Diverse Particle Selection for High-Dimensional Inference in Graphical Models

Erik Sudderth

UC Irvine Computer Science

Collaborators:

- ***Particle Max-Product:*** *Jason Pacheco, MIT*
- ***Human Pose:*** *Silvia Zuffi & Michael Black, MPI Tubingen*

Related papers at ICML 2014 & ICML 2015

High-Dimensional Inference

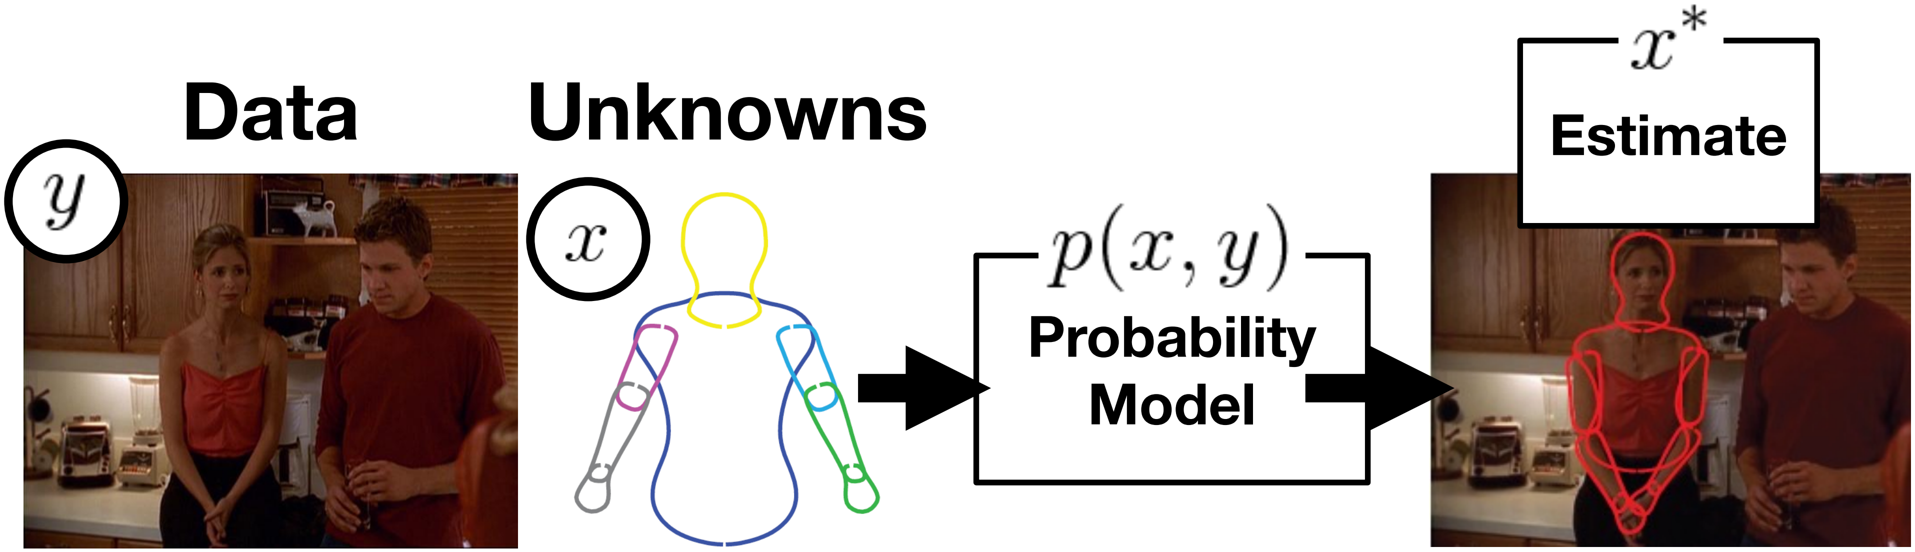

Discrete Unknowns	Efficient inference based on *combinatorial optimization*
Continuous Unknowns	Unless we make unrealistic model approximations, *no efficient general solutions. Standard gradient-based optimization is ineffective.*

Continuous Inference Problems

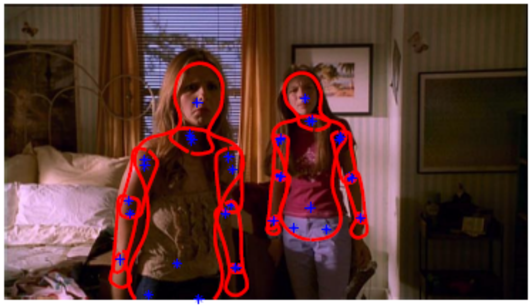

Human pose
estimation & tracking

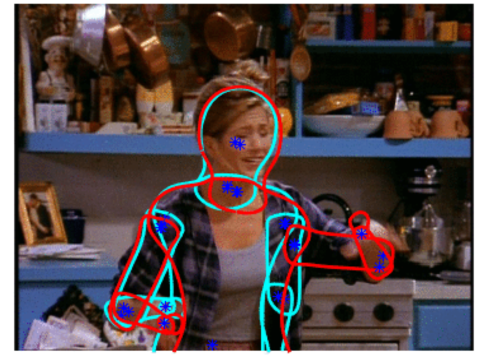

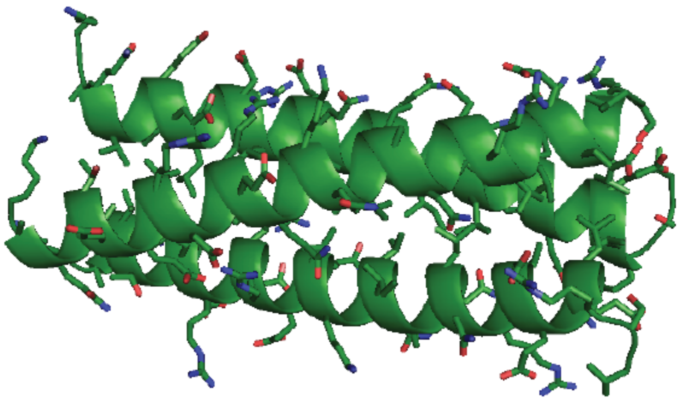

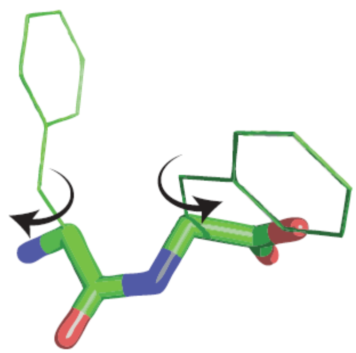

Protein structure &
side chain prediction

Robot motion &
vehicle path planning

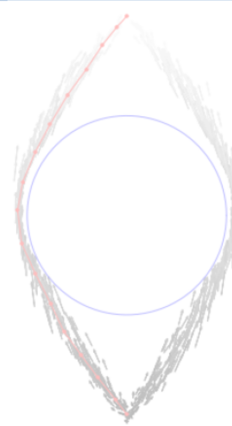

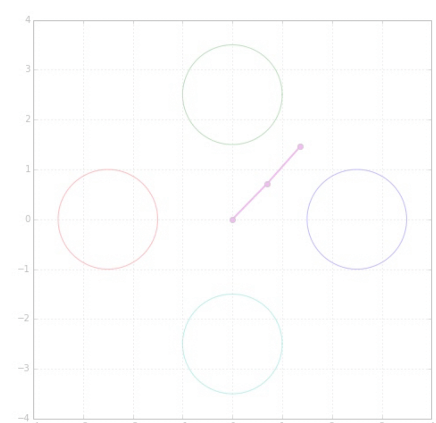

Maximum a Posteriori (MAP)

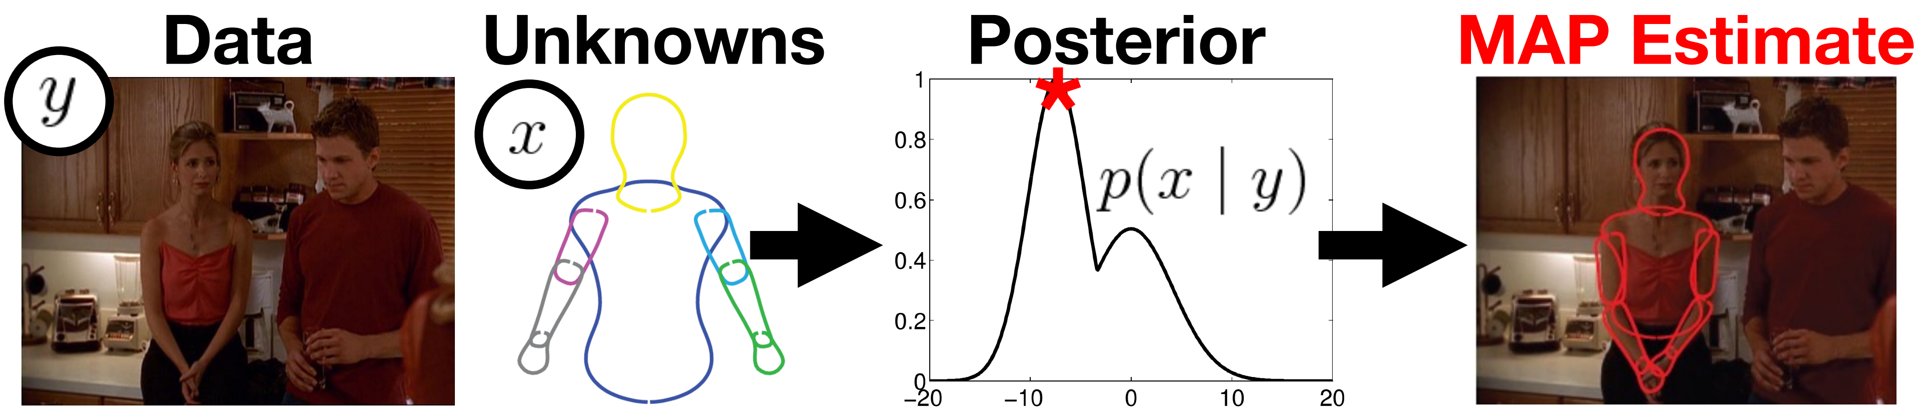

Posterior often intractable and multimodal complicating exact MAP inference:

$$x^* = \underset{x}{\operatorname{argmax}}\, p(x \mid y)$$

Maximum a Posteriori (MAP)

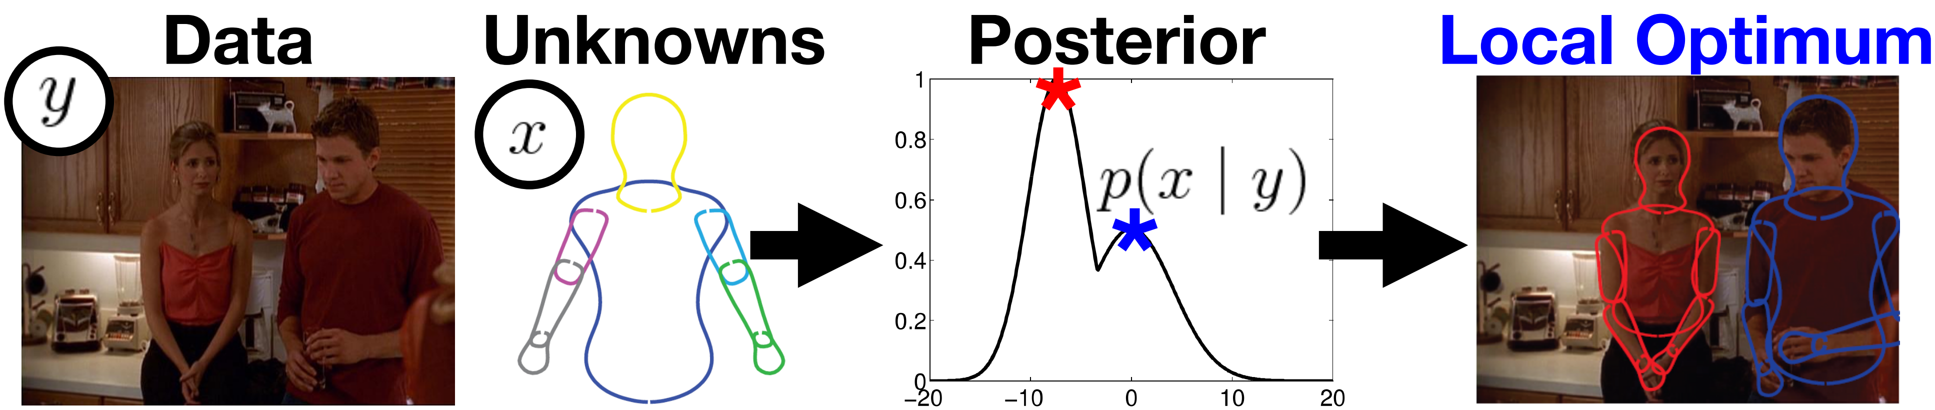

Posterior often intractable and multimodal complicating exact MAP inference:

$$x^* = \underset{x}{\operatorname{argmax}}\, p(x \mid y)$$

Local optima can be useful when models are inaccurate or data are noisy.

Goal

Develop maximum a posteriori (MAP) inference algorithms for continuous probability models that:

- Apply to *any pairwise graphical model*, even if model is complex (highly non-Gaussian)
- Are *black-box* (no gradients required)
- Will *reliably* infer multiple local optima

Pairwise Graphical Models

$$p(x) \propto \prod_{s \in \mathcal{V}} \psi_s(x_s) \prod_{(s,t) \in \mathcal{E}} \psi_{st}(x_s, x_t) \qquad x_s \in \mathbb{R}^d$$

- Nodes are continuous random variables
- Potentials encode statistical relationships
- Edges indicate direct, pairwise energetic interactions

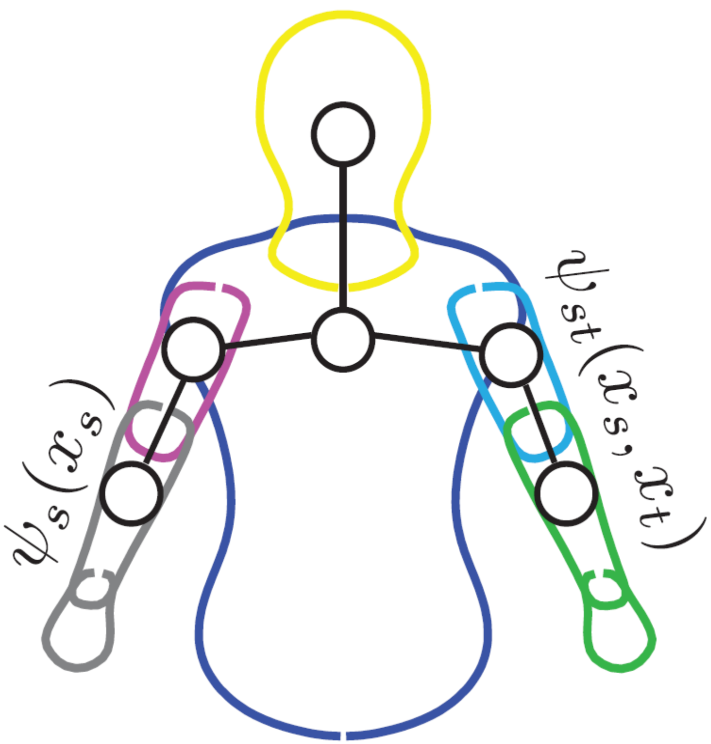

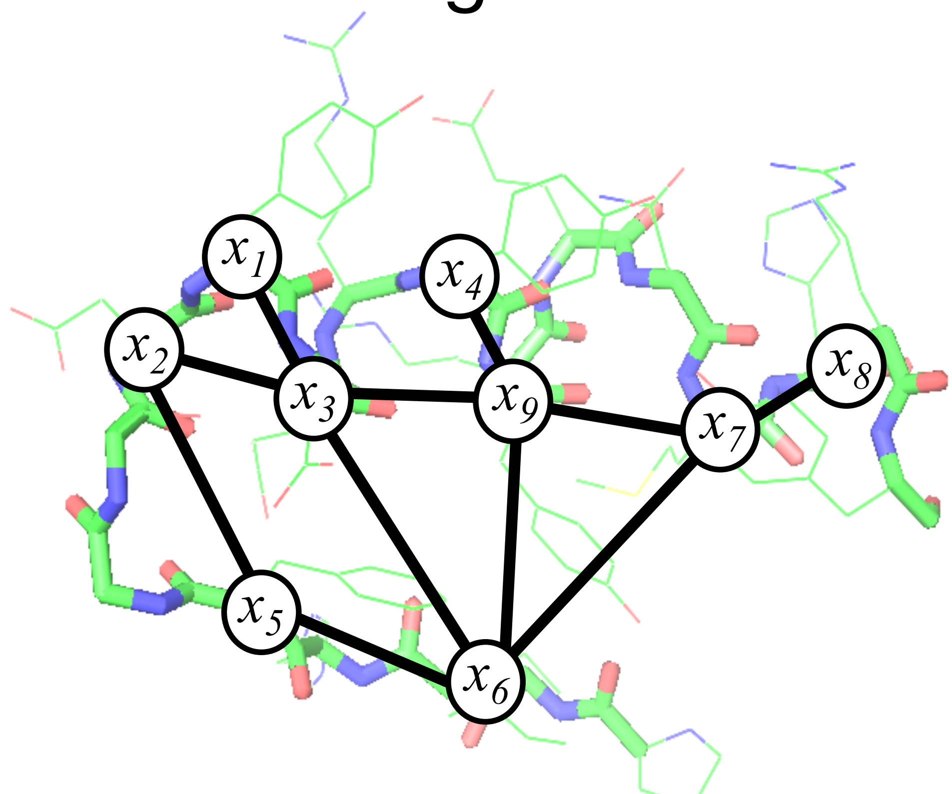

Message Passing on Trees

Global MAP inference decomposes into local computations via graph structure…

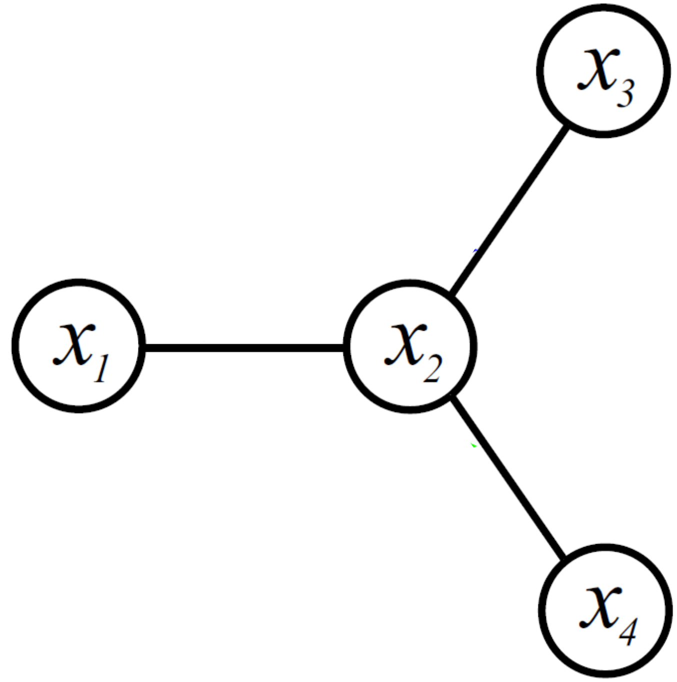

$$\max_x p(x) \propto \max_x \psi_{12}(x_1, x_2)\psi_{23}(x_2, x_3)\psi_{24}(x_2, x_4)$$

$$= \underbrace{\max_{x_1} \max_{x_2} \psi_{12}(x_1, x_2) \underbrace{\left[\max_{x_3} \psi_{23}(x_2, x_3)\right]}_{m_{32}(x_2)} \underbrace{\left[\max_{x_4} \psi_{24}(x_2, x_4)\right]}_{m_{42}(x_2)}}_{m_{21}(x_1)}$$

Max-Product Belief Propagation

Finding max-marginals via message-passing

$$q_s(x_s) = \max_{x_{t \neq s}} p(x_s, x_{t \neq s}) \propto \psi_s(x_s) \prod_{t \in \Gamma(s)} m_{ts}(x_s)$$

$$m_{ts}(x_s) = \max_{x_t} \psi_{st}(x_s, x_t)\psi_t(x_t) \prod_{k \in \Gamma(t) \setminus s} m_{kt}(x_t)$$

Why max-marginals?

- Directly encode global MAP
- Other modes important: models approximate, data uncertain

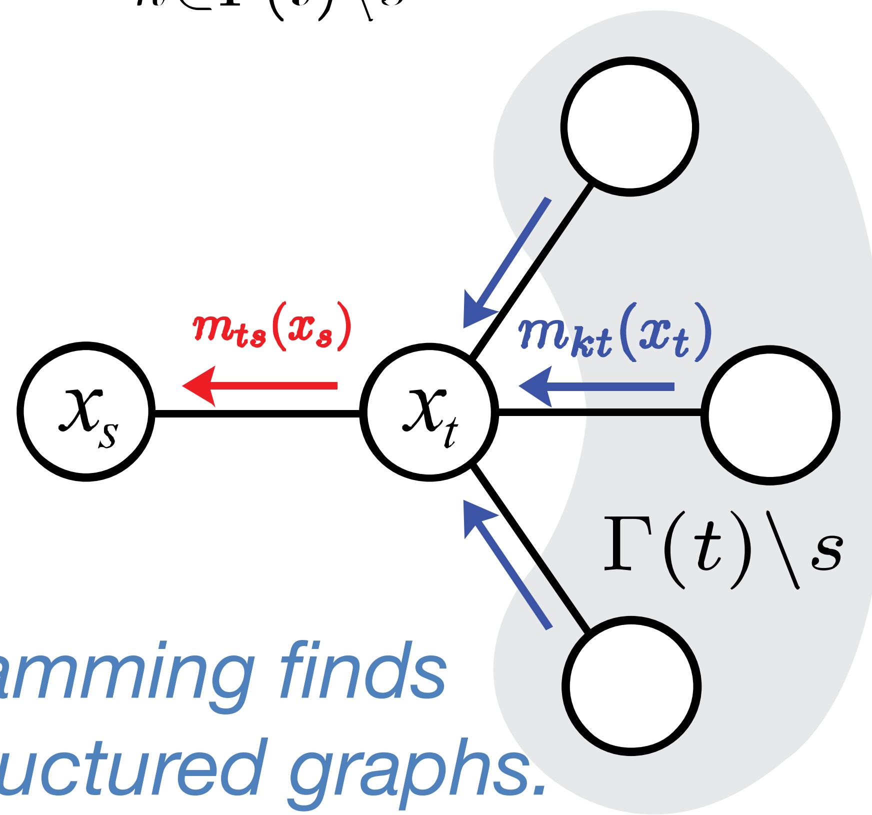

Max-product dynamic programming finds exact max-marginals on tree-structured graphs.

