMM solution color representing the dot product of the corresponding vectors after alignment. Red color means total alignment; taken from [449].
37. MR images of postnatal mouse brains. Coronal MR T_2 and DT images were shown at the level of anterior commissure. Images have been aligned to ensure proper orientation and position. Intensity in MR T_2 images is related to tissue properties, such as the content of myelin protein. MR diffusion tensor images reveal local tissue orientation. Both types of images were utilized in our mouse brain developmental study. The color scheme for diffusion tensor images was illustrated by color arrows in the figure, with red for local tissue whose orientation is anterior-to-posterior (perpendicular to the current plane), green for horizontal orientation, and blue for vertical orientation. The scalar bar represents 1 mm. Structure abbreviations are—ac: anterior commissure; cc: corpus callosum; CPu: caudate putamen; CX: cortex.
38. Panel 1 shows white matter tracts, anterior part of anterior commissure (aca), posterior (acp), cerebral peduncle (cb), hippocampus (H), and hippocampal commissure (hc). Panels 2,3, and 4 show visualization of landmarks on white matter tracts (green dots) and landmarks inside hippocampus (yellow structures) (blue dots), outer brain surface (orange) with boundary of white matter and cortical region shown in green, and landmarks depicted as intersections of white lines on the green structure.
39. LDDMM of P7 mouse brain images to P30 images. Coronal MR images of P7 (A) and P30 (C) were shown with brain boundaries marked by orange curves. Deformed P7 (B) was overlaid with the brain boundary of P30. White arrows in A, B, and C are pointing to hippocampus (H). The transformation was visualized as color-coded vector plot (D) and Jacobian plot (F), with enlarged local areas shown in E and G, respectively. For vector plots, color of vectors represents the magnitude of local displacement. For Jacobian plot, color on deformed grids represents changes in local volume. Surface of P7, deformed P7, and P30 mouse brains were visualized in H, from top to bottom. Vector maps show how transformations deform a local voxel from its original location by taking the original location as the start of the vector and the destination as the end; length and color of vectors reflect the magnitude of tissue displacement. The Jacobian map shows local volume changes; a value greater than unity corresponds to volume increase, a value less than unity corresponds to volume loss, and unity for no volume change.
40. Outer-shape changes of mouse hippocampus measured using LDDMM. The displacement due to growth of hippocampus was color coded and the direction of growth was visualized by 3D glyphs.

41. Series of volumes depicting aging in Huntington's disease patient over time as manifest in the Caudate. Data taken from the laboratory of Dr. Elizabeth Aylward of the University of Washington.
42. 2D Face Cartoon landmark LDDMM .
43. Non-linearity artifact in panel 4 overcome by initial momentum approach shown in panel 5. Results taken from Vaillant [284].
44. First two eigenmodes of 3D PCA applied to the Morphable Faces database of 100 faces. Row 1 shows front views; row 2 shows side views. Column 1 shows template. Columns 2 and 3 show deformation via eigenfunction 1. Columns 4 and 5 show deformation via eigenfunction 2.
45. First two eigenmodes of 3D PCA applied to hippocampi from Randy Buckner of Washington University. Column 1 shows front and side view. Top row shows deformation via eigenfunction 1. Bottom row shows deformation via eigenfunction 2.
46. Left column panel 1 shows a whole MRI-MPRAGE image volume with a section through the brain delineating the surface of the hippocampus. Left column panel 2 shows the triangulated graph representing the mean state of the hippocampus. The data are courtesy of Dr. John Csernansky of the Department of Psychiatry at Washington University. Right column four panels show the first four surface harmonics visualized through deformation of the template (panel 2). Courtesy Sarang Joshi, Ph.D. thesis [242].
47. Top row: Panels 1–3 show maps of the initial hippocampus to three in a population of 30 patients, $M_0 \circ g_1, M_0 \circ g_2, M_0 \circ g_3$ studied in Csernansky et al. [162]. Panel 4 shows the composite template generated from the average of 30 maps, $M_0 \circ \bar{g}$. Bottom left column shows the first two eigenshapes of the left and right hippocampus generated from a population of maps of normals and schizophrenics. Lower right panels show random instances of left and right hippocampi generated via the empirically estimated covariance function for the Gaussian random field on the hippocampus surface. Taken from Joshi [242].
