Optimal, Robust Information Fusion in Uncertain Environments

MURI Kickoff Meeting

Integrated Fusion, Performance Prediction, and Sensor Management for Automatic Target Exploitation

Alan S. Willsky

July 21, 2006

What is needed: An expressive, flexible, and powerful framework

- **Capable of capturing uncertain and complex sensor-target relationships**
  - Among a multitude of different observables and objects being sensed
- **Capable of incorporating complex relationships about the objects being sensed**
  - Context, behavior patterns
- **Admitting scalable, distributed fusion algorithms**
- **Admitting effective approaches to learning or discovering key relationships**
- **Providing the "glue" from front-end processing to sensor management**

Our choice: Graphical Models

- Extremely flexible and expressive framework
  - Allows the possibility of capturing (or learning) relationships among features, object parts, objects, object behavior, and context
    - E.g., constraints or relationships among parts, spatial and spatio-temporal relationships among objects, etc.
  - Natural framework to consider distributed fusion
- While we can't beat the dealer (NP-Hard is NP-Hard),
  - The flexibility and structure of graphical models provides the potential for developing scalable, approximate algorithms

Graphical Models 101

- $\boldsymbol{G} = (V, E)$ = a graph
  - V = Set of vertices
  - $E \subset V \times V$ = Set of edges
  - **C** = Set of cliques
- Markovianity on $\boldsymbol{G}$ (Hammersley-Clifford)

$$P(\{x_s \mid s \in V\}) \propto \prod_{C \in \mathbf{C}} \psi_C(x_C)$$

Objectives

Estimation: Compute $P_s(x_s)$

Optimization: Argmax $P(\{x_s \mid s \in V\})$

Algorithms that do this on trees

- *Message-passing* algorithms for "estimation" (marginal computation)
  - Two-sweep algorithms (leaves-root-leaves)
    - For linear/Gaussian models, these are the generalizations of Kalman filters and smoothers
  - Belief propagation, sum-product algorithm
    - Non-directional (no root; all nodes are equal)
    - Lots of freedom in message scheduling
- *Message-passing* algorithms for "optimization" (MAP estimation)
  - Two sweep: Generalization of Viterbi/dynamic programming
  - Max-product algorithm

In pictures...

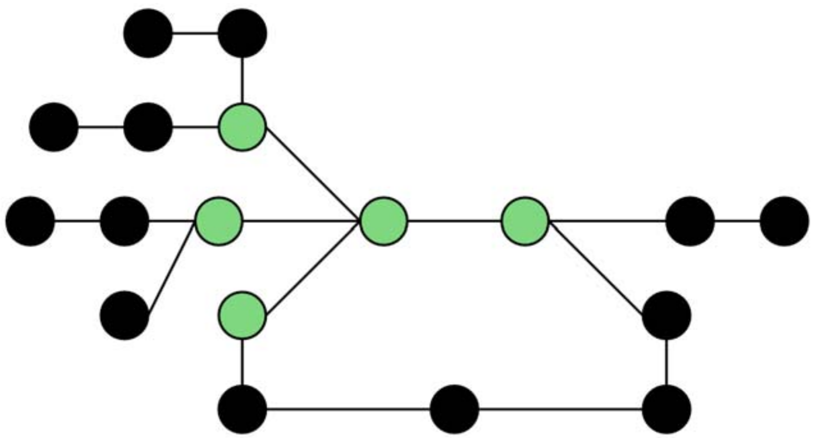

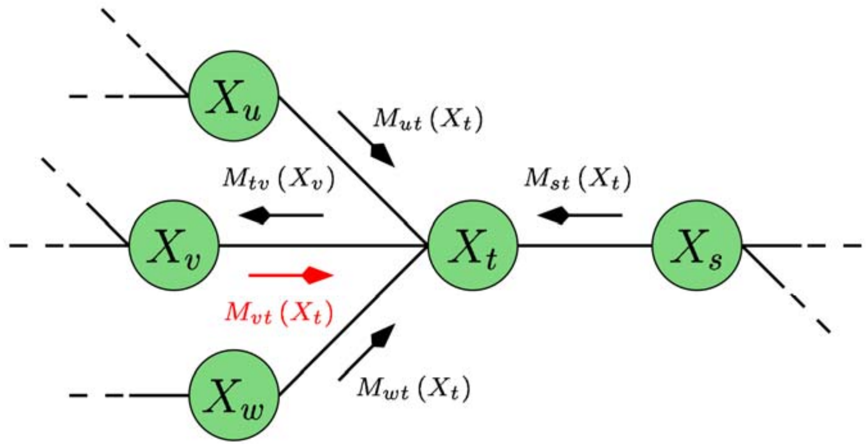
X_u
X_v
X_t
X_s
X_w
$M_{ut}(X_t)$
$M_{tv}(X_v)$
$M_{st}(X_t)$
$M_{vt}(X_t)$
$M_{wt}(X_t)$

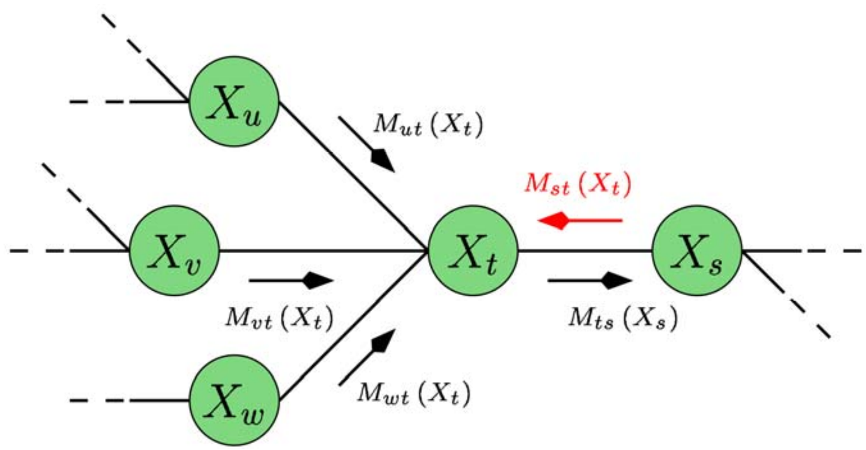
X_u
X_v
X_t
X_s
X_w
$M_{ut}(X_t)$
$M_{st}(X_t)$
$M_{vt}(X_t)$
$M_{ts}(X_s)$
$M_{wt}(X_t)$

Comments and questions - I

- Other natural inference problems can be thought of as hypothesis testing on such models
  - Estimating potentials
  - Discovering/estimating “links”
  - Distributed inference in comms-limited environments
- Performing inference tasks such as these
  - Wonderfully scalable if the graphs are trees
  - NP-Hard in general if they are not (which is the case for essentially all problems in this MURI)
  - Beating the dealer in this case is crucial

Comments and questions - II

- The Glue: Graphical models can capture relationships among observables and objects
  - Allowing object hypotheses to influence front-end processing
    - Think PEMS but with more sophisticated feedback from object hypotheses to front-end processing
  - Allowing object hypotheses to influence what measurements should be made
    - Again, think PEMS, with an expanded notion of what is involved in “search”

The Three Big Questions

- Why should we believe that graphical models can capture things of interest to this MURI?
- Why should we believe that it is possible to develop tractable and useful fusion algorithms based on such models?
- What are we going to do?

Question #1: Why we think that graphical models are relevant—A few f'rinstances

- Multi-target tracking and data association
- Tracking of dynamically coupled and constrained objects
- Flexible framework for learning features-to-parts-to-objects models for object recognition in complex scenes
- Graphical Models for Shapes and their projections
- Robust association of heterogeneous signals and discovery of "links" among observed (possibly dynamic) variables
- Significant theoretical advances
  - Performance bounds
  - Substantial generalization of particle filtering
  - New algorithms

Data Association and Multi-Target Tracking and ID

- Great flexibility in representing data association as a problem of inference on a graphical model
- Distributed fusion naturally leads to sensor- and region-oriented framework
  - Nodes ⇔ Sensors, Regions, Groups of targets seen by common set of sensors
  - Variables ⇔ Assignments of measurements to regions, targets to regions, measurements to targets
- Very different from standard MHT
- Leads naturally to comms-sensitive message-passing and very efficient iterative algorithms

Illustrating comms-sensitive message-passing dynamics

Organized network data association

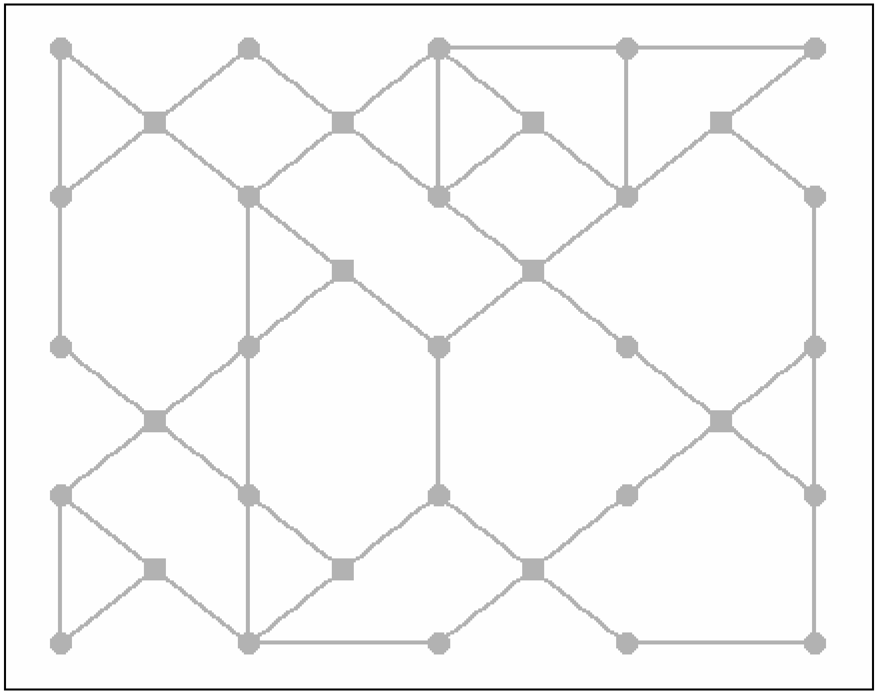

Self-organization with region-based representation

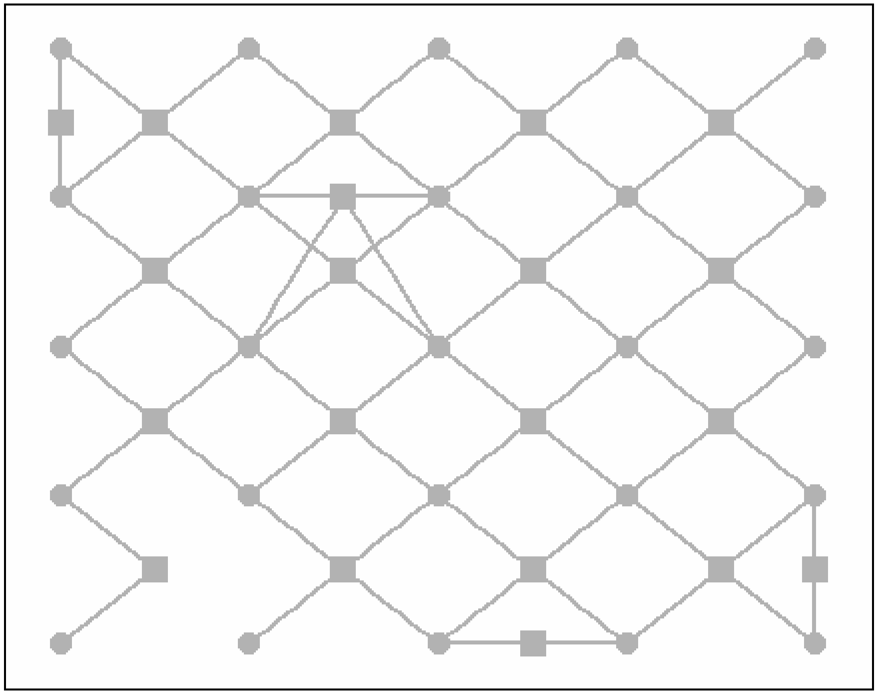

Tracking over time and beating the MHT-combinatorial dealer

- Add nodes that allow us to separate target dynamics from discrete data associations

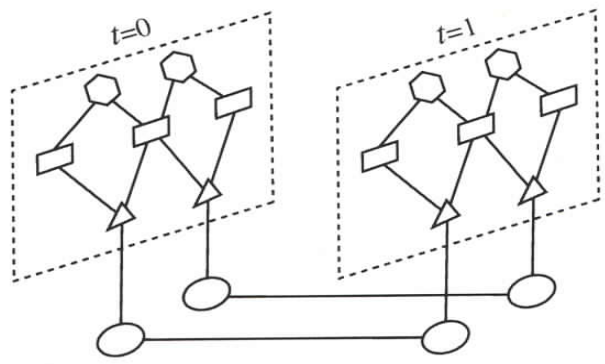

- Perform explicit data association within each frame (using evidence from other frames)
- Stitch across time through temporal dynamics

Dynamic fusion in complex, constrained contexts

Example: Graphical Model for Shape-Tracking with Level Sets

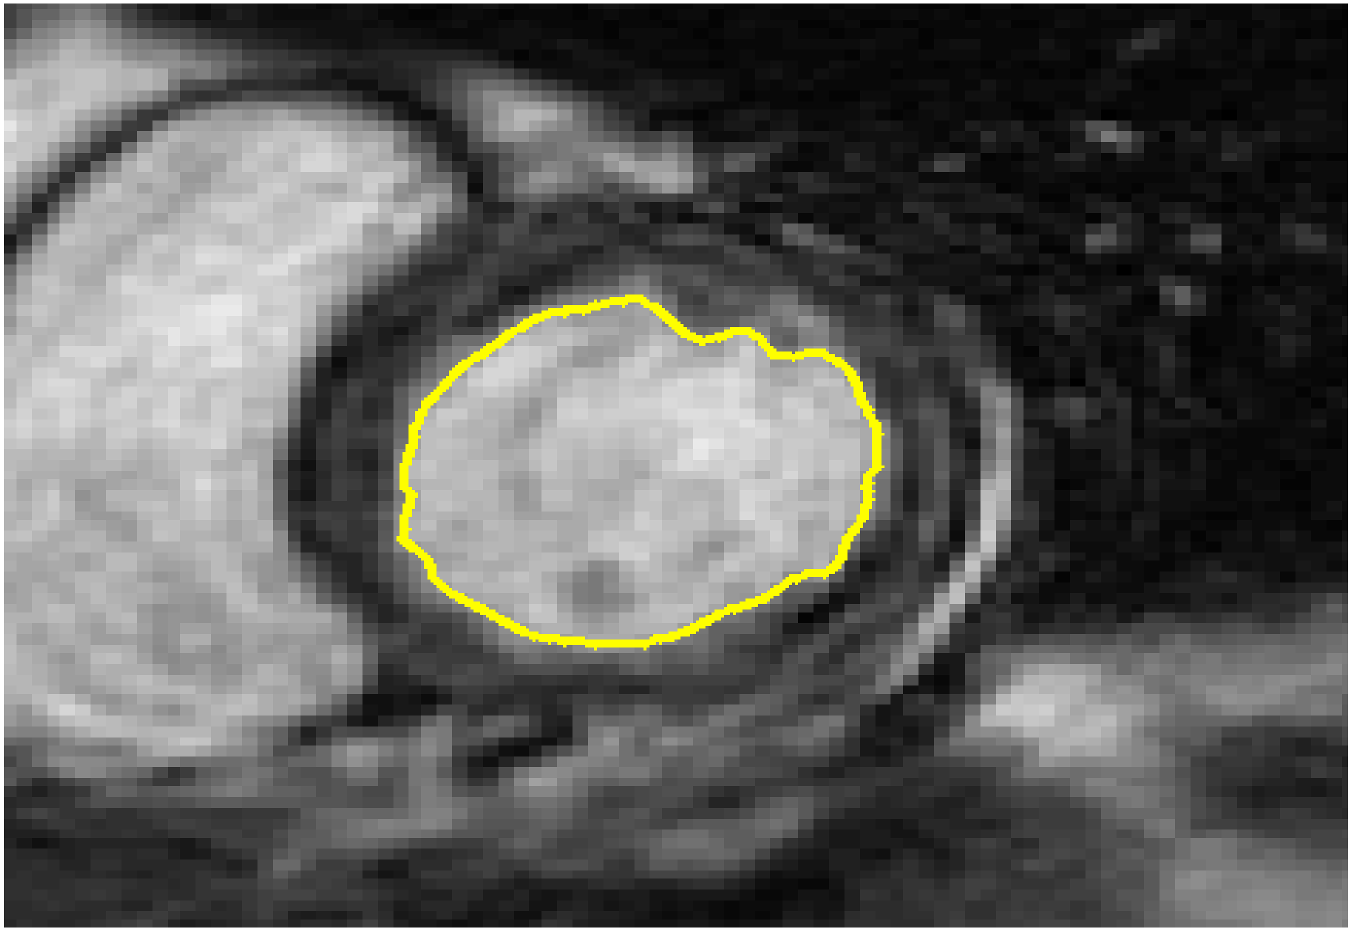

Hierarchical Graphical Models: From Scene/context to objects to parts/shape to features

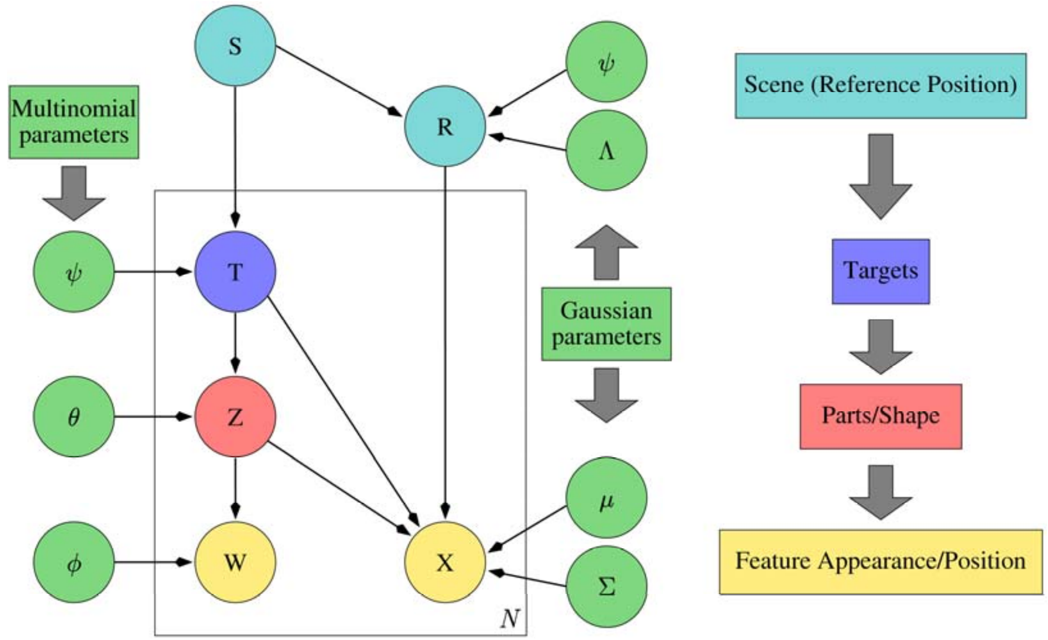

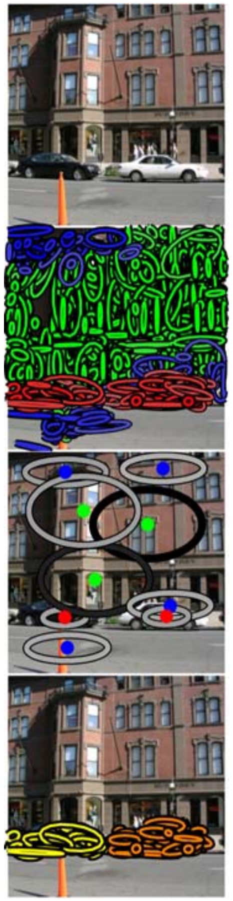

Determining graphical structure: Capturing statistical links among objects

- Object 3 tries to interpose itself between objects 1 and 2.
- The graph describes the state (position) dependency structure.

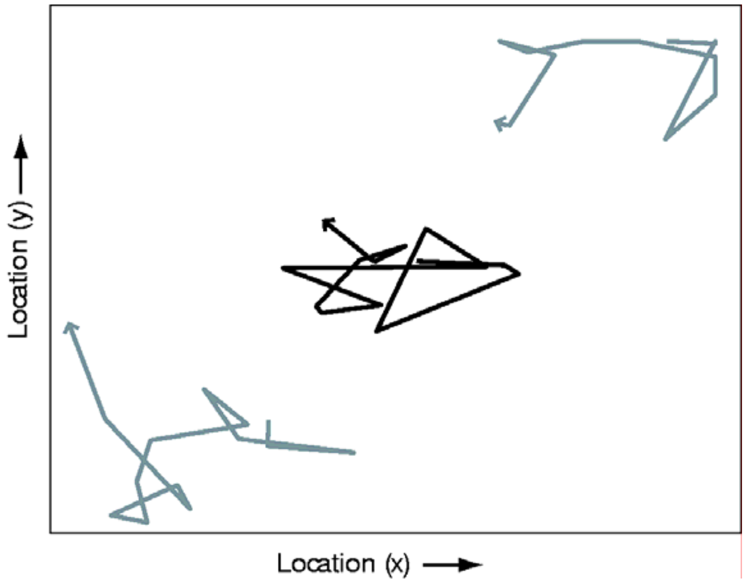

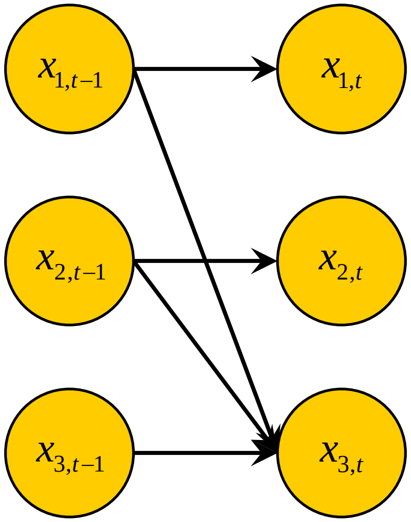

Modeling Group Interactions

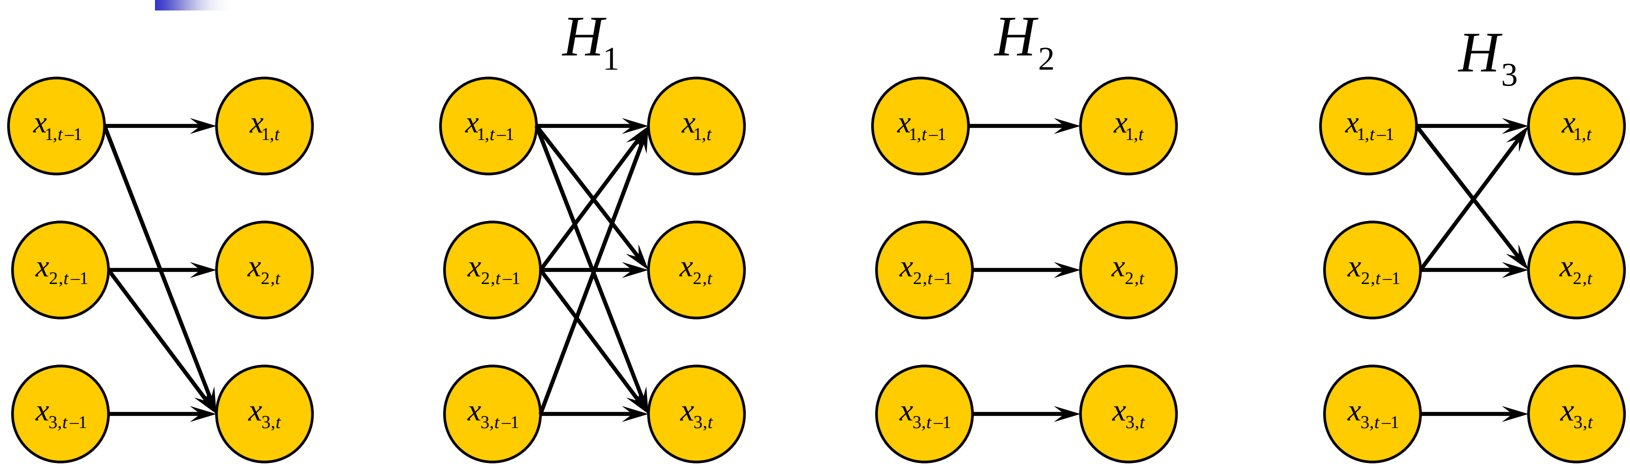

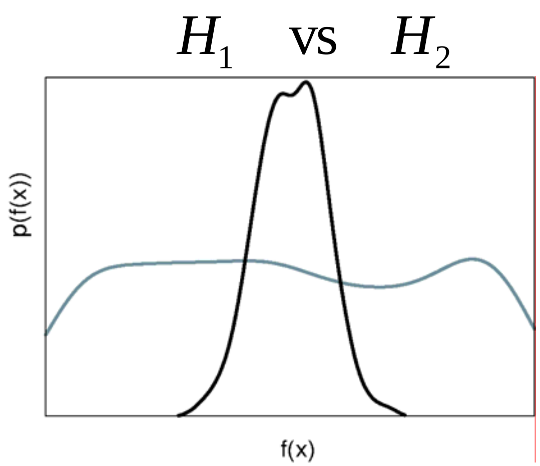

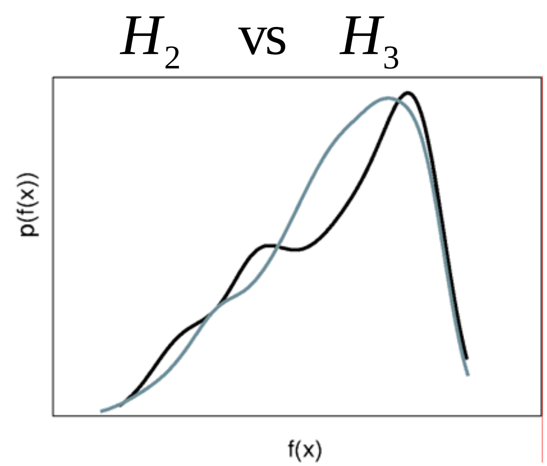

High-Dimensional, Multi-modal Data Association

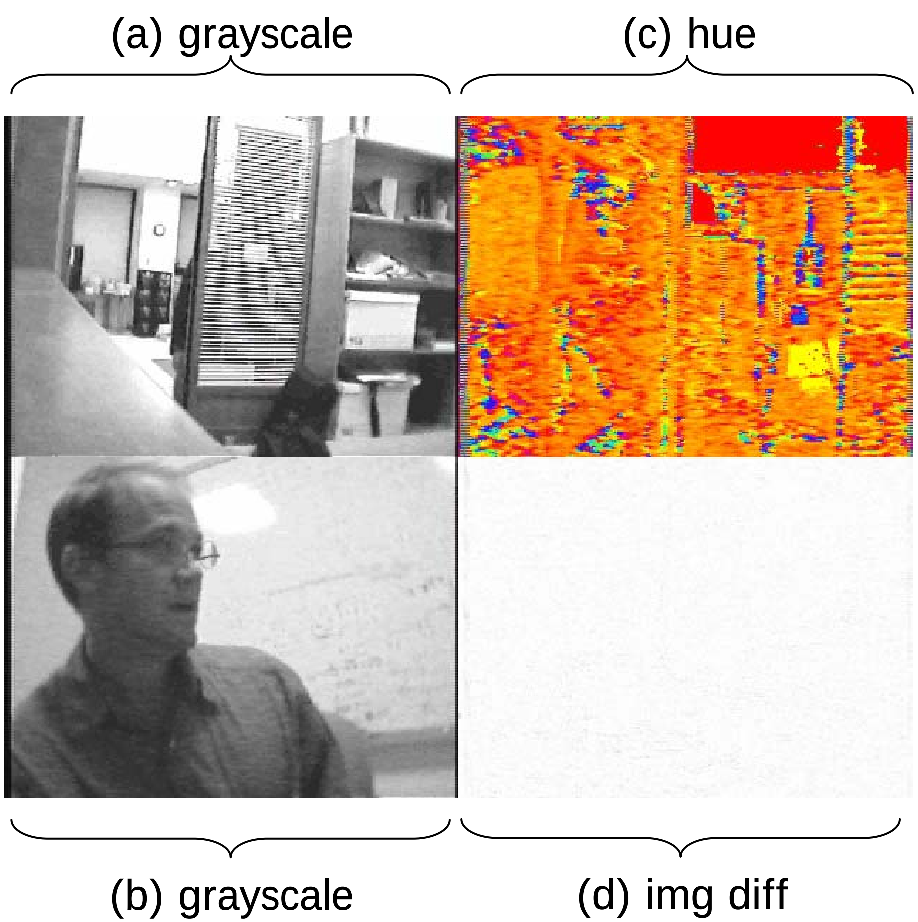

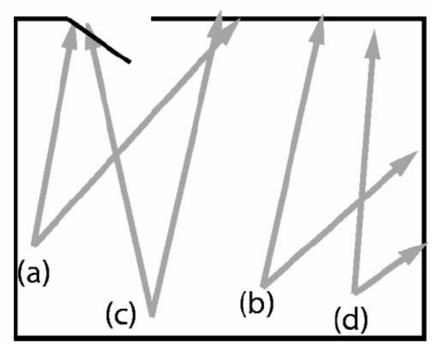

Data Association Example

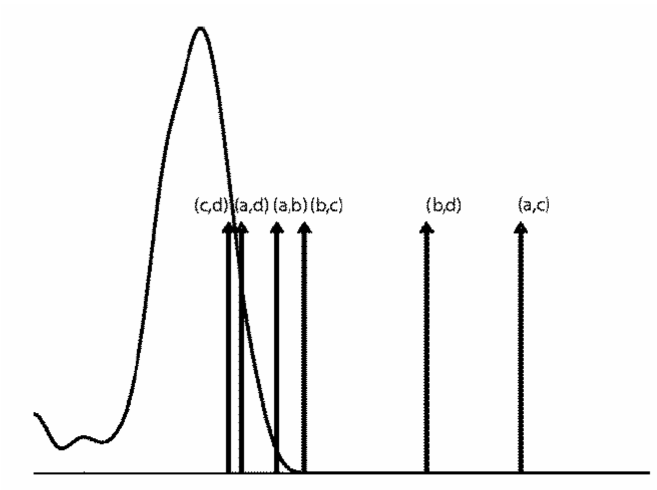

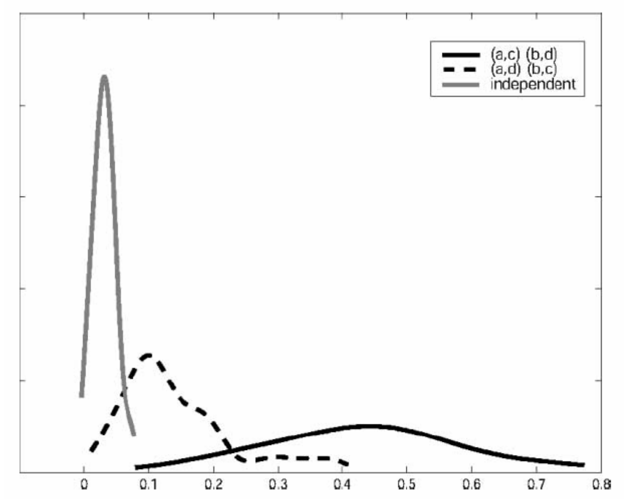

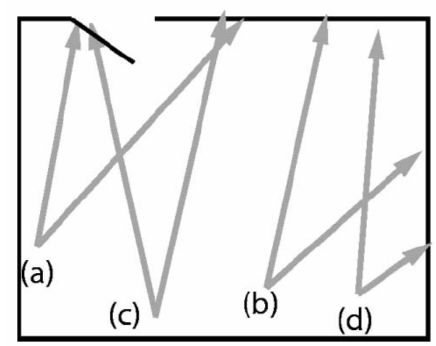

- LLR estimates for pair-wise associations (left)
  - Compared to the distribution over the null hypothesis
- Distribution of full association (middle)
  - Incorrect association likelihood shows some global scene dependence (e.g. due to common lighting changes)

Question #2: Why we think we can beat the dealer

- The previous examples are ***not*** small examples
- But there is much more that needs to be done to deal with the exponential complexity of the challenges we hope to confront
  - But there is hope...

Recursive Cavity Modeling: Remote Sensing Application

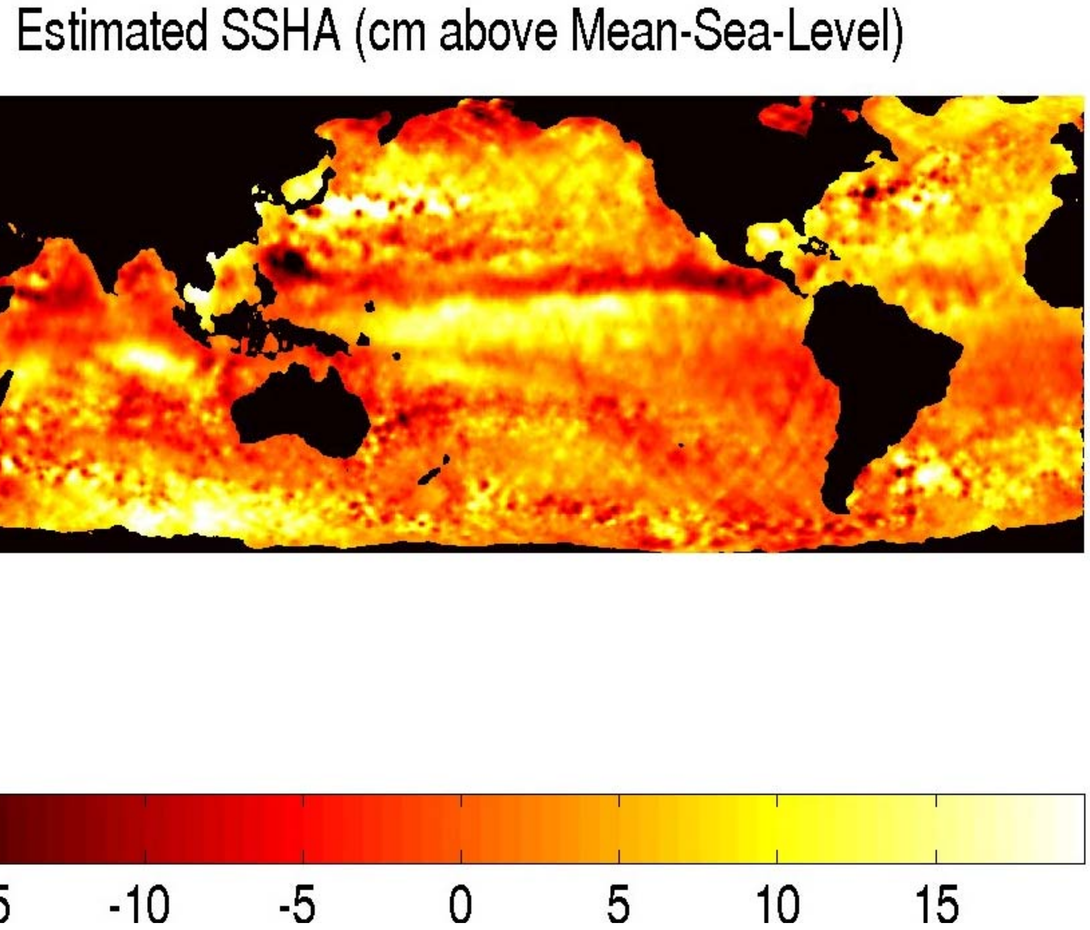

Question #3: What we'll be doing I: Topics where we're up and running

- Scalable, broadly applicable inference algorithms
  - Build on the foundation we have
  - Provide performance bounds/guarantees
  - Provide framework for distributed implementation
    - Including building/learning collaborative fusion algorithms
- Graphical-model-based methods for sensor fusion for tracking, and identification
  - Will also allow us to incorporate and demonstrate the value of context
    - Graphical models to capture multi-object motion patterns associated with suspicious/threatening activities
    - Graphical models to capture relationships among features-parts-objects

Question #3: What we'll be doing II: Topics that are on the horizon

- Learning model structure
  - Discovering links (e.g., detecting coordination)
  - Exploiting and extending advances in learning (e.g., information-theoretic and manifold-learning methods) to build robust models for multimodal fusion
- Driving the front end
  - Higher-level hypotheses drive signal processing (for feature extraction and to answer "queries")
  - For example: high-level information on shape, scene, objectives/performance used to guide choice of sparse representation dictionaries
    - Think PEMS: If we're looking for an anisotropic scatterer in a particular location, guide the front end to do this

Question #3: What we'll be doing III: Topics that are on the horizon (cont.)

- Informing resource management
  - Using informational structure of a graphical model to decide what evidence to gather
    - What messages should be sent to inform specific hypotheses (information pull rather than push)?
    - What measurements should be taken to reduce critical sources of uncertainty (e.g., at particular graphical nodes)?
    - In distributed processing, dealing with handoff of inference responsibilities
  - Answering these requires information- and graph-theoretic performance bounds