Articulated Pose Estimation

[Zuffi et al., CVPR 2012]

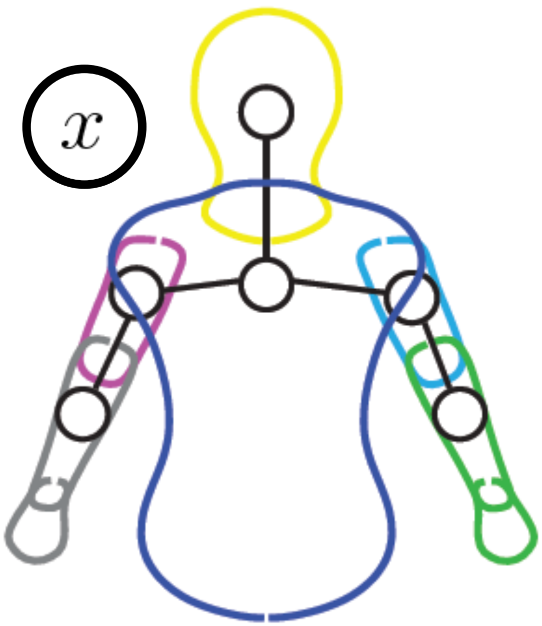

$$p(x,y) \propto \prod_{s\in\mathcal{V}} \psi_s(x_s, y) \prod_{(s,t)\in\mathcal{E}} \psi_{st}(x_s, x_t)$$

Complicated Likelihood

Non-Gaussian Compatibility

Deformable Structures (DS):

Continuous state $x_s \in \mathcal{X}_s$ for part

shape, *location*, *orientation*, *scale*.

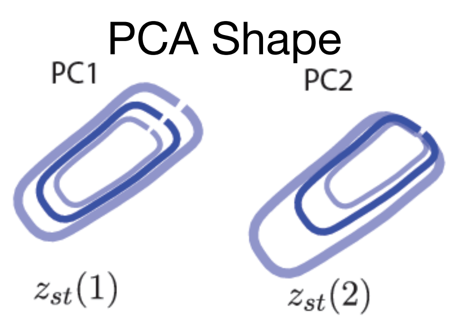

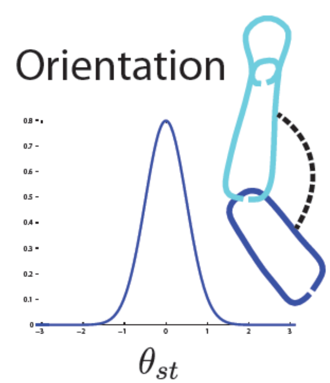

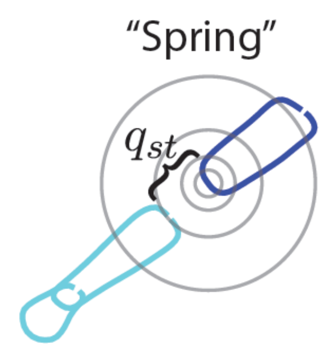

Poses & Discrete Energies

Fischler & Elschlager, 1973

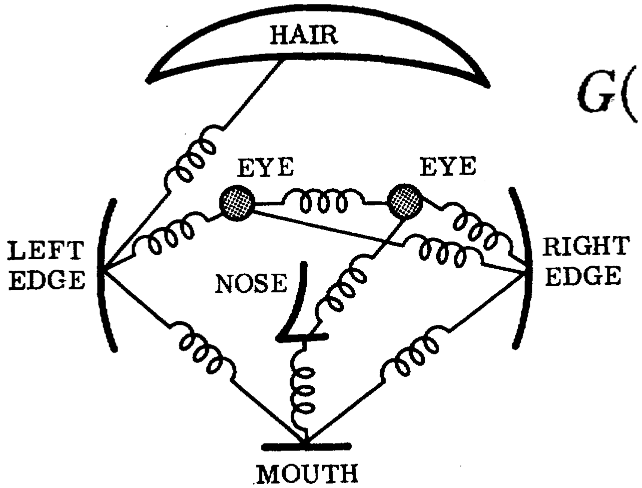

$$G(X_p) = \sum_{i=1}^{p} \sum_{j=1}^{i} g_{ij}(x_i, x_j)$$

$x_i - x_j = (z_i - z_j,\ y_i - y_j)$	$g_{ij}(x_i - x_j)$
1,0	0
2,0	1
otherwise	∞

Localize object by minimizing cost or energy defined by synthetic springs.

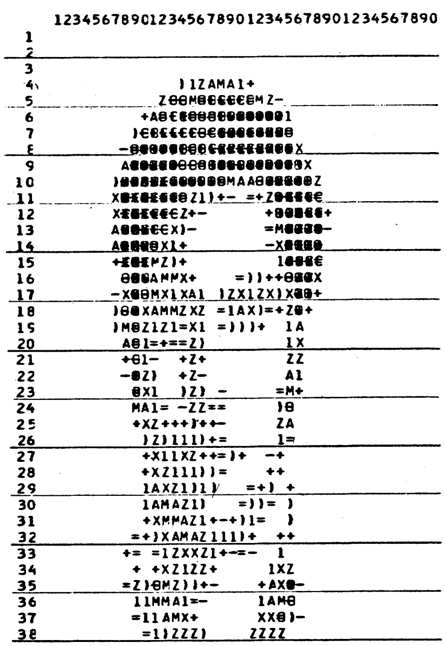

Original picture.

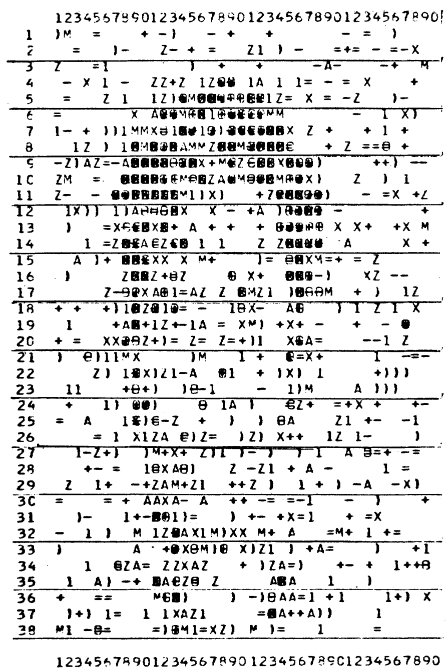

Noisy picture (sensed scene) as used in experiment.

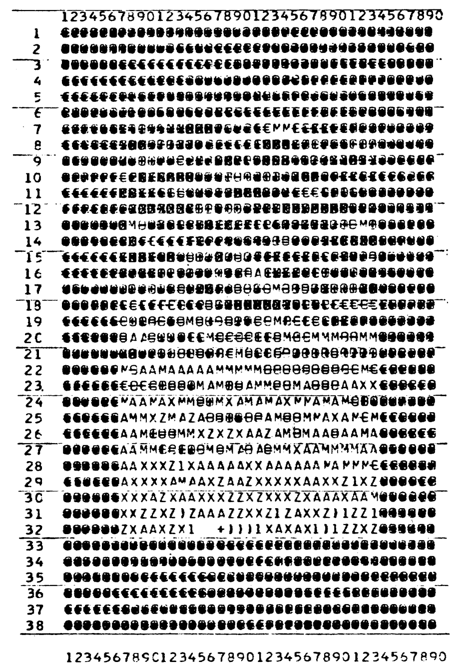

L(EV)A for nose. (Density at a point is proportional to probability that nose is present at that location.)

HAIR WAS LOCATED AT (8, 21)
L/EDGE WAS LOCATED AT (17, 11)
R/EDGE WAS LOCATED AT (17, 25)
L/EYE WAS LOCATED AT (17, 14)
R/EYE WAS LOCATED AT (17, 20)
NOSE WAS LOCATED AT (21, 16)
MOUTH WAS LOCATED AT (23, 16)

Poses & Discrete Probabilities

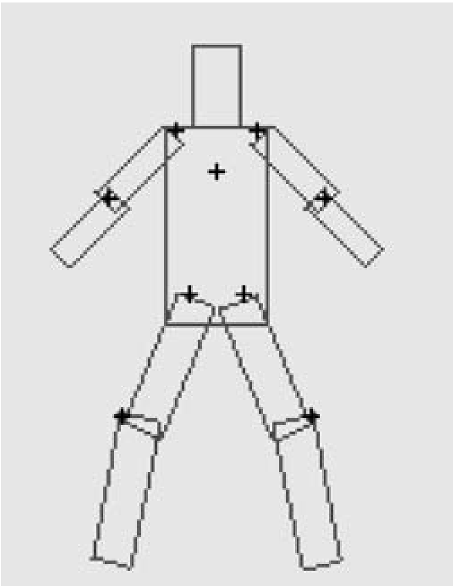

Felzenszwalb & Huttenlocher, 2005

Localize object via MAP estimate in pairwise MRF with rigid geometry.

$$p(L \mid \theta) = \frac{\prod_{(v_i, v_j) \in E} p(l_i, l_j \mid \theta)}{\prod_{v_i \in V} p(l_i \mid \theta)^{\deg v_i - 1}}$$

$$p(I \mid L, \theta) = p(I \mid L, u) \propto \prod_{i=1}^{n} p(I \mid l_i, u_i).$$

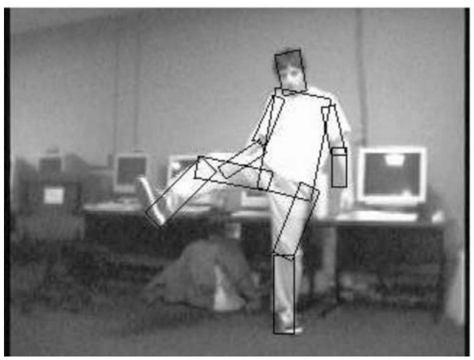

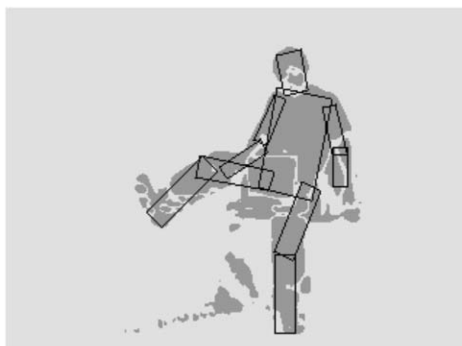

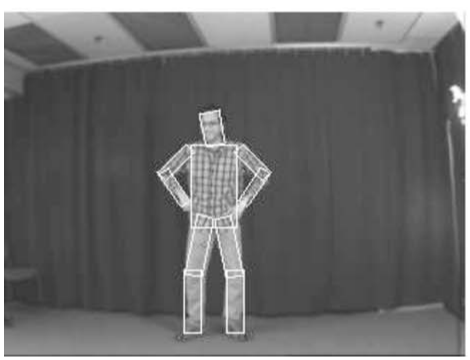

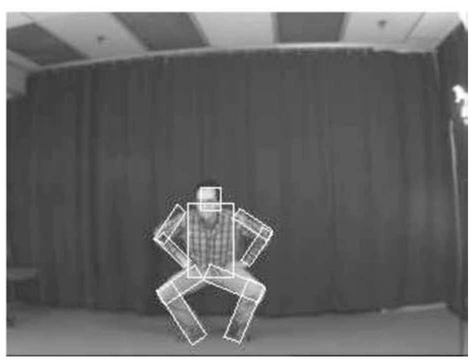

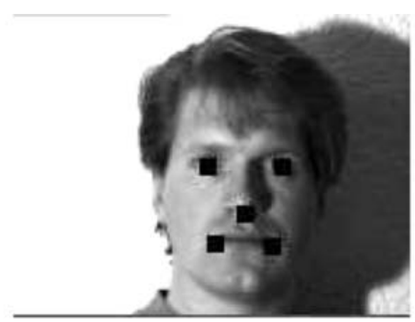

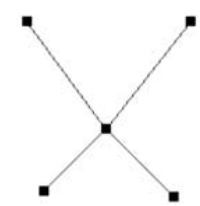

SCAPE

Shape Completion and Animation of People, Anguelov et al. 2004

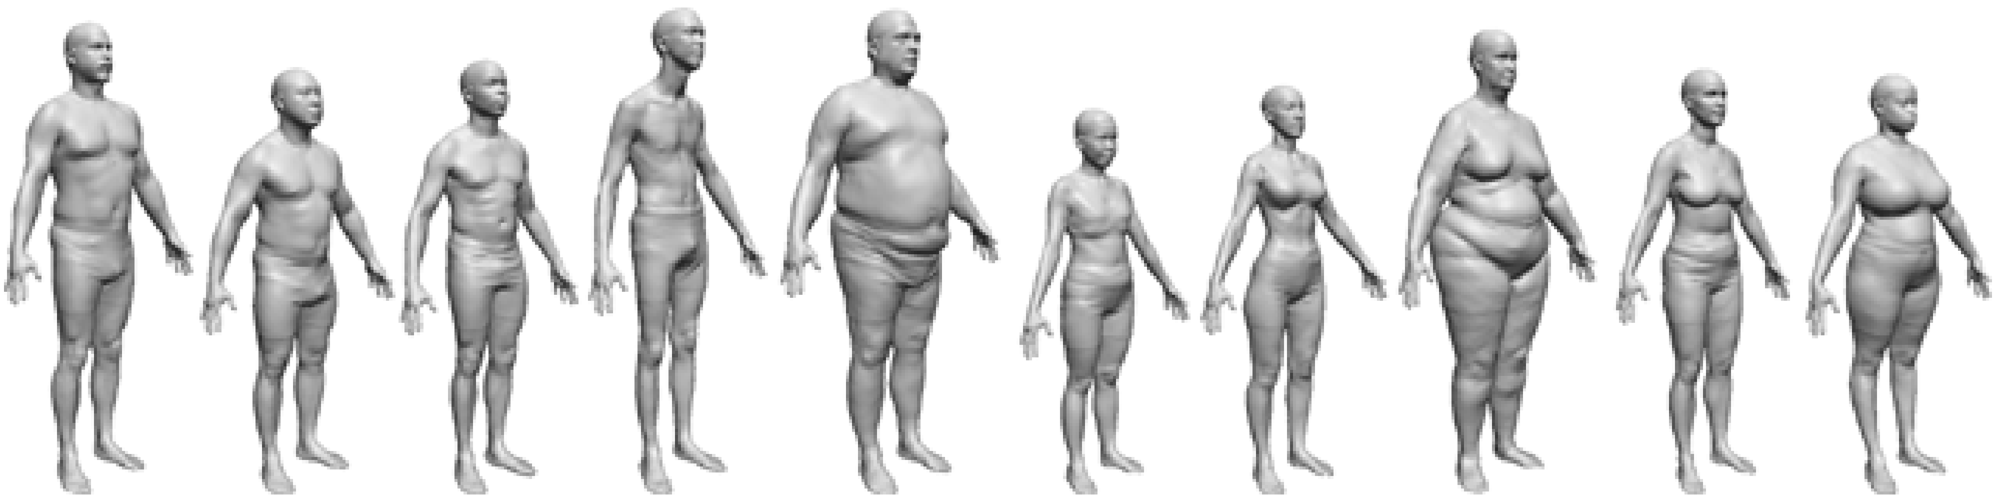

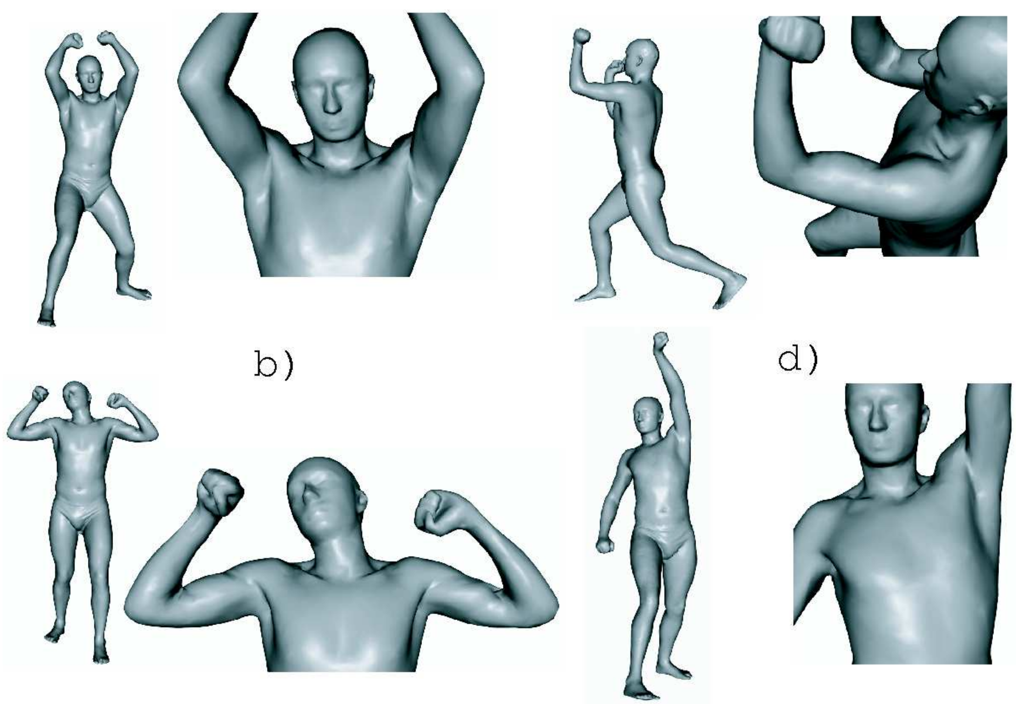

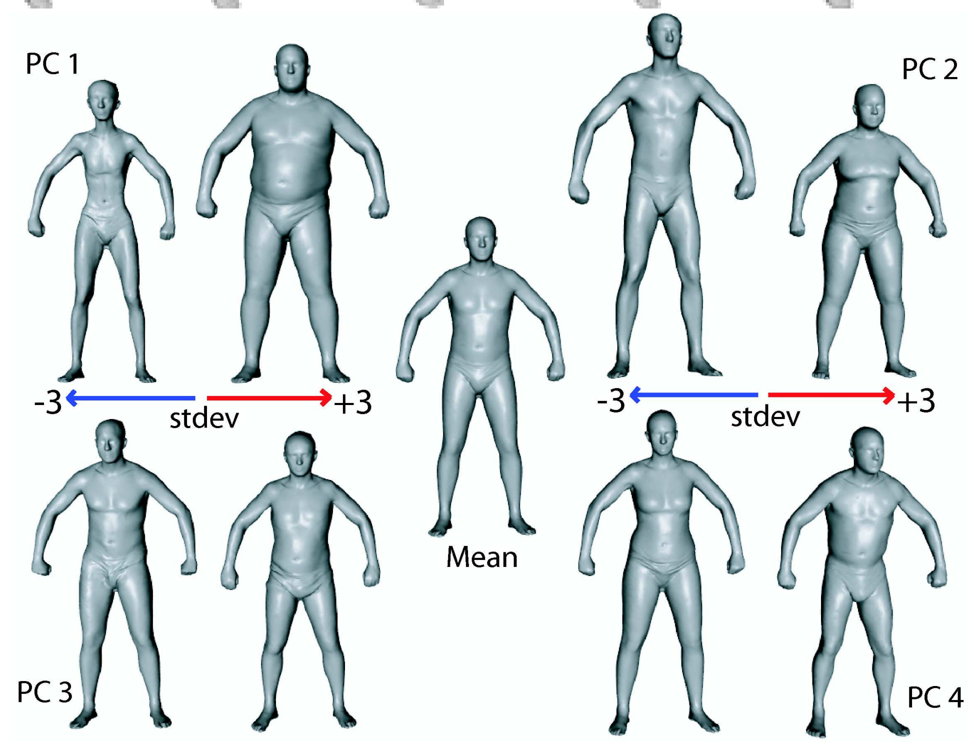

Deformable Structures

Zuffi, Freifeld, & Black, CVPR 2012

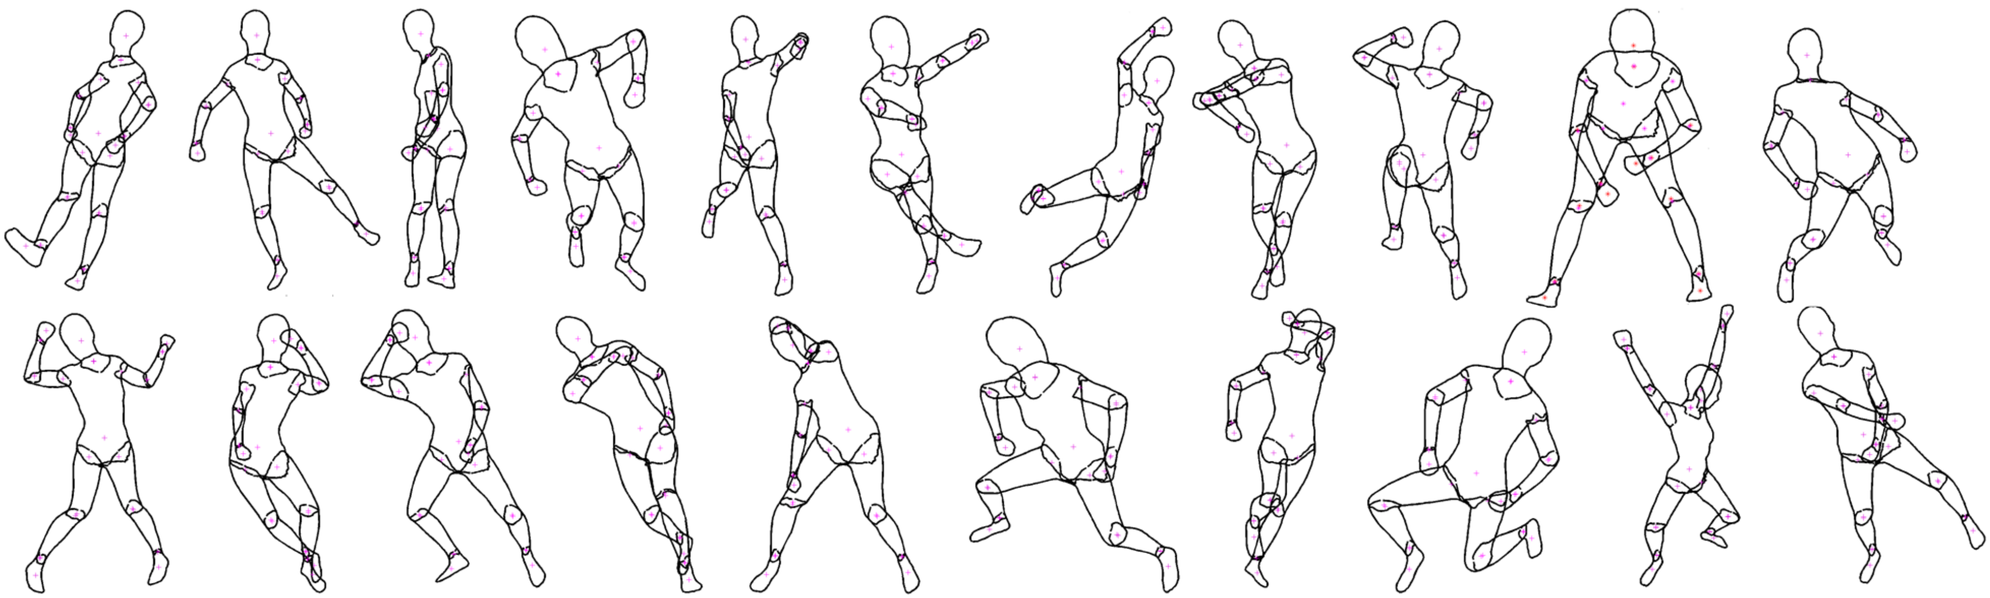

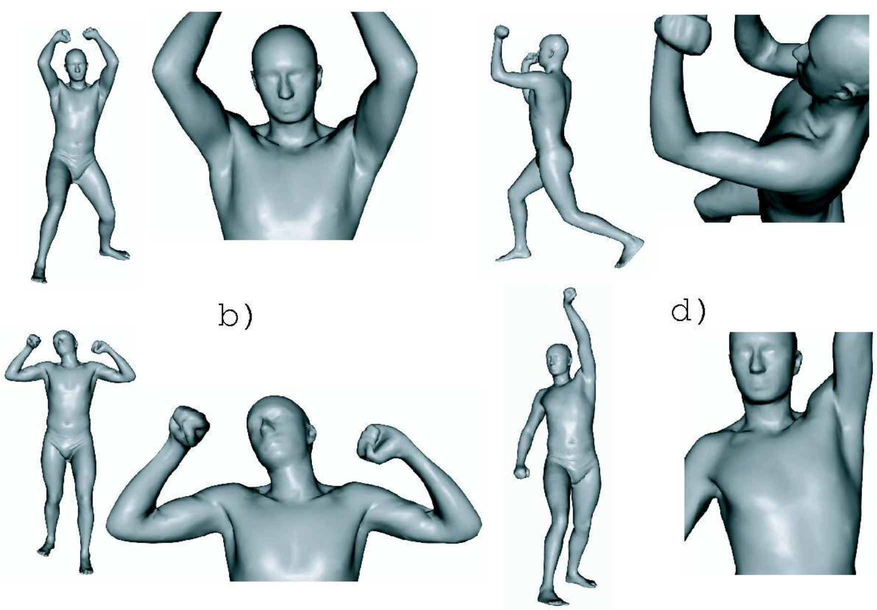

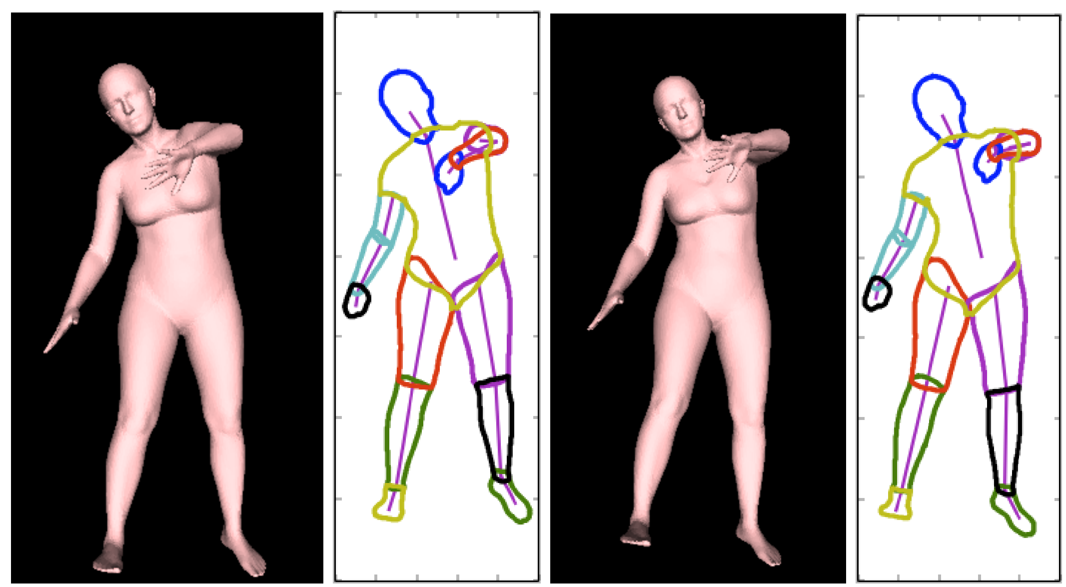

Deformable Structures

Zuffi, Freifeld, & Black, CVPR 2012

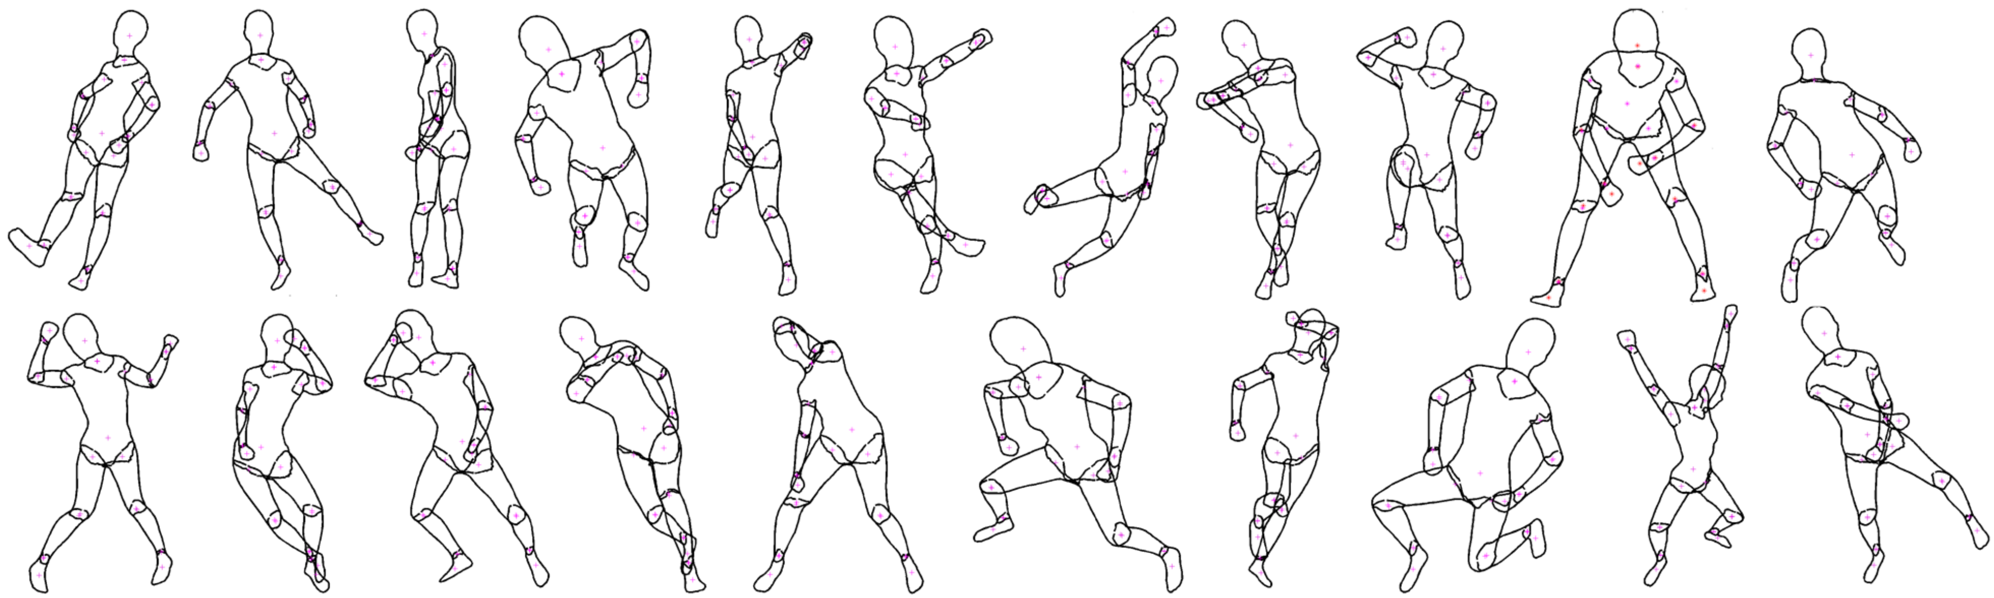

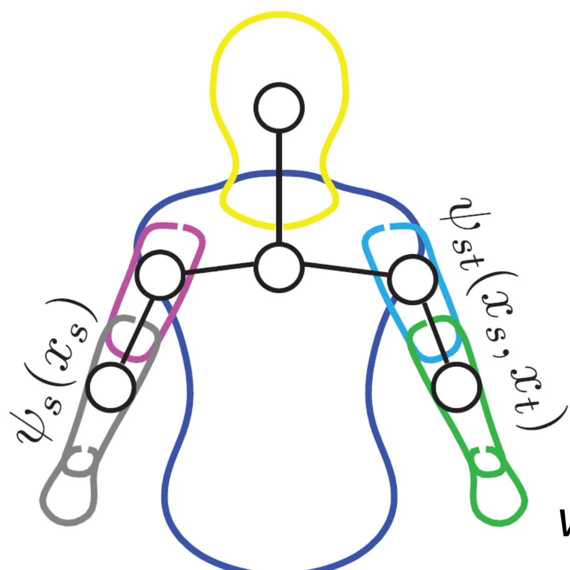

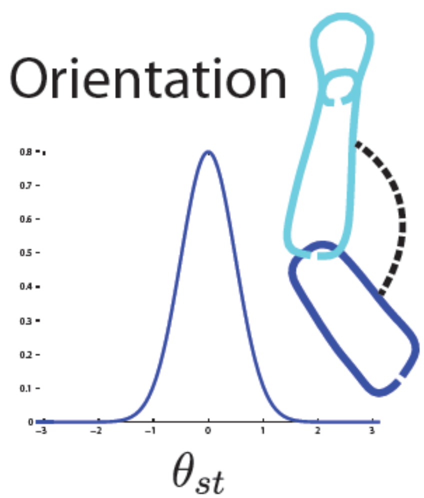

von Mises distribution on relative orientation

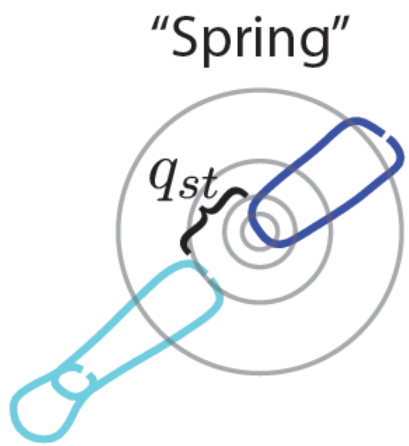

Gaussian distribution on relative position

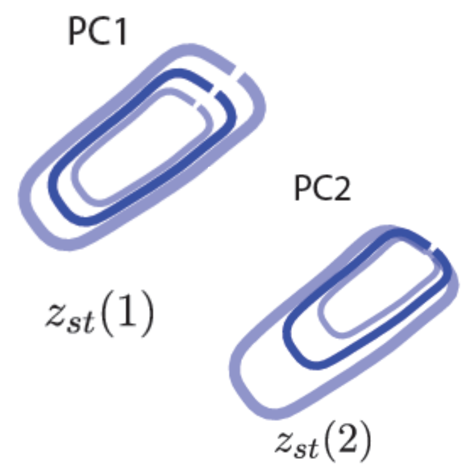

PCA model of part shape

Max-Product Belief Propagation

Discrete

$$x \in \{1, \ldots, N\}^D$$

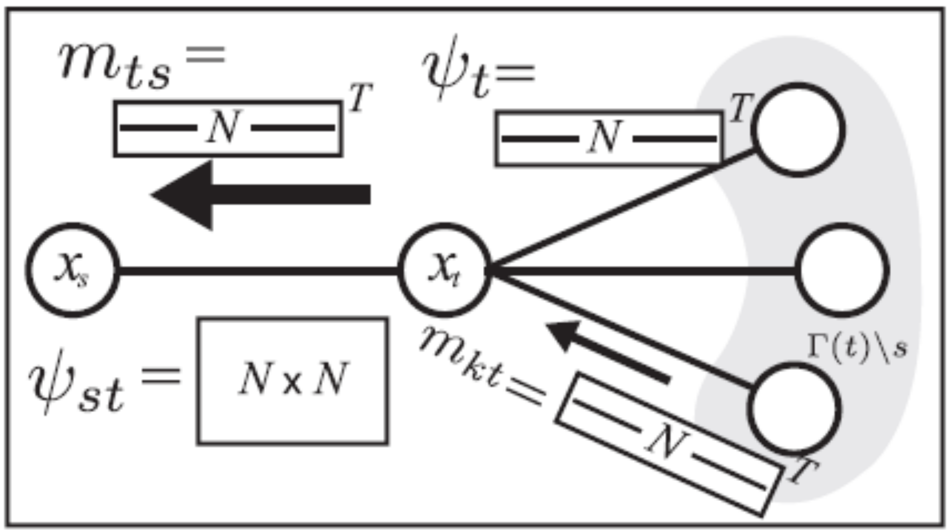

Message Update:

$$m_{ts} = \max_{x_t} \boxed{\psi_{st}} \; \boxed{\psi_t} \prod \boxed{m_{kt}}$$

Matrix-vector multiplication and discrete maximization.

Continuous

$$x \in \mathcal{R}^D$$

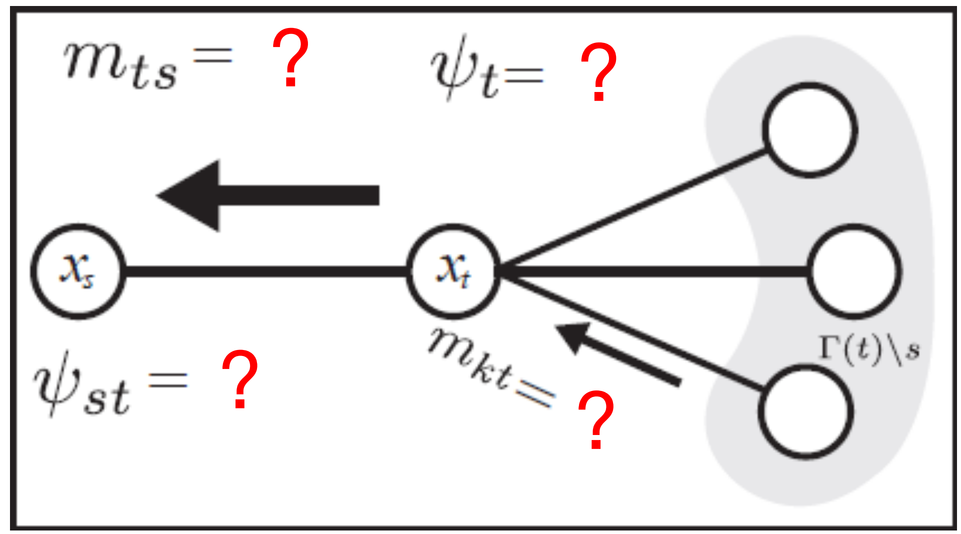

Message Update:

$$m_{ts}(x_s) = \ldots$$

$$\max_{x_t} \psi_{st}(x_s, x_t)\psi_t(x_t) \prod m_{kt}(x_t)$$

Messages are functions with no analytic form. Nonlinear optimization.

Regular Discretization Infeasible

Approximate continuous max-product messages over regular *grid* of points?

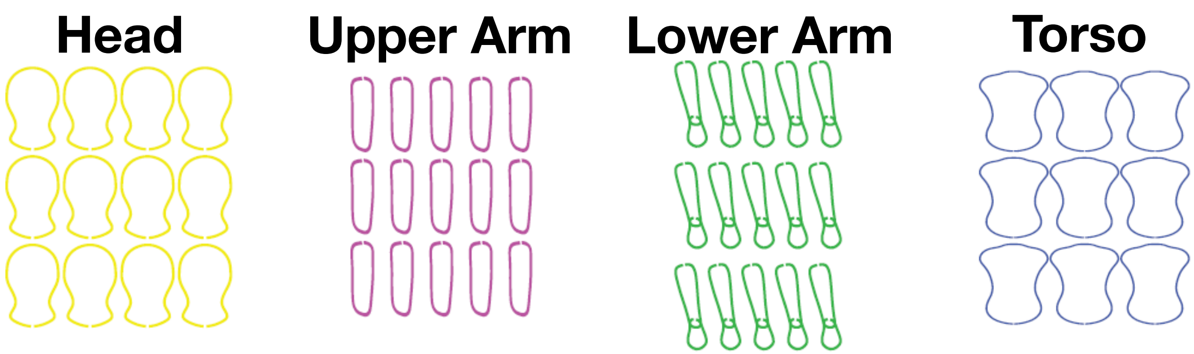

<u>Example: Torso</u>

- ~10 dimensions.
- 10 grid points per dimension
- 10 Million points!

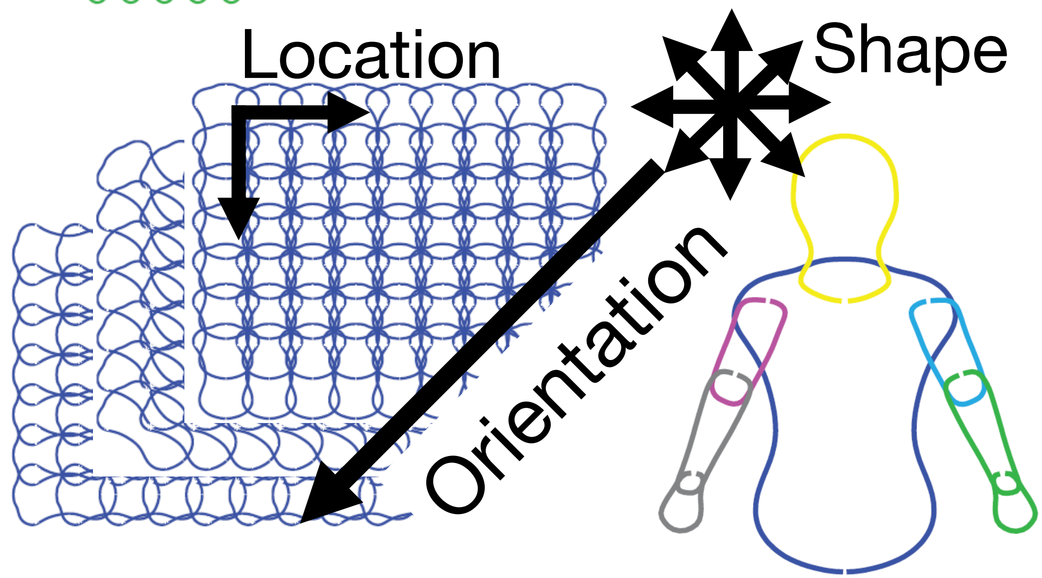

Infeasible for high dimensional models.

Pose Tracking Particle Filter?

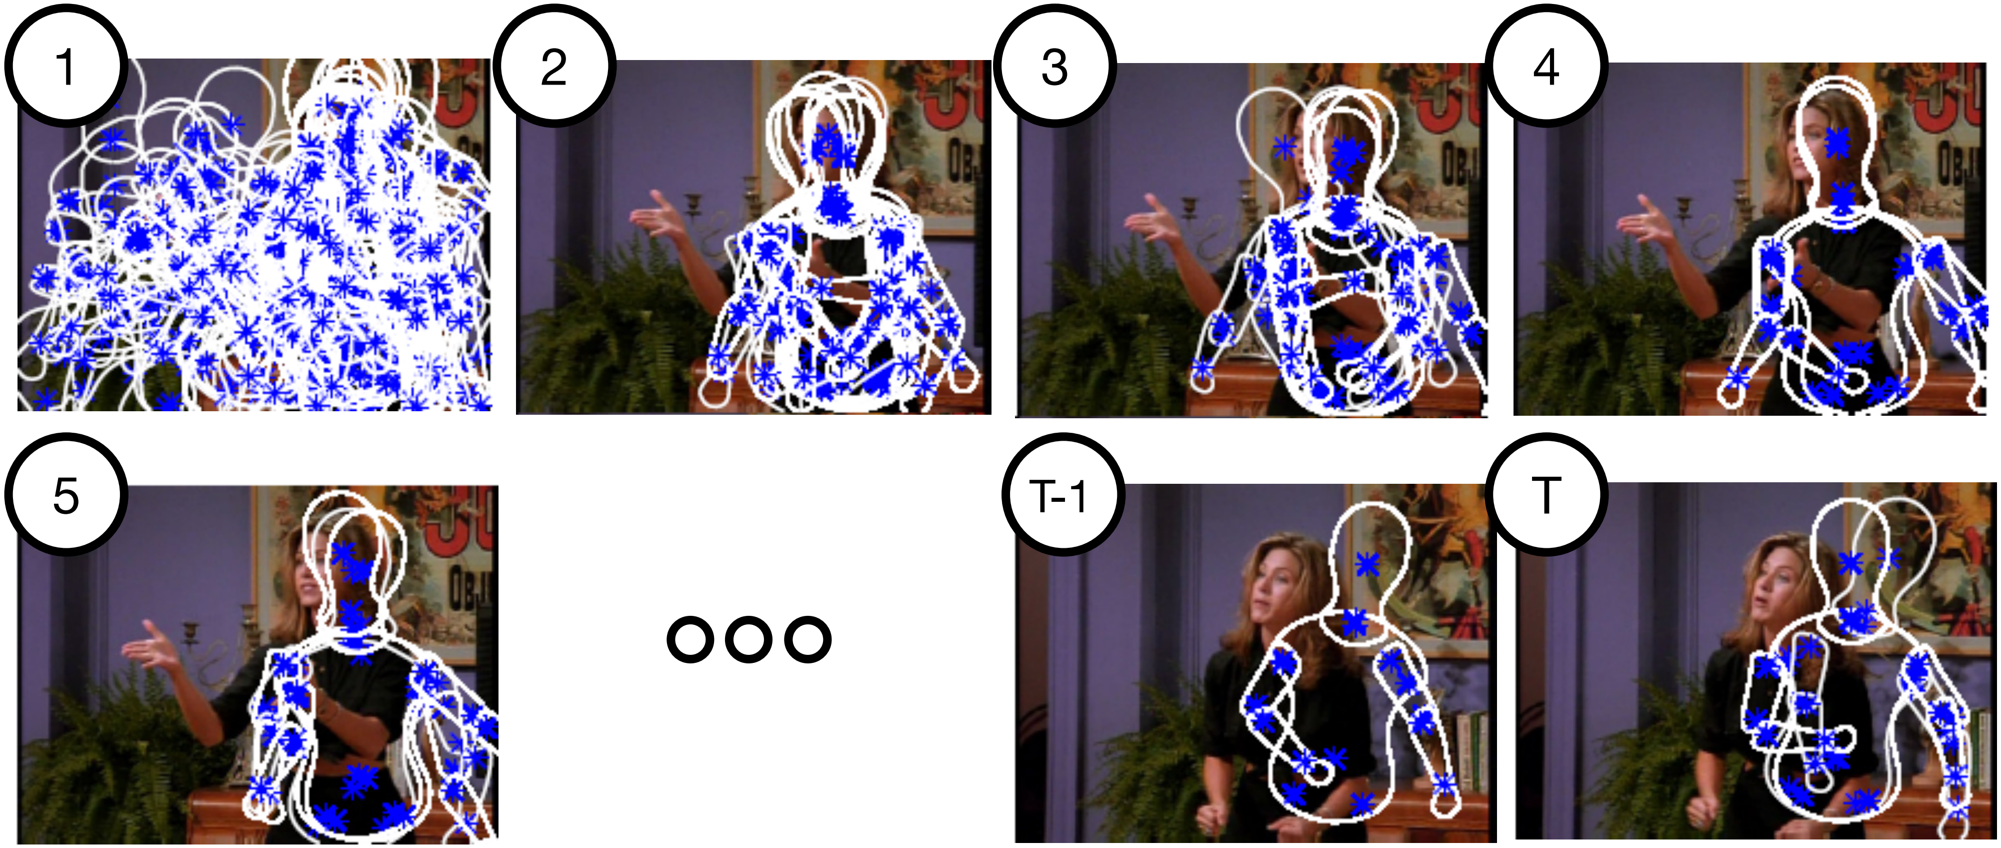

CONDENSATION algorithm [Isard & Blake, 1998]

- Particles degenerate over time
- Resampling reduces effective number of particles
- Extension beyond time series models non-trivial

Particle Representations

Particle filter:

- ➢ Each particle is a full joint instantiation

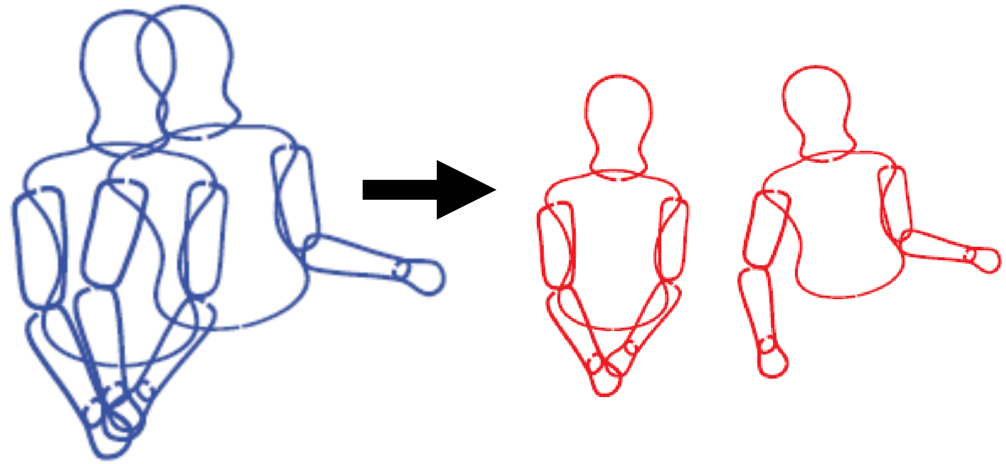

Max-Product:

- ➢ Each particle is a single variable node (part)
- ➢ Efficiently enumerates all combinations

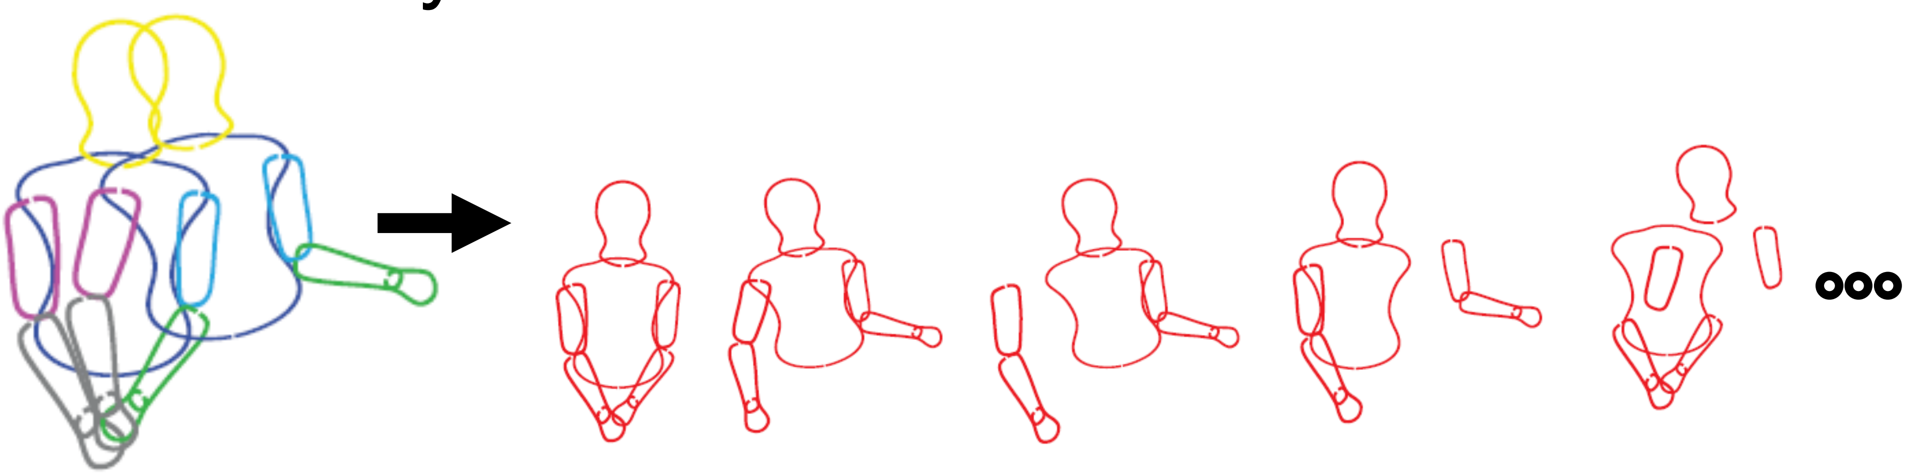

Particle Max-Product (PMP)

Combine particle filter ideas with max-product more effectively.

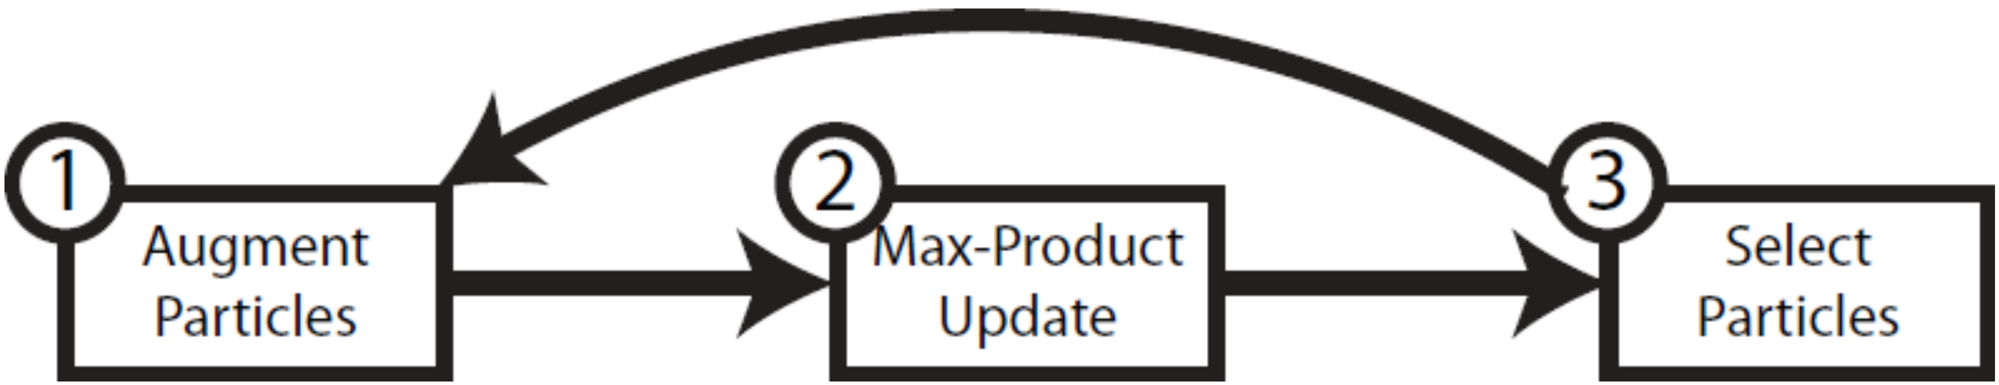

Particle approximation of continuous max-product (MP) messages.

Particle Max-Product (PMP)

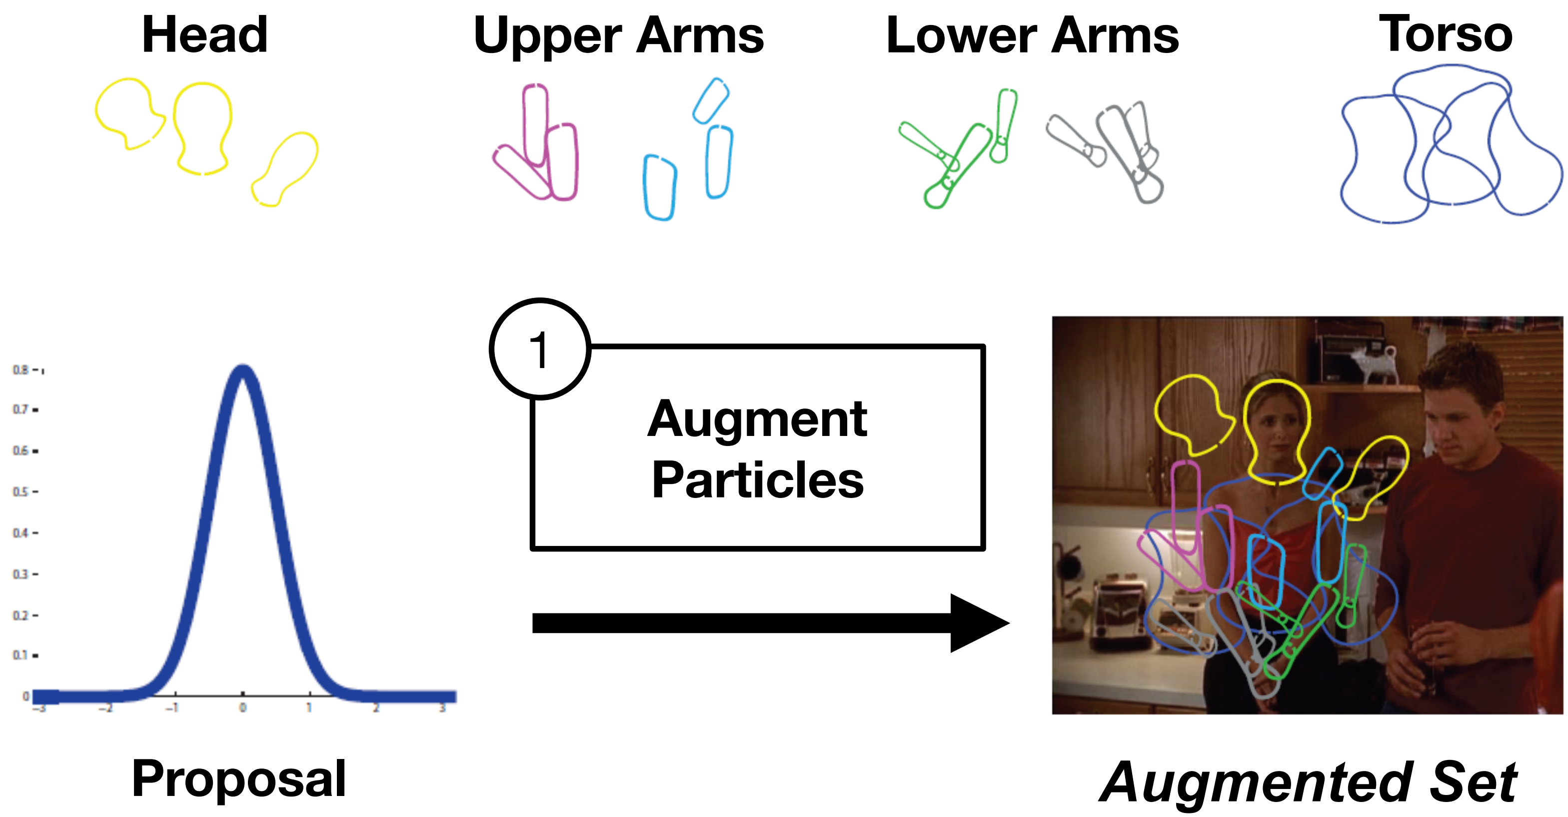

Sample new hypotheses at <u>every node</u> to grow particle set.

Particle Max-Product (PMP)

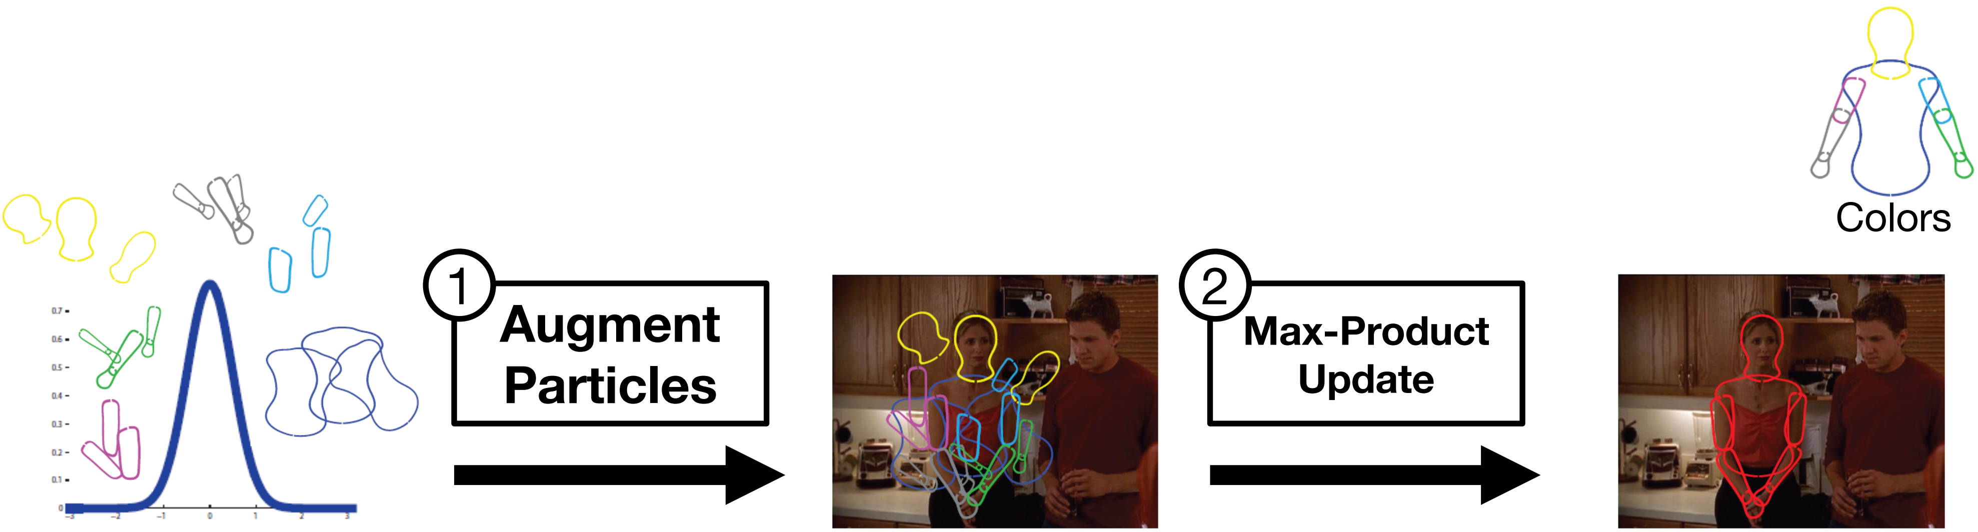

$$m_{ts}(x_s) \propto \max_{x_t} \psi_{st}(x_s, x_t)\psi_t(x_t) \prod_{k\in\Gamma(t)\setminus s} m_{kt}(x_t)$$

Update MP messages on augmented particles.

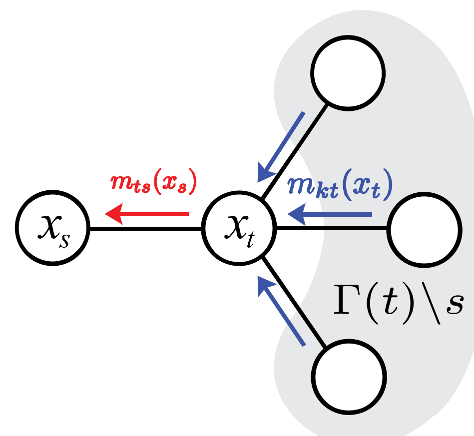

Particle Max-Product (PMP)

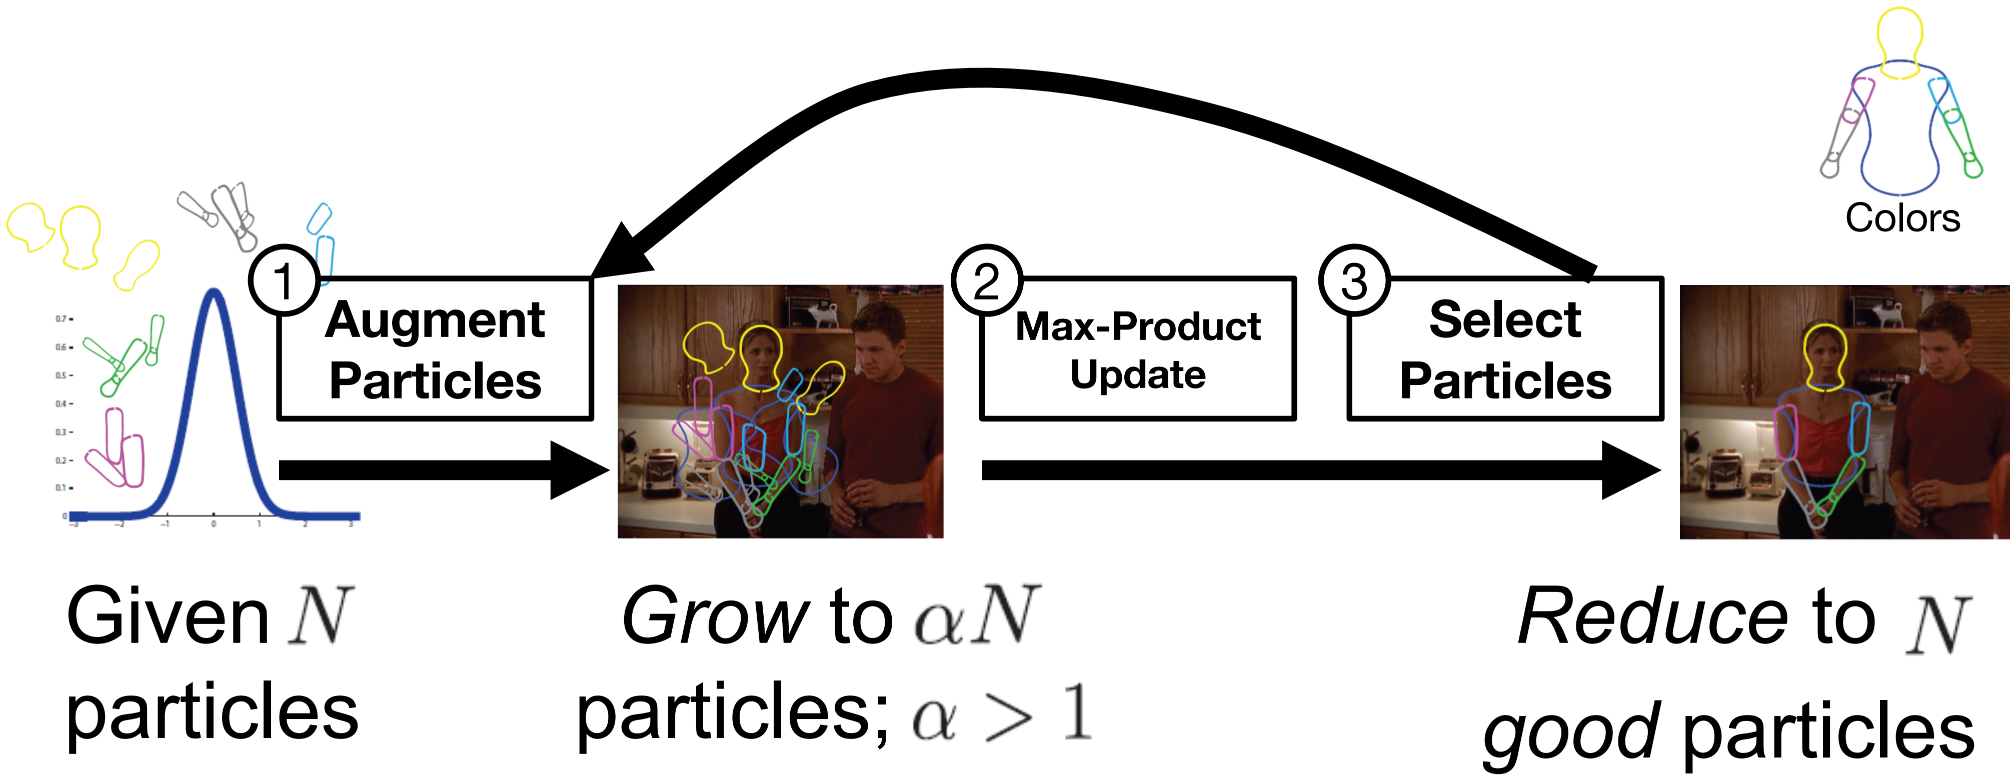

Select subset of *good* particles & repeat

Need a particle selection method…

Deformable Structures for Silhouettes

ICML

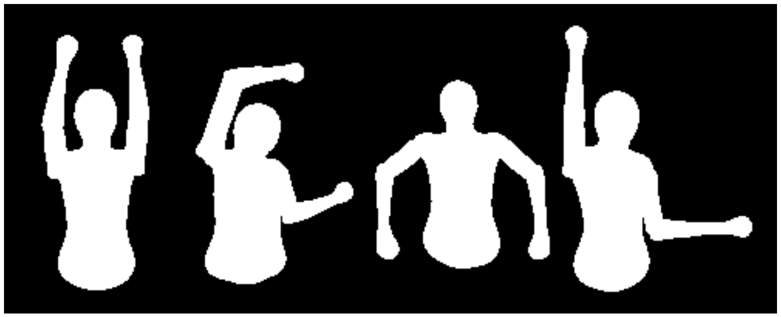

Chamfer Distance Likelihood

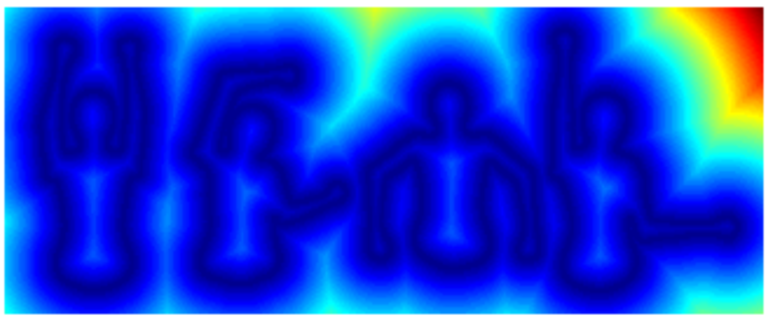

Random Initialization

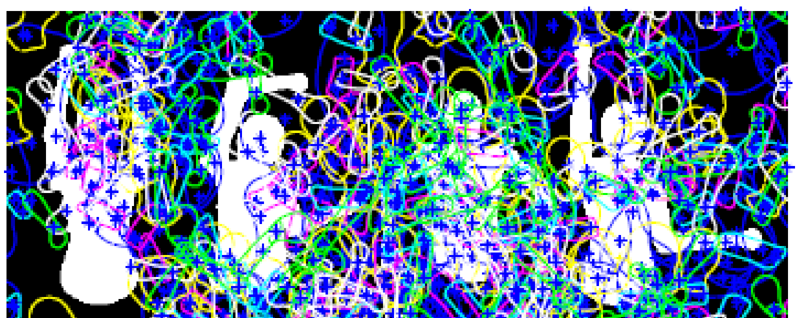

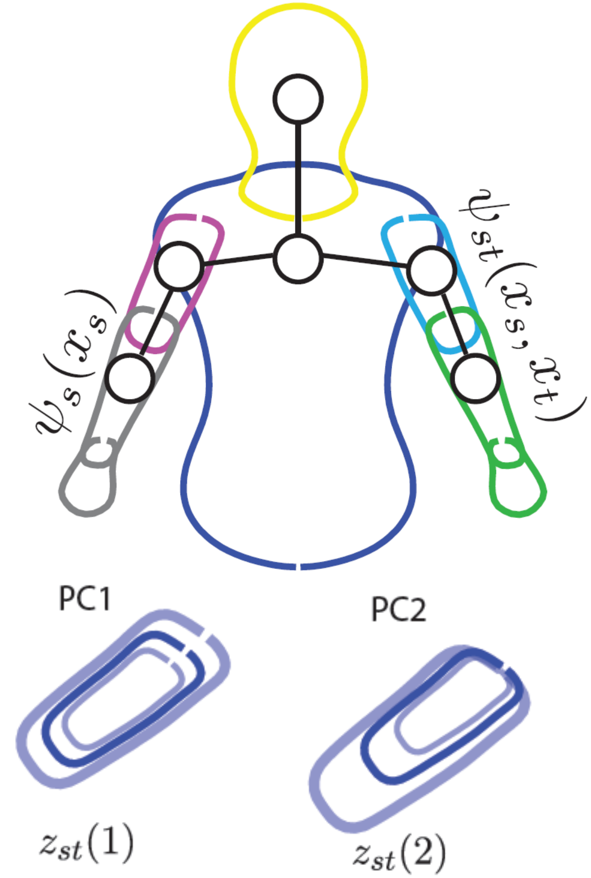

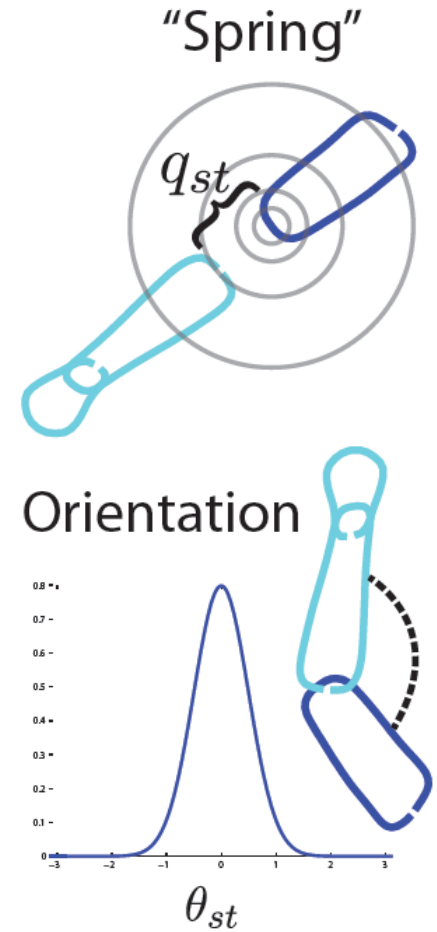

Inference Goals:

- Accurately localize all 4 people
- Reliably find global MAP (the "M")

Greedy Particle Max-Product

G-PMP: *Trinh & McAllester 2009*

- **Select:** Discard all current particles except "MAP"
- **Augment:** Propose new particles by perturbing MAP (Gaussian "random walk")

Example Runs

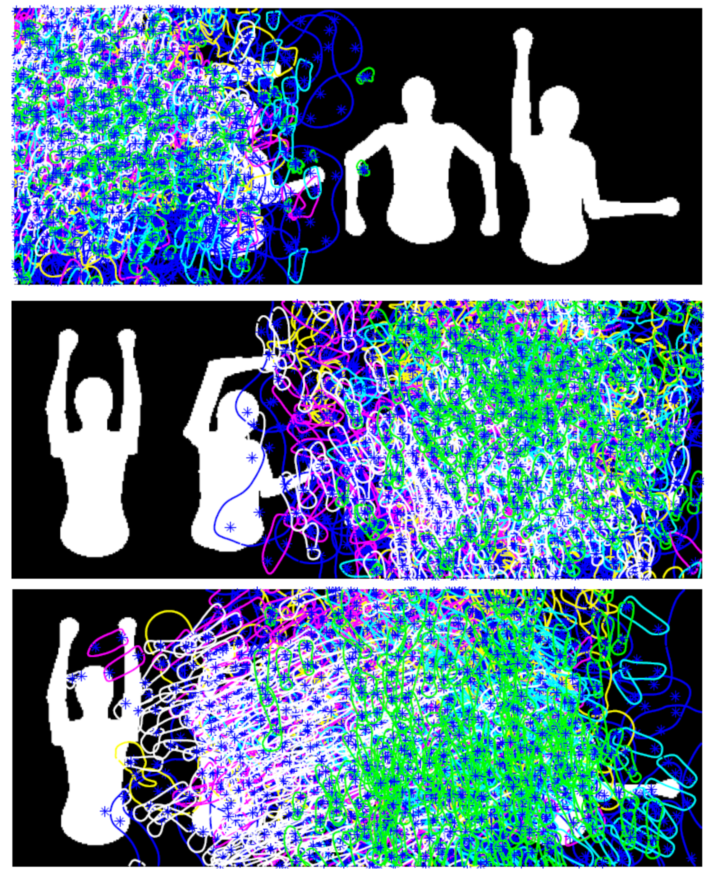

Colors

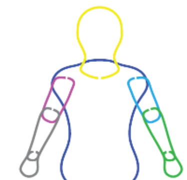

Particles *degenerate* to a single mode. Discovered mode is *very sensitive to initialization*, and is often not the true MAP.

Top-Mode Particle Max-Product

T-PMP: *Generalization of PatchMatch BP, Besse et al. 2012*

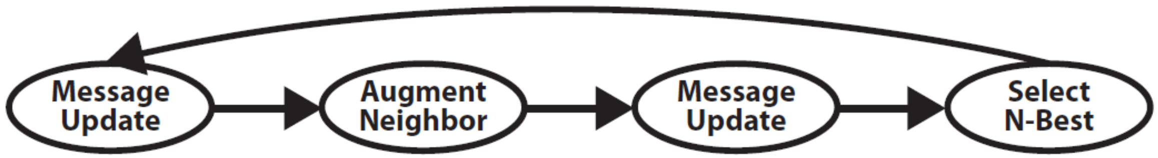

Example Runs

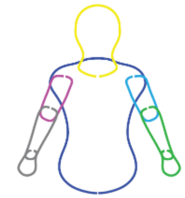

Colors

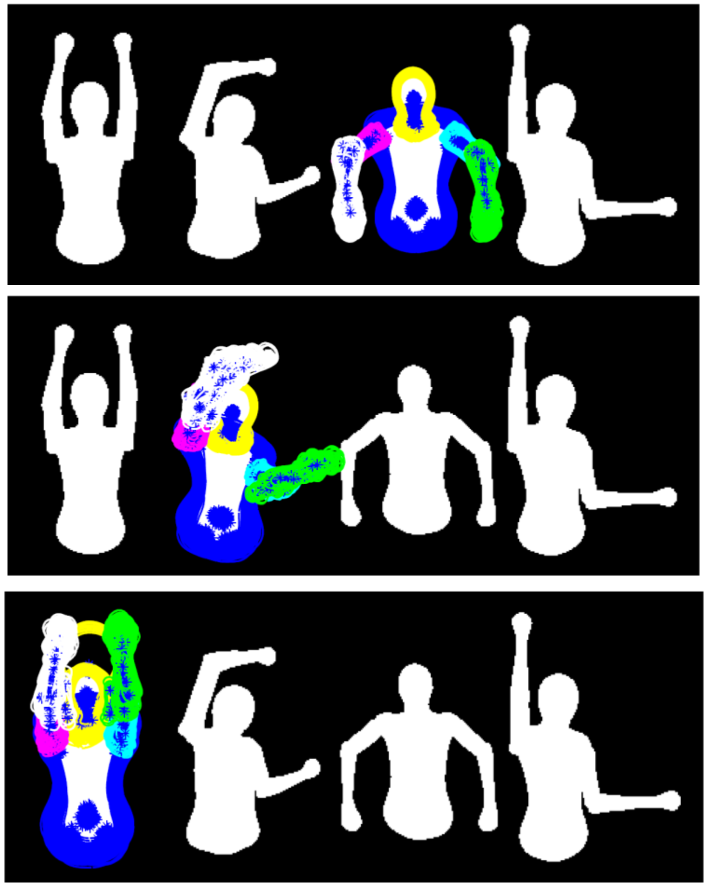

- **Augment:** Propose new particles from neighbors
- **Select:** Sort max-marginals and keep top *N* particles

Particles *degenerate* to a single mode. Discovered mode is *sensitive to initialization*, and is often not the true MAP.

Diverse Particle Selection

GOAL: Maintain *diversity* in particles.

Initial Particles

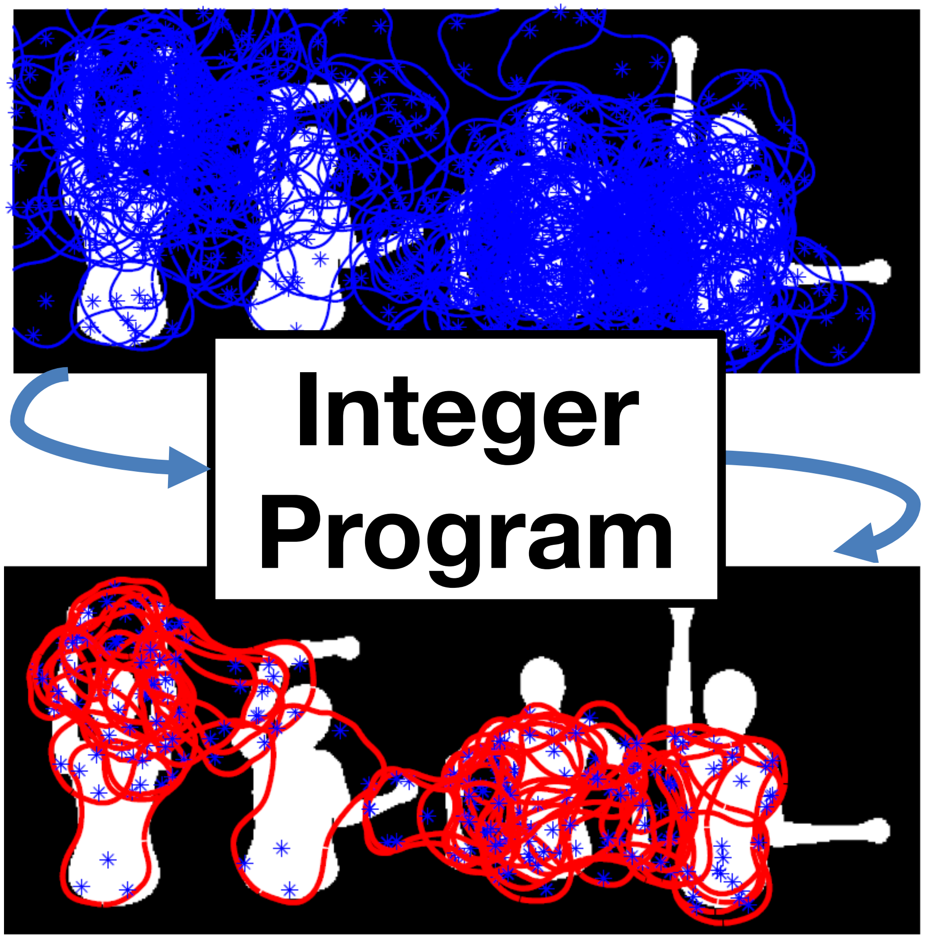

Diverse Selection

Integer Program (IP) solved with efficient greedy approximation:

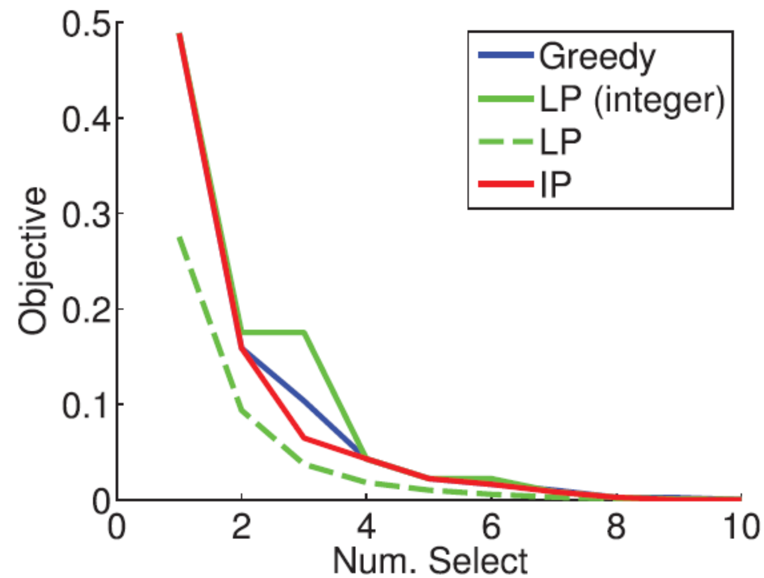

LP : Linear Program relaxation

IP: Optimal solution by brute force

Greedy: Efficient approximation

Continuous Message

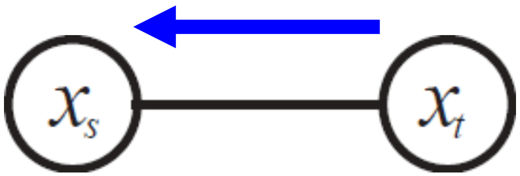

$$m_{ts}(x_s) = \max_{x_t} \psi_{st}(x_s, x_t)\psi_t(x_t)$$

Joint Distribution

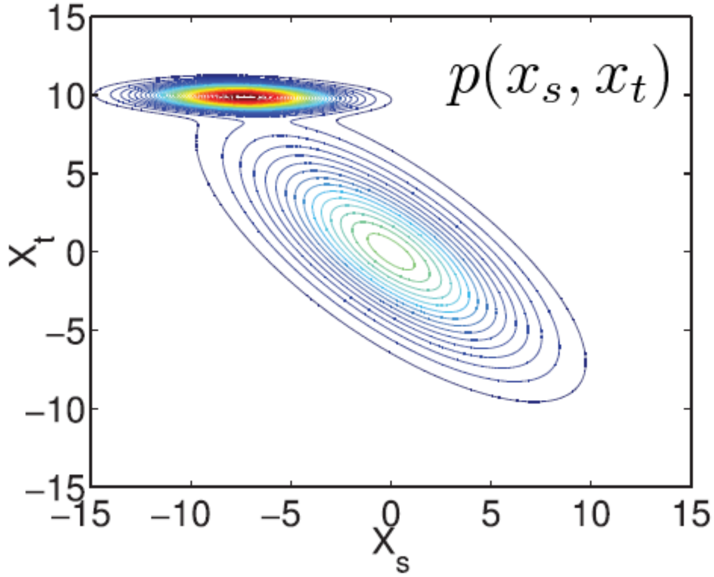

Message

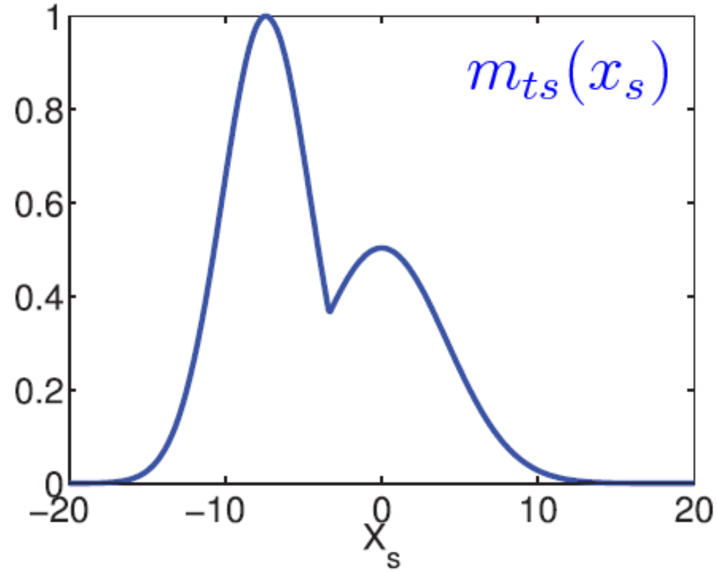

Model is a mixture of 2 Gaussians.

Discrete Message

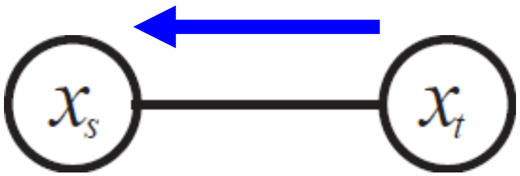

$$m_{ts} = \max_{x_t} \boxed{\psi_{st}} \; \psi_t$$

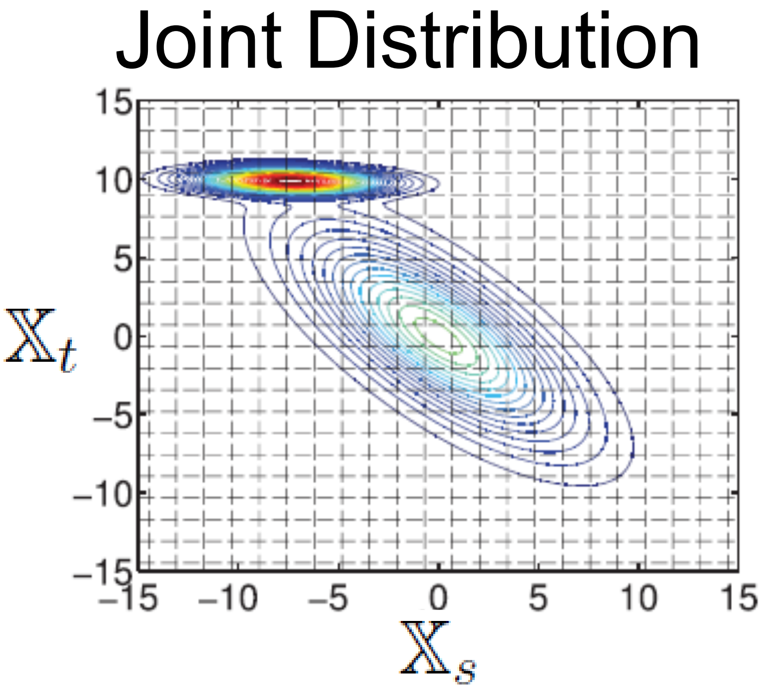

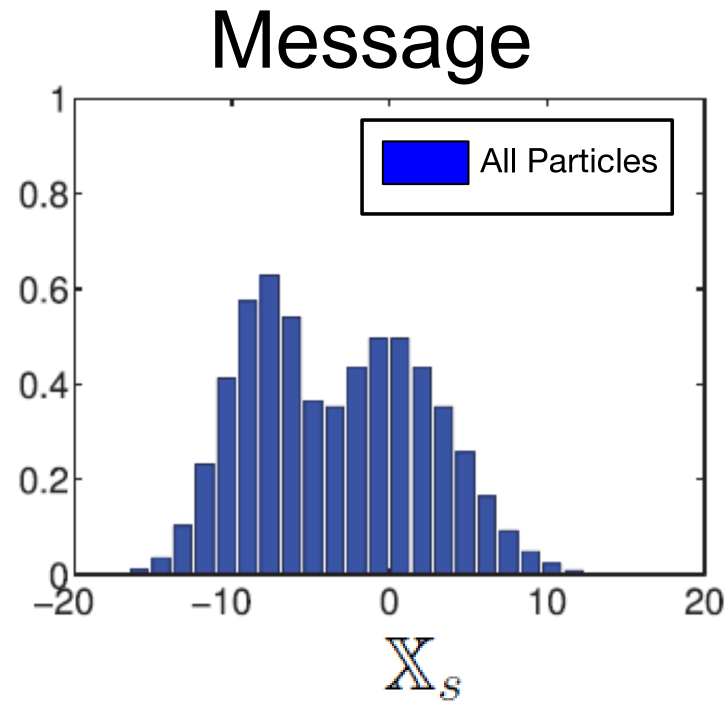

Regular grid of 50 states gives discretization:

$$\mathbb{X} = \{x^{(1)}, \ldots, x^{(50)}\}$$

Particle Selection

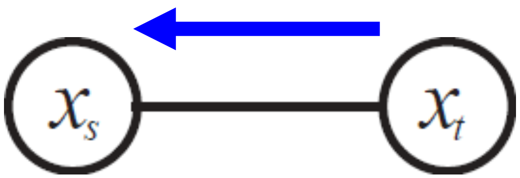

$$\hat{m}_{ts}(z) = \max_{x_t} \boxed{z} \boxed{\psi_{st}} \left[\psi_t\right]$$

Selection vector

Joint Distribution

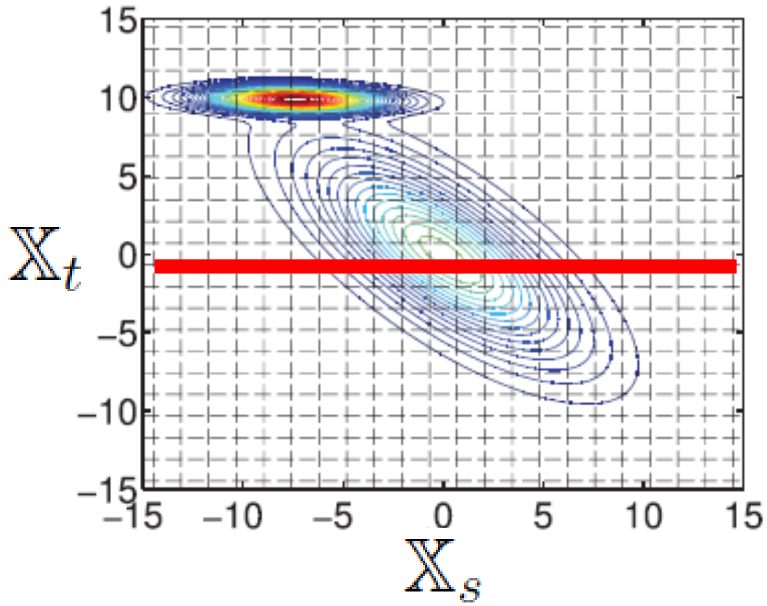

Message

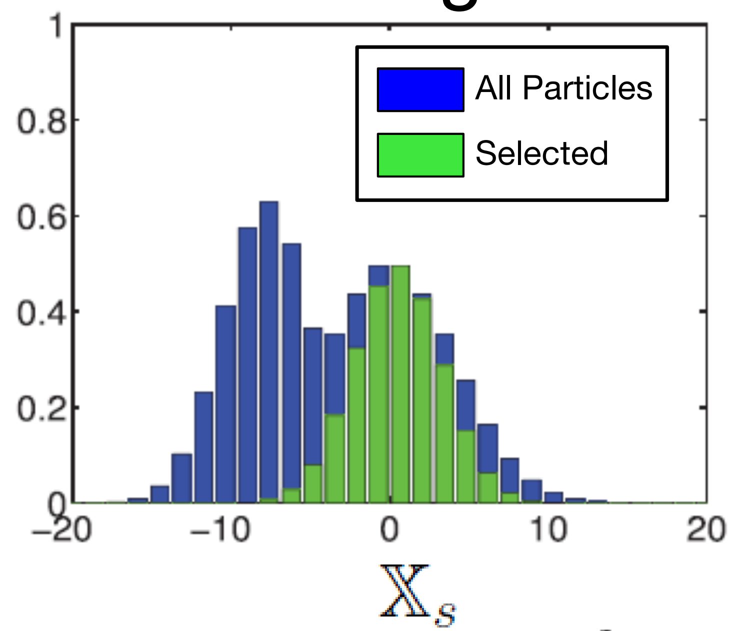

- Indicator vector controls state selection: $z \in \{0,1\}^{50}$
- $z(i) = 1$ indicates selected states (red line)

Particle Selection

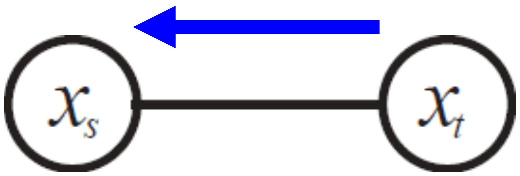

$$\hat{m}_{ts}(z) = \max_{x_t} \boxed{z} \boxed{\psi_{st}} \left[\psi_t \right]$$

Joint Distribution

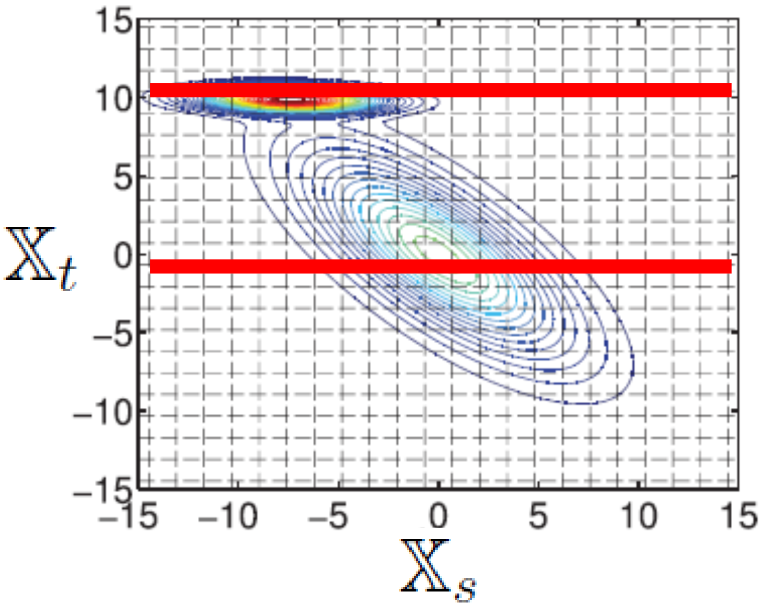

Message

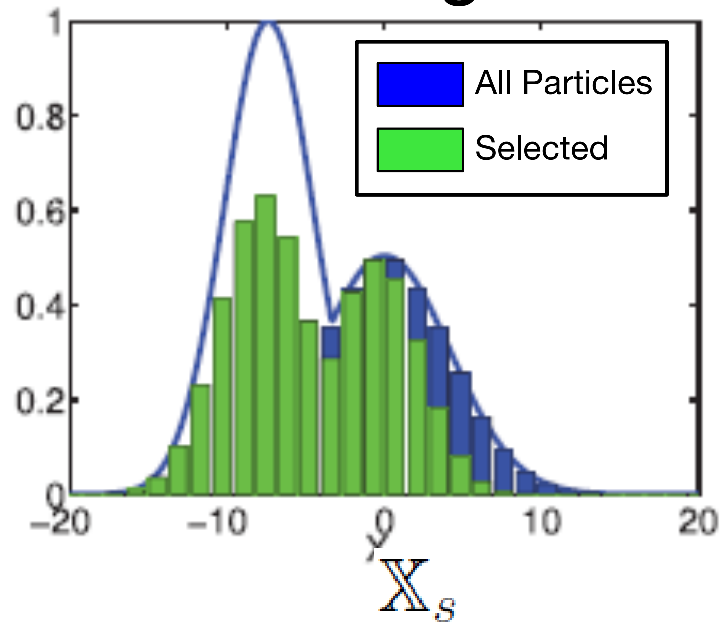

Adding states reduces distortion between discrete message vectors.

Diverse Particle Selection

Minimize total message distortion:

$$\underset{z}{\text{minimize}} \sum_{s \in \Gamma(t)} \sum_{a=1}^{\alpha N} (m_{ts}(a) - \hat{m}_{ts}(a, z))$$

$$\text{subject to } \|z\|_1 \leq N, \; z \in \{0, 1\}^{\alpha N}$$

✗ **NP-hard**

✓ **Submodular**

Good approximation qualities.

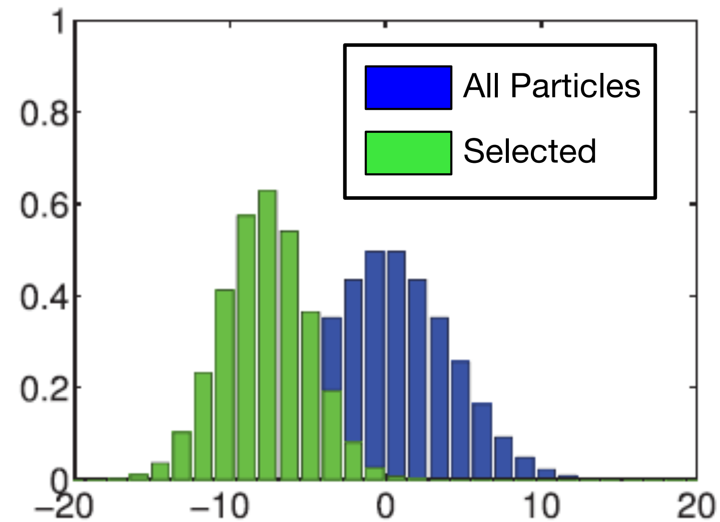

Submodularity

Set function $f : 2^Z \to \mathbb{R}$ is submodular iff diminishing marginal gains.

$$f(Y \cup \{e\}) - f(Y) \geq f(X \cup \{e\}) - f(X)$$

Margin

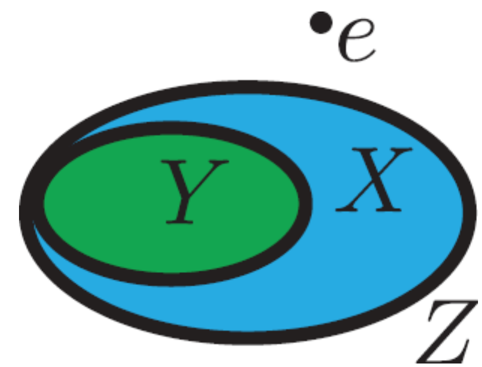

Diverse particle selection IP equivalent to submodular maximization.

- Efficient greedy approximation
- Within $(1 - 1/e) \approx 63\%$ of optimal

Greedy Particle Selection

$$\underset{z}{\text{minimize}} \sum_{s \in \Gamma(t)} \sum_{a=1}^{\alpha N} \left(m_{ts}(a) - \hat{m}_{ts}(a, z) \right)$$

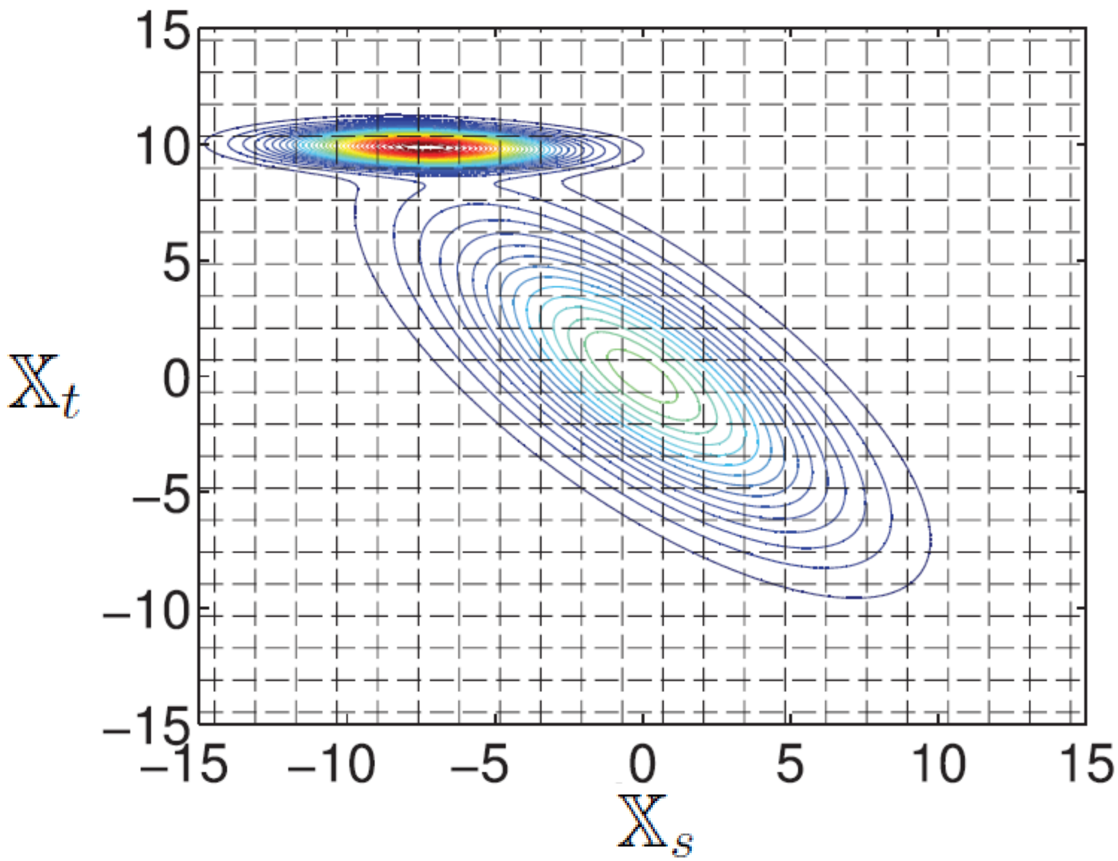

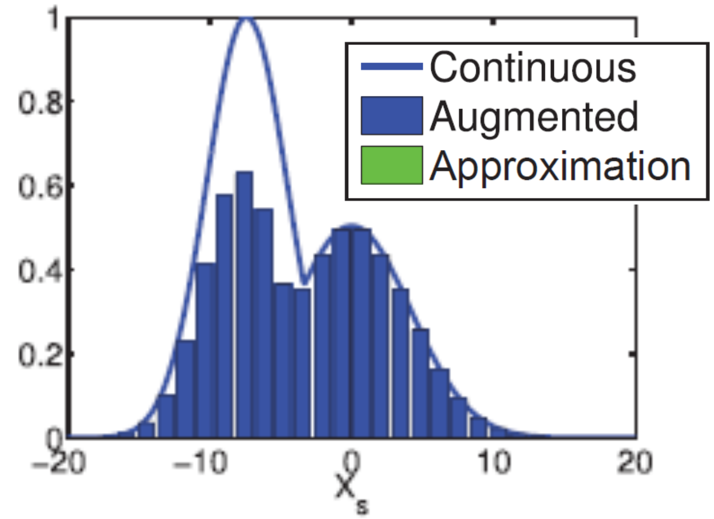

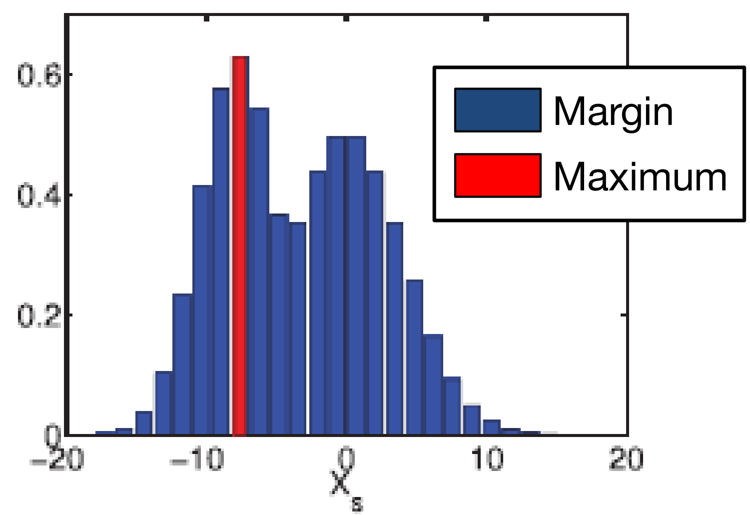

Greedy Particle Selection

$$\underset{z}{\text{minimize}} \sum_{s \in \Gamma(t)} \sum_{a=1}^{\alpha N} \left(m_{ts}(a) - \hat{m}_{ts}(a, z) \right)$$

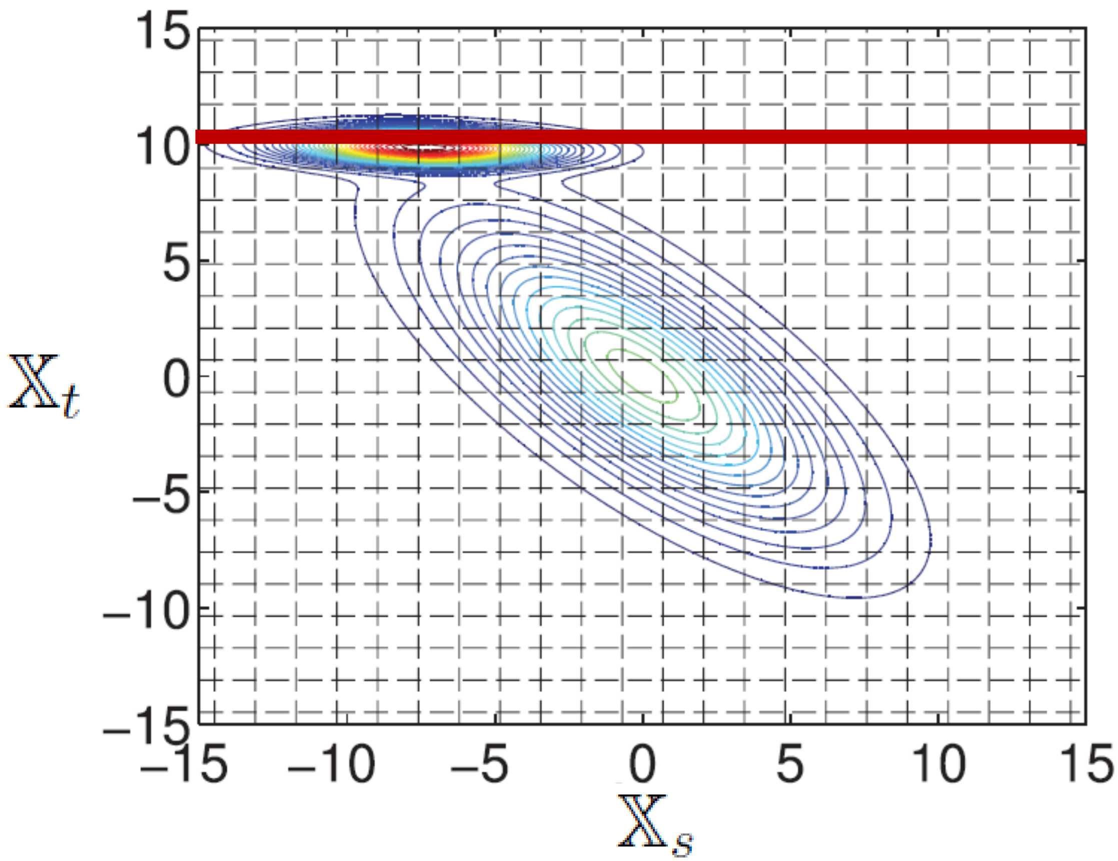

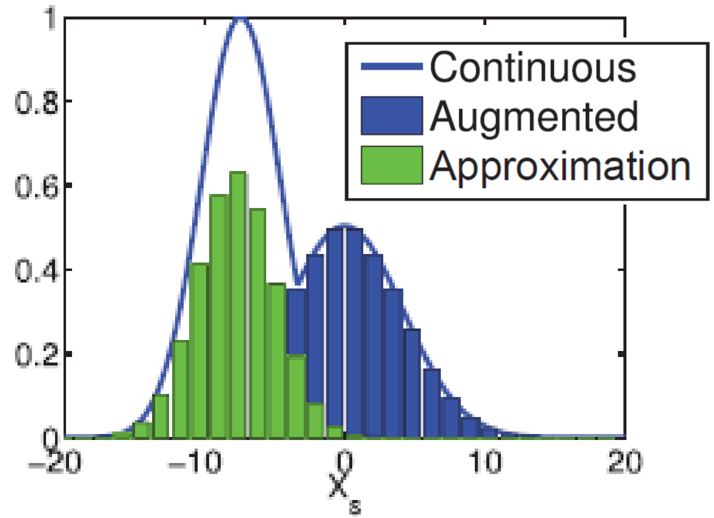

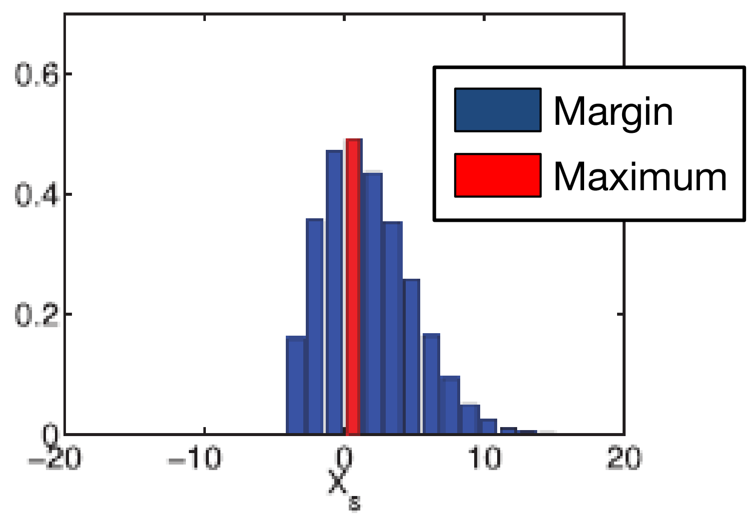

Greedy Particle Selection

$$\underset{z}{\text{minimize}} \sum_{s\in\Gamma(t)} \sum_{a=1}^{\alpha N} \left(m_{ts}(a) - \hat{m}_{ts}(a,z) \right)$$

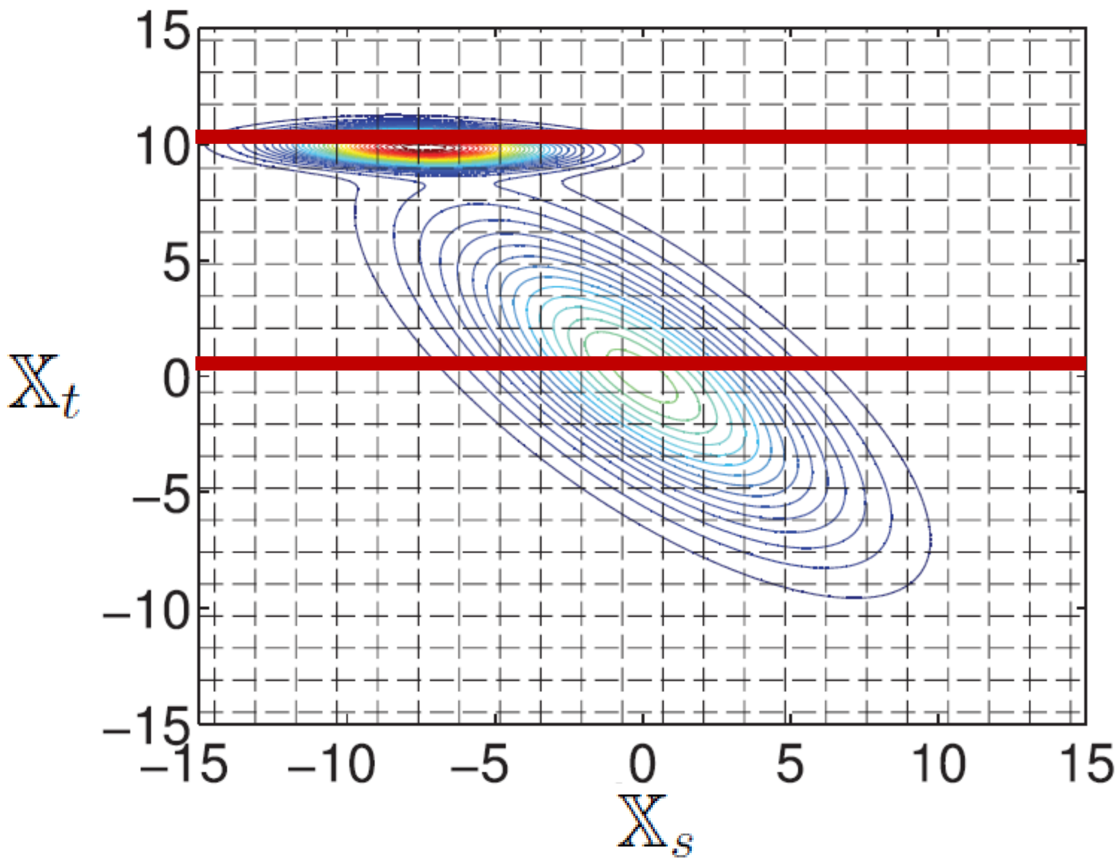

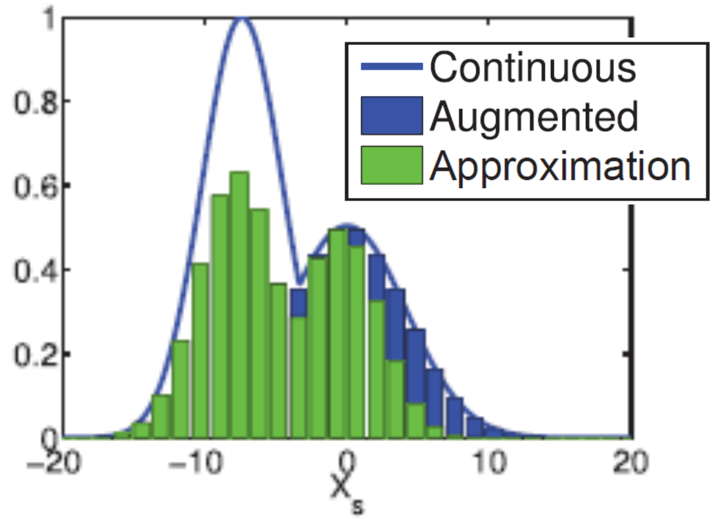

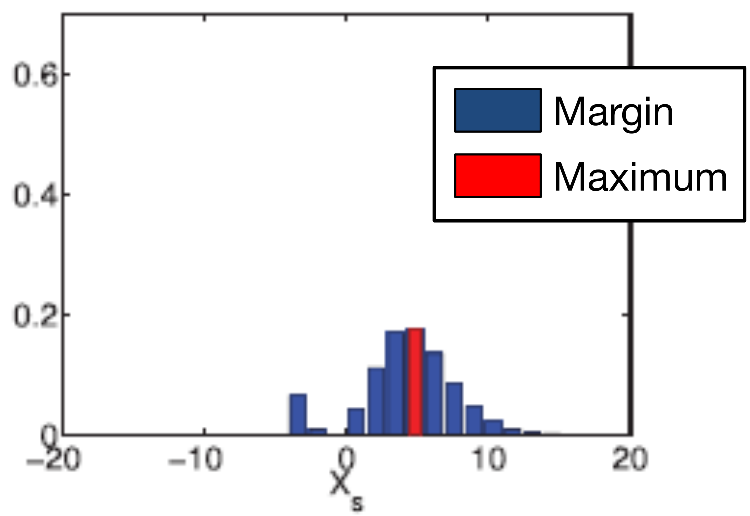

Greedy Particle Selection

$$\underset{z}{\text{minimize}} \sum_{s \in \Gamma(t)} \sum_{a=1}^{\alpha N} \left(m_{ts}(a) - \hat{m}_{ts}(a, z) \right)$$

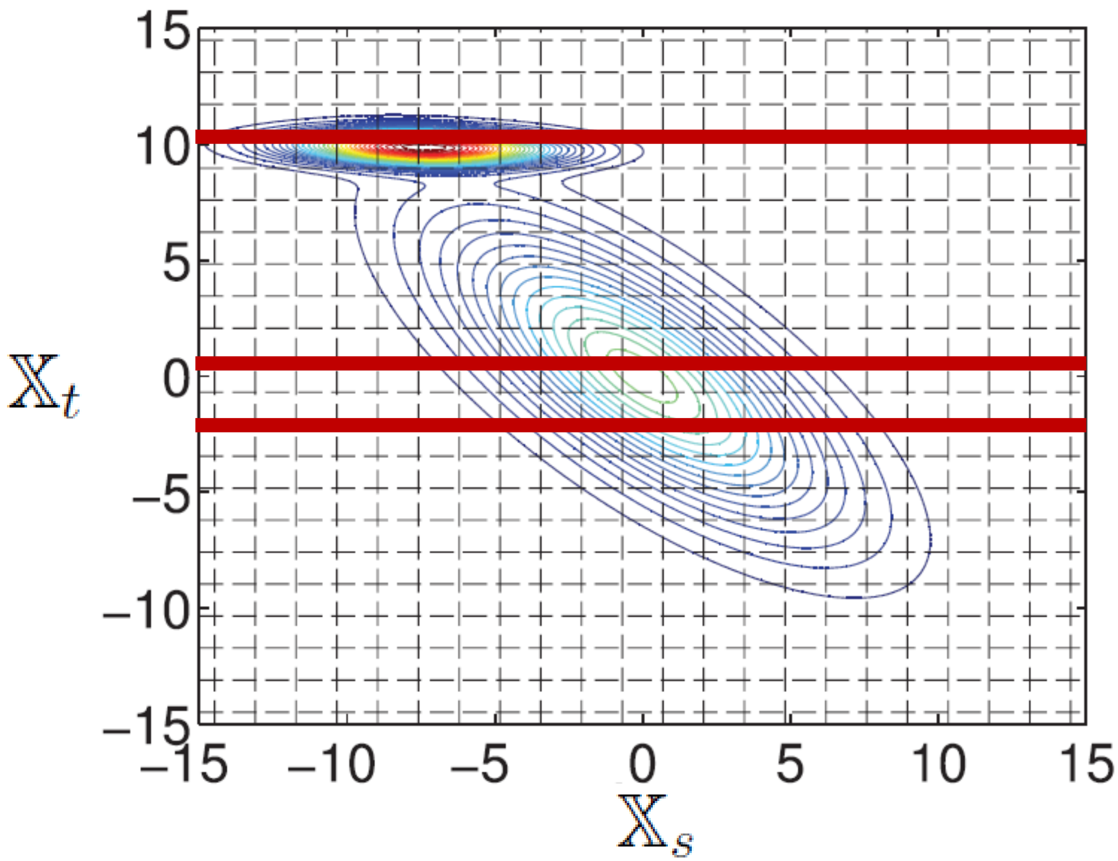

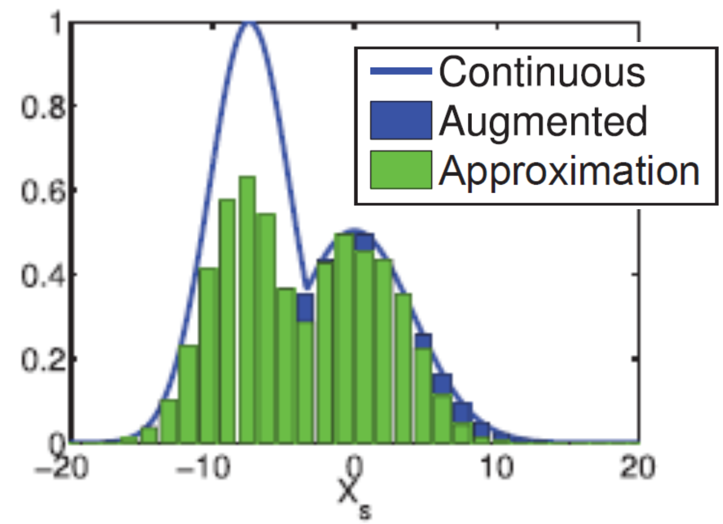

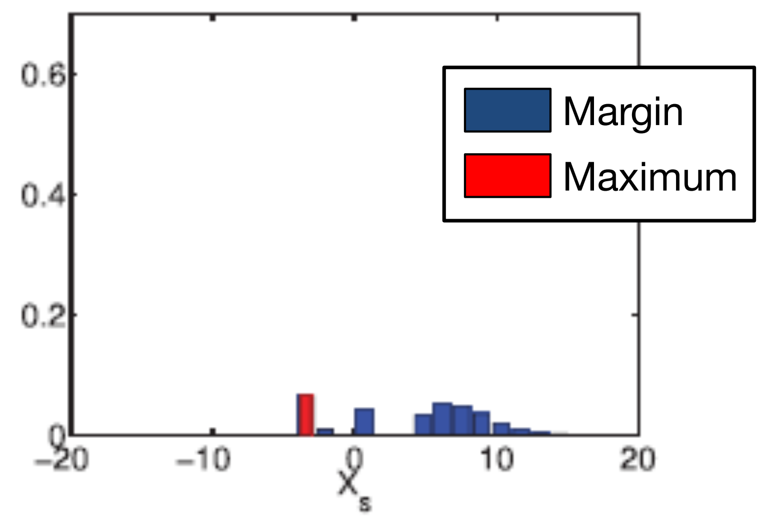

Diverse Particle Max-Product (D-PMP)

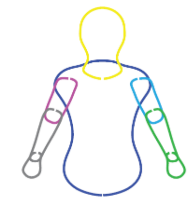

Example Runs

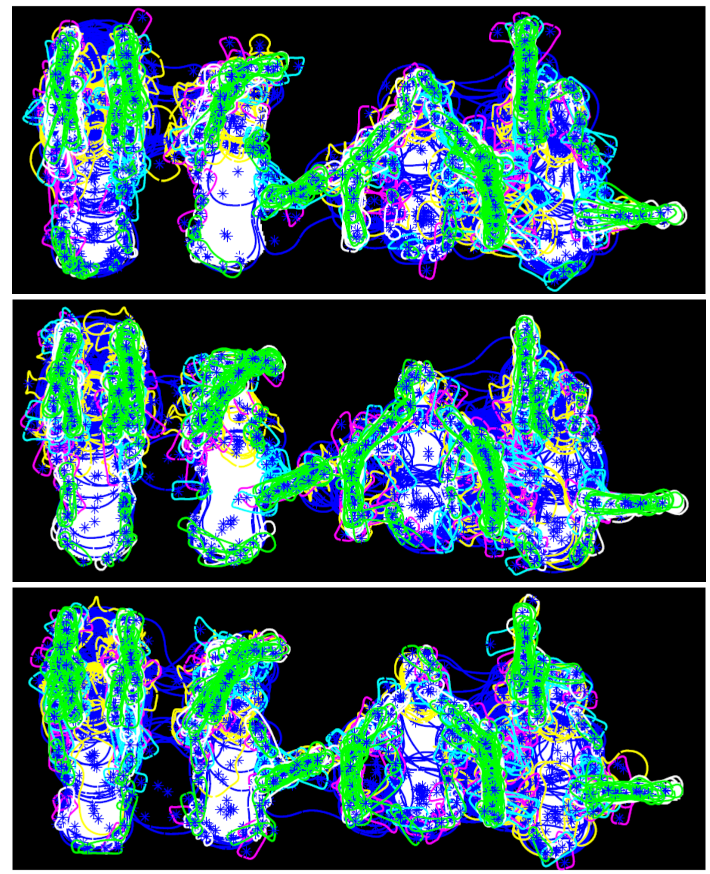

- ➢ No explicit diversity constraint
- ➢ Objective encourages diversity
- ➢ Efficient *"lazy"* greedy algorithm
- ➢ Bounds on optimality

Avoids particle degeneracies by maintaining *ensemble* of *diverse solutions* near local modes.

[Pacheco et al., ICML 2014]

Discovering Multiple Hypotheses

Prior Work Specialized to Discrete Graphical Models

M-Best MAP *[Nilsson 1998; Yanover and Weiss 2003]*

- ➢ Produce *M* solutions with highest joint probability
- ➢ Typically, these are minor variations of a single mode

Diverse M-Best MAP *[Batra et al. 2012]*

- ➢ Externally specified metric used to find probable hypotheses separated by some distance threshold
- ➢ Specialized to discrete models, and requires tuning of metrics/thresholds for each graphical model

Diverse Particle Max-Product

- ➢ Tractable for high-dimensional state spaces
- ➢ Notion of "distance" arises automatically from model

Synthetic Images: ICML Puppets

True MAP

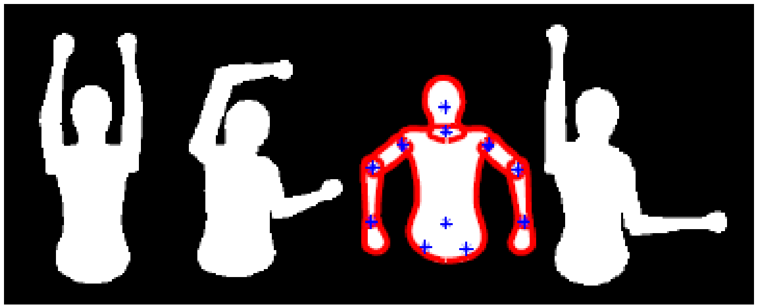

Random Initialization

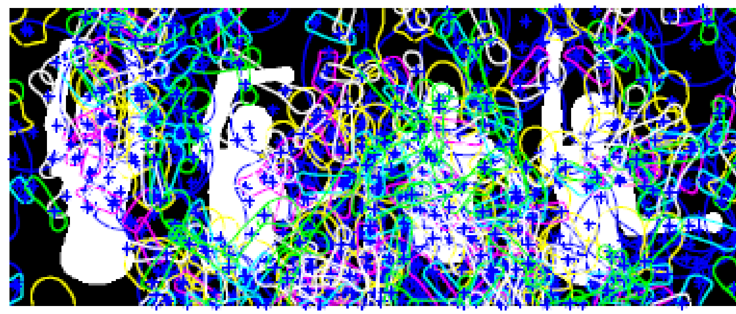

Pose Error of MAP Estimate

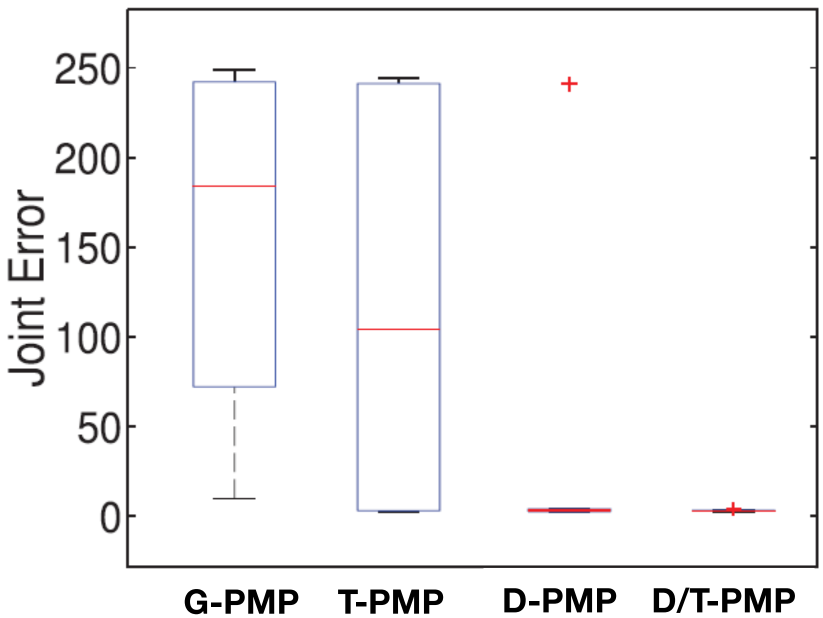

Log Probability of MAP Estimate

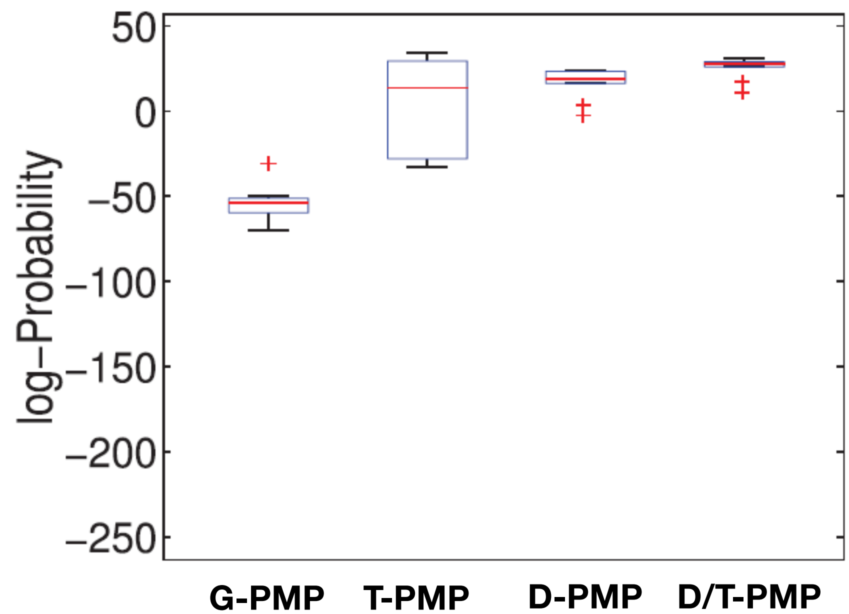

Box plots summarize results from 10 random initializations.

Real Images (Single Person)

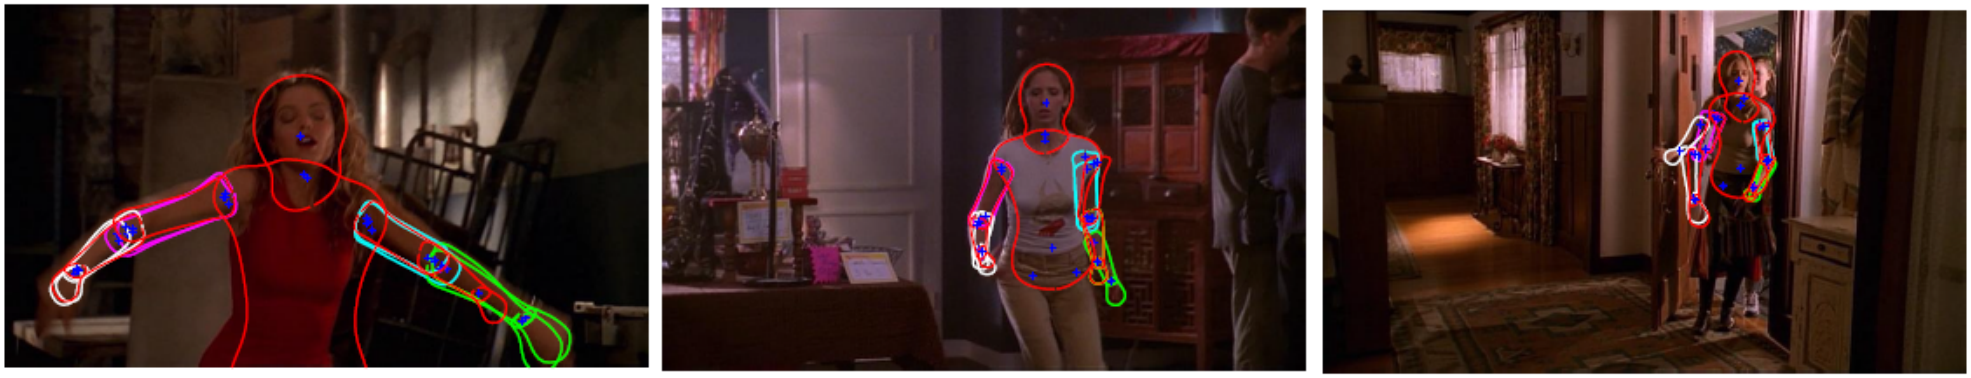

Top 3 arm hypotheses MAP estimate, 2nd and 3rd modes for upper arm (magenta, cyan), lower arm (green, white).

- “Buffy” dataset [Ferrari et al. 2008].
- Detections versus number of ranked hypotheses.
- Baseline: Flexible Mixture of Parts (FMP) [Yang & Ramanan 2013; Park & Ramanan 2011]

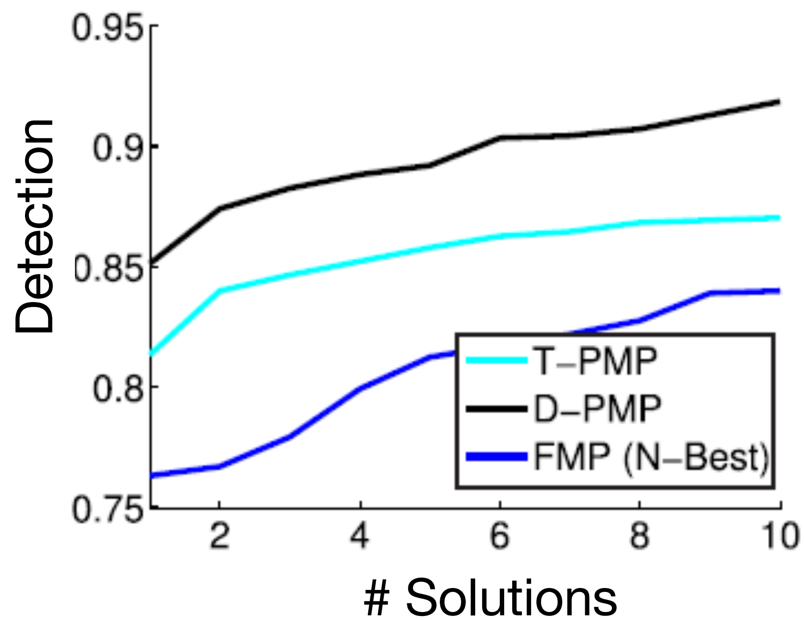

[Pacheco, Zuffi, Black & Sudderth, ICML 2014]

Real Images (Multiple People)

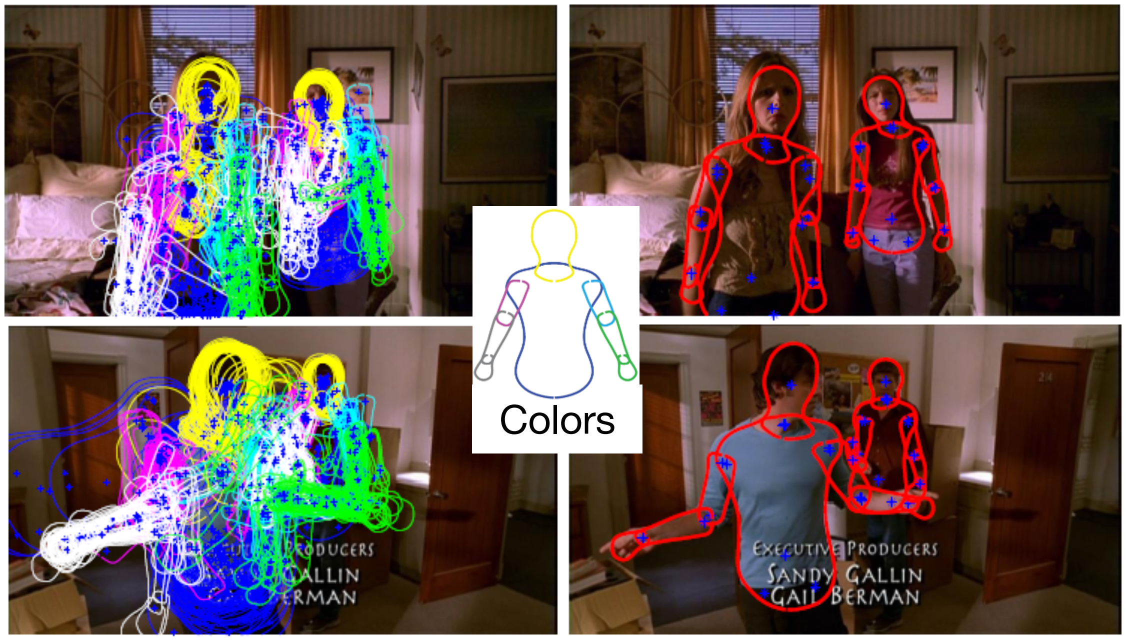

D-PMP Particles **Mode Estimates**

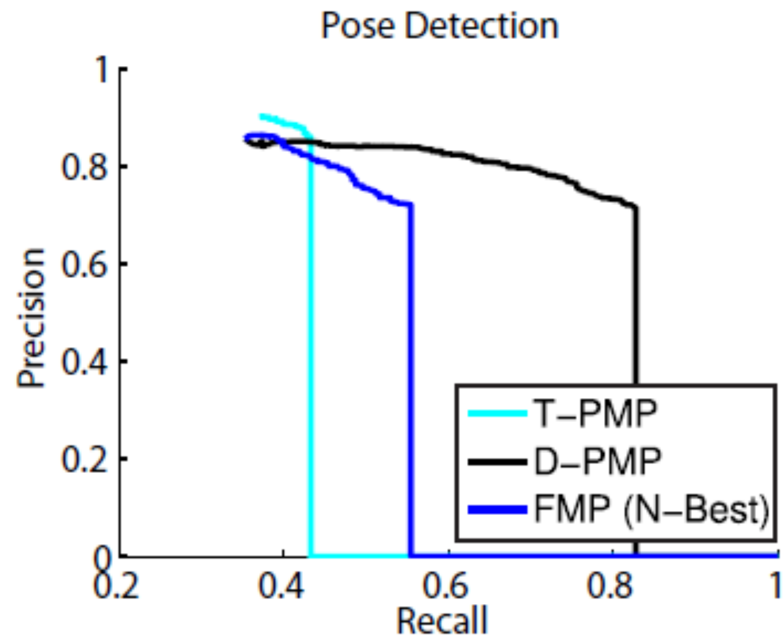

Precision-Recall for multi-person frames:

T-PMP : High precision, low recall, particles on one figure
D-PMP : Outperforms **FMP** and other particle methods

Note: G-PMP not reported due to poor performance.

[Pacheco, Zuffi, Black & Sudderth, ICML 2014]

D-PMP for 3D Mesh Alignment

Independent work by Zuffi & Black, appeared at CVPR 2015.

Articulated Pose Tracking

Prior work fails to show improvement by incorporating motion model.

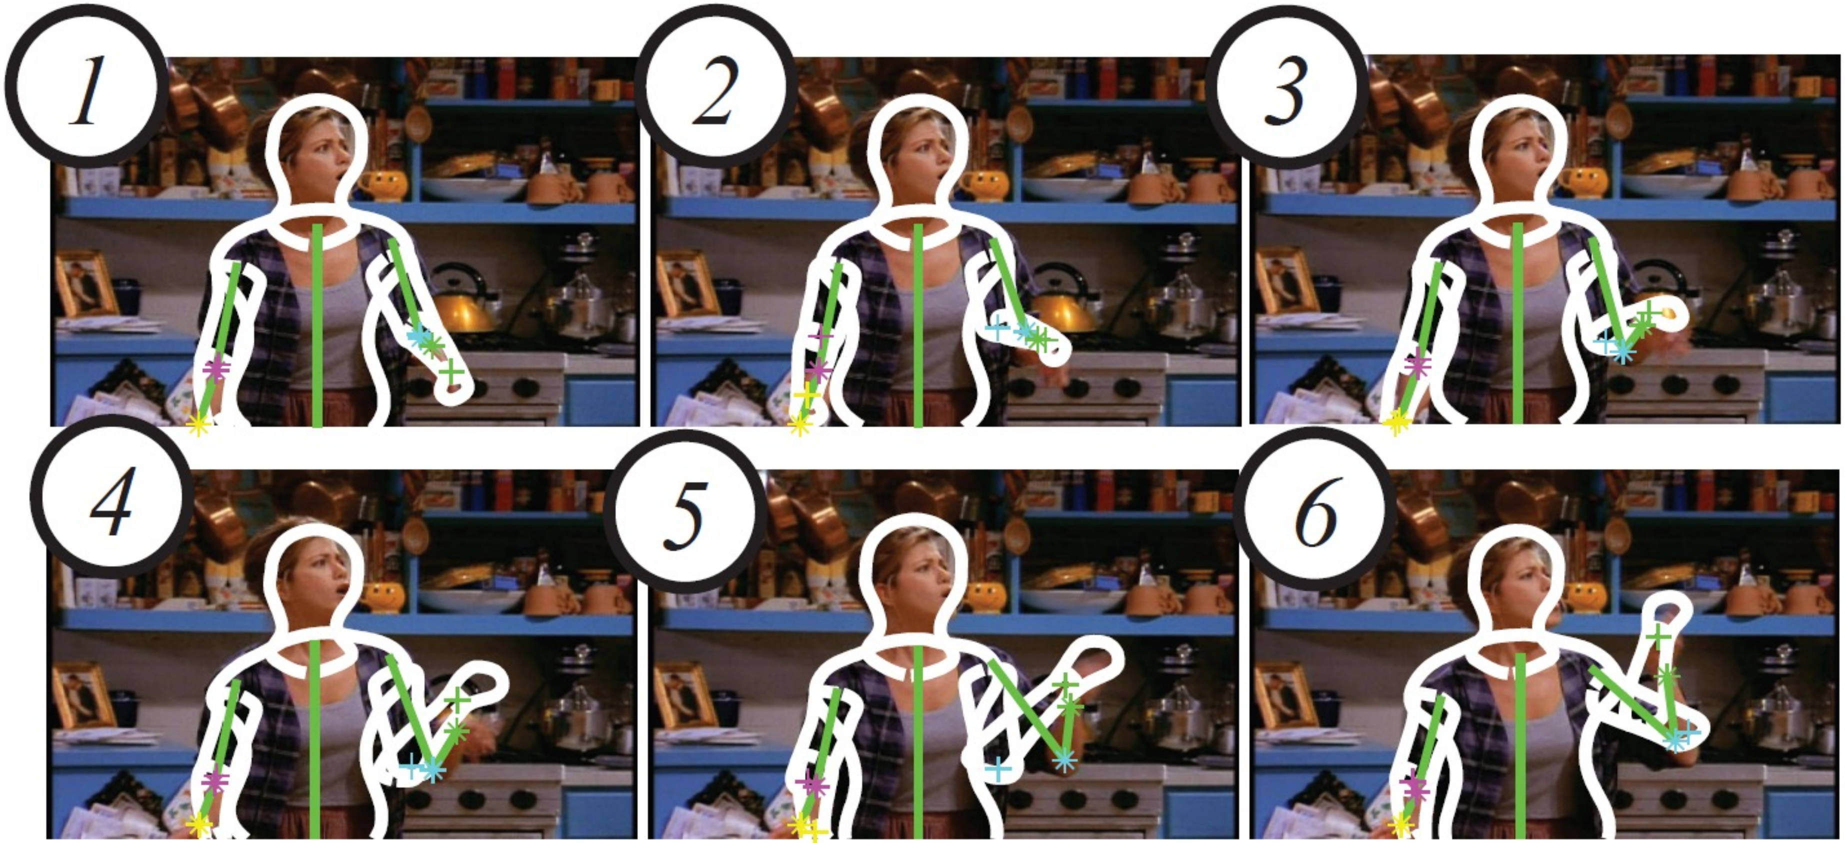

This is a failure of inference…

Articulated Pose Tracking

Data and Optical Flow

Extension of the *Flowing Puppets* model [Zuffi et al., 2013]

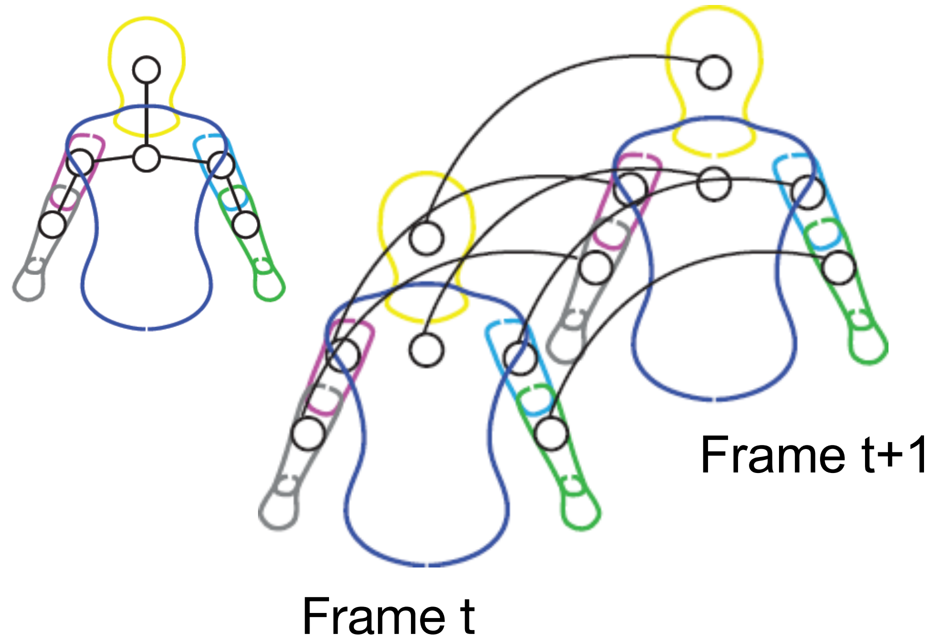

Part Likelihood

Gradients: Encode object and motion boundaries via HOG / HOF.

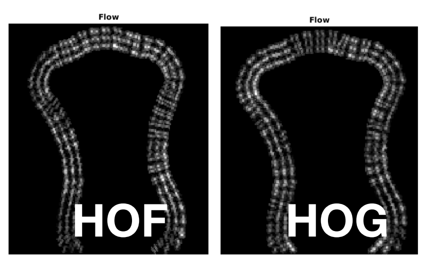

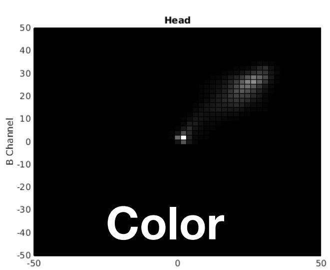

Appearance: 2D histogram of A/B color channels in L*a*b* space. Luminance ignored.

Prior

Part Motion: Scale mixture captures heavy tailed statistics of motion between frames.

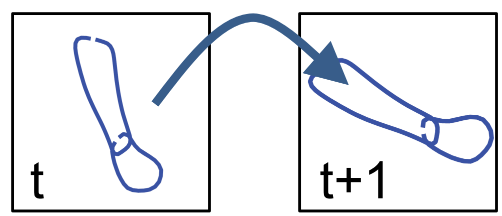

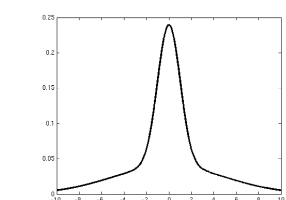

Structural prior identical to DS.

Loopy Max-Product BP

Many interesting models exhibit *cyclic* dependency structure...

Loopy Max-Product BP:
Iteratively update until converged.

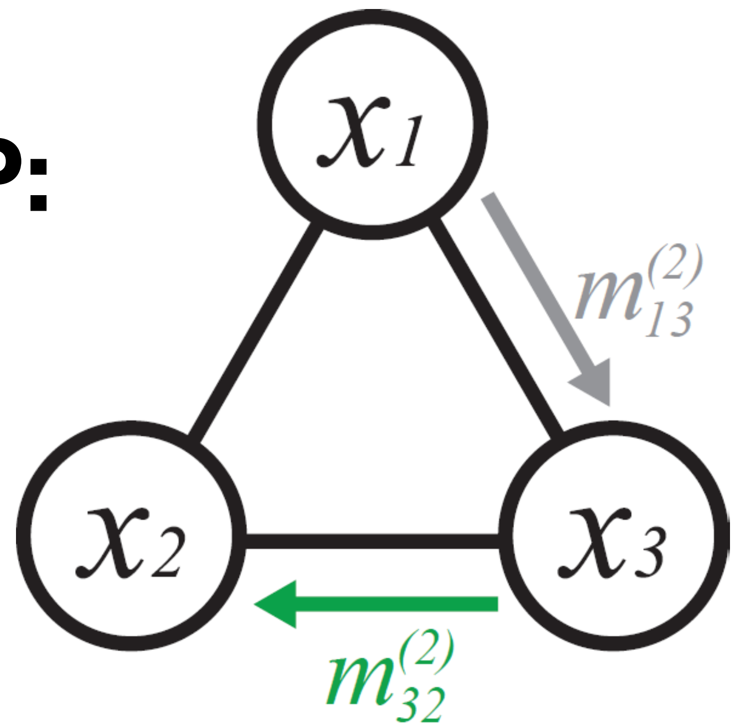

State-of-the-art decoding for error correcting codes but may perform poorly in general.

MAP Probability Bound

Spanning Tree Distribution

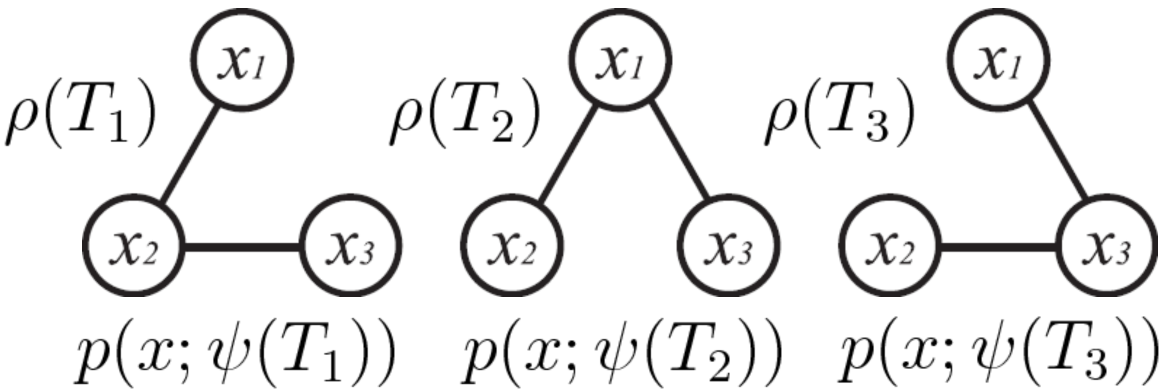

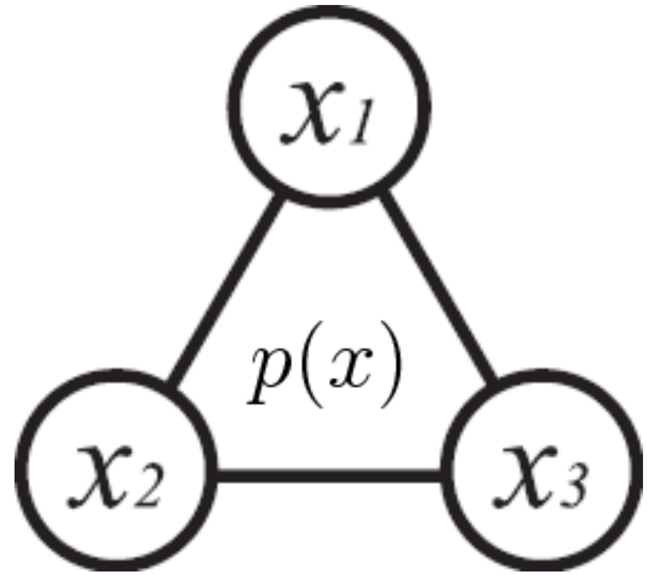

Bound MAP via Jensen's Inequality:

$$\max_x \log p(x) \le \sum_T \rho(T) \max_x \log p(x;\psi(T))$$

Dual Problem:

$$\underset{\psi}{\text{minimize}} \sum_T \rho(T) \max_x \log p(x;\psi(T))$$

[Wainwright et al., 2005]

Reweighted Max-Product (RMP)

Edge Appearance

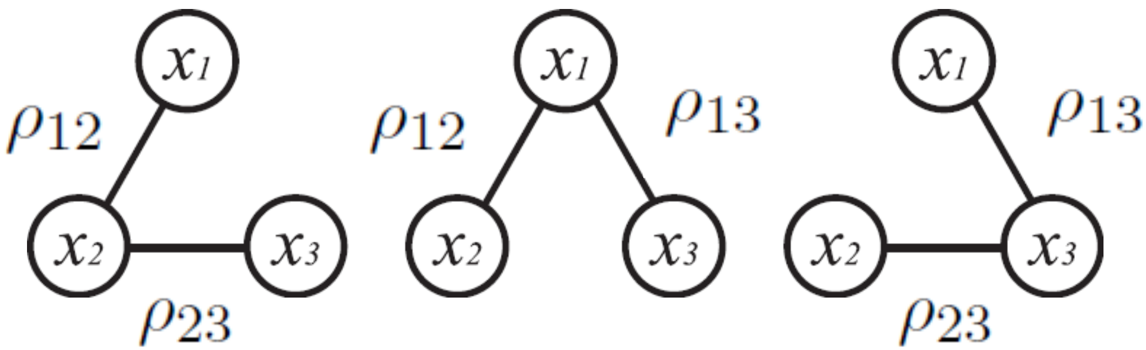

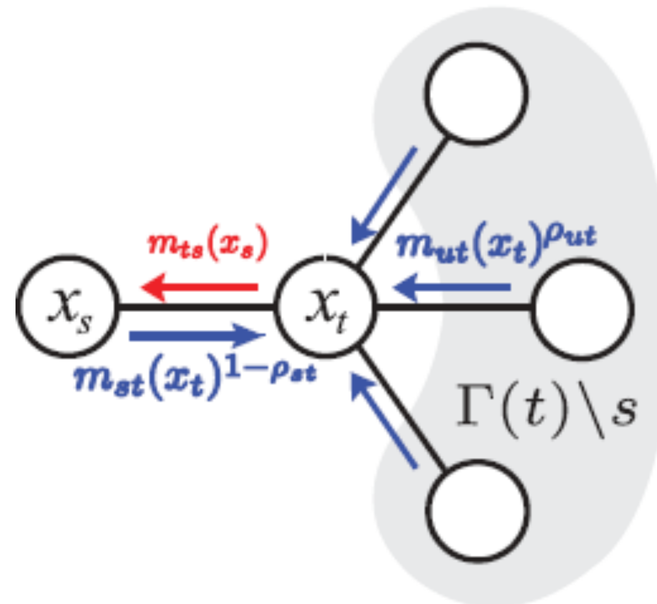

$$\color{red}{m_{ts}(x_s)} = \max_{x_t} \psi_t(x_t)\psi_{st}(x_s, x_t)^{\frac{1}{\rho_{st}}} \frac{\prod_{u\in\Gamma(t)\backslash s} \color{blue}{m_{ut}(x_t)^{\rho_{ut}}}}{\color{blue}{m_{st}(x_t)^{1-\rho_{st}}}}$$

Solve dual problem via *reweighted* message passing

[Wainwright et al., 2005]

RMP Bound Tightness

Pseudo-Max-Marginal distribution:

$$\nu_s(x_s) \propto \psi_s(x_s) \prod_{u \in \Gamma(s)} m_{us}(x_s)^{\rho_{us}} \approx q_s(x_s)$$

Consistent maximizer:

$$x_s^* = \operatorname*{argmax}_{x_s} \nu_s(x_s)$$

$$(x_s^*, x_t^*) = \operatorname*{argmax}_{x_s, x_t} \nu_{st}(x_s, x_t)$$

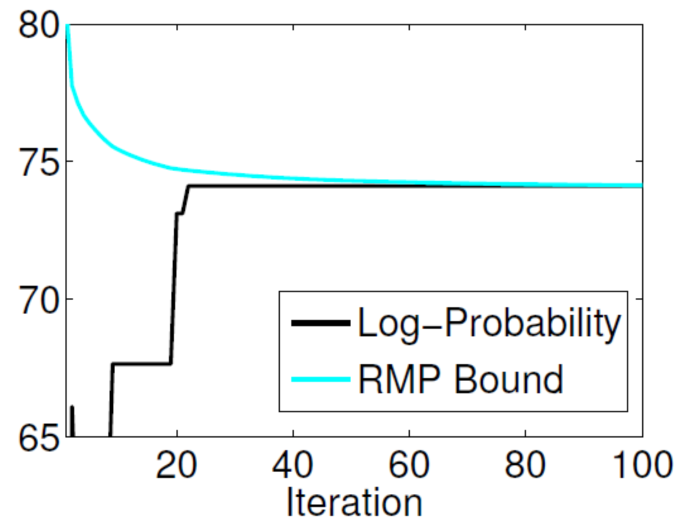

RMP bound tight and x^* global MAP:

$$\max_x \log p(x) = \sum_T \rho(T) \max_x \log p(x; \psi(T))$$

Reweighted BP & Stereo Vision

Left

Right

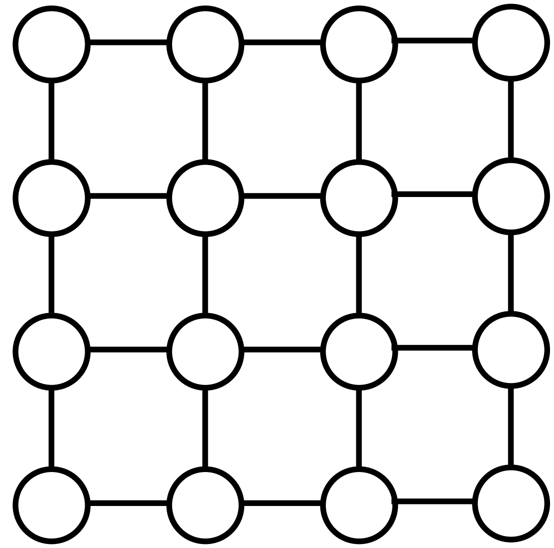