48. Top row: Schizophrenia. Panel 1 shows difference of hippocampal surface patterns between the control and schizophrenia groups visualized as z -scores on the mean surface of the control group. Inward surface deformations due to schizophrenia are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 2 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,3,4,6,10,15}. Multivariate ANOVA indicates significant between-group difference: $p = 0.0028$ ($F = 4.73$, $df = 1, 28$). Bottom row: Asymmetry in Schizophrenia. Panel 3 shows difference of hippocampal surface asymmetry patterns between the control and schizophrenia groups visualized as z -scores on the mean flipped right-side surface of the control group. Inward surface deformations due to differences in asymmetry are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 4 shows statistical testing of asymmetry pattern differences: log-likelihood ratios for a linear combination of basis functions {2,12,17}. Multivariate ANOVA indicates significant between-group difference in asymmetry: $p = 0.0029$ ($F = 6.03$, $df = 3, 26$).
49. Top row: Alzheimer's Disease. Panel 1 shows difference of hippocampal surface patterns between the elderly control (CDR 0) and AD (CDR 0.5) groups visualized as z -scores on the mean surface of the elderly group. Inward surface deformations due to AD are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 2 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,5}. Multivariate ANOVA indicates significant between-group difference: $p = 0.0002$ ($F = 11.4$, $df = 2, 33$). Bottom row: Normal Aging. Panel 3 shows difference of hippocampal surface patterns between the younger and the elderly control (CDR 0) groups visualized as z -scores on the mean surface of the younger group. Inward surface deformations due to aging are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 4 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,2}. Multivariate ANOVA indicates significant between-group difference: $p < 0.0001$ ($F = 348$, $df = 2, 30$).

50. AD Progression. Panel 1 shows deformations from baseline to follow-up for the CDR 0.5 AD group. Inward surface deformations due to AD are visualized in colder colors, outwards in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 2 shows deformations from baseline to follow-up for the CDR 0 control group. Inward surface deformations due to AD are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Panel 3 shows follow-up-versus-baseline “spread” of the between-group inward surface deformation patterns, shown as Wilcoxon’s sign rank test map on the CDR 0 mean surface. Areas of significant ($p < 0.05$) inward deformation at baseline of CDR 0.5 group are shown in turquoise color, representing 38% of total hippocampal surface area. By follow-up, areas of significant inward deformation have increased to 47% of total hippocampal surface area. The increased affected areas are shown in purple color. Areas of non-significant surface deformation are shown in green color. Panel 4 shows statistical testing of shape differences: log-likelihood ratios for a linear combination of basis functions {1,2,4,11}. Multivariate ANOVA of the first 12 basis functions indicates significant between-group difference: $p = 0.014$ ($F = 2.66$, $df = 12,31$).
51. Depression. Panel shows differences of hippocampal surface patterns between the control and depression groups visualized as perpendicular displacements on the mean surface of the control group. Inward surface deformations due to depression are visualized in colder colors, outward in warmer colors, and areas which are not deformed in neutral yellow to green colors. Multivariate ANOVA of the first 10 basis functions indicates significant between-group difference: $p < 0.0001$ ($F = 34.1$, $df = 10,58$).
52. Empirical distribution \hat{F} from randomized Fisher’s test with 10,000 group permutations, between the group of control and schizophrenia subjects. Basis vectors {2,12,17} are selected. The $p = 0.0031$ value shown is calculated from Eqn. (17.25). Also shown are (i) $\hat{F}(T^2)$ value (solid blue line) of the control-versus-schizophrenia group comparison; (ii) theoretical F -distribution (solid red curve) with (3,26) degrees of freedom superimposed on the empirical distribution; and (iii) $p = 0.05$ (red dotted line) and $p = 0.01$ (red dot-dash line) for reference.
53. Figure showing heteromorphic deformations. The left panel shows an expansion or “push out”; the right panel shows a “push in”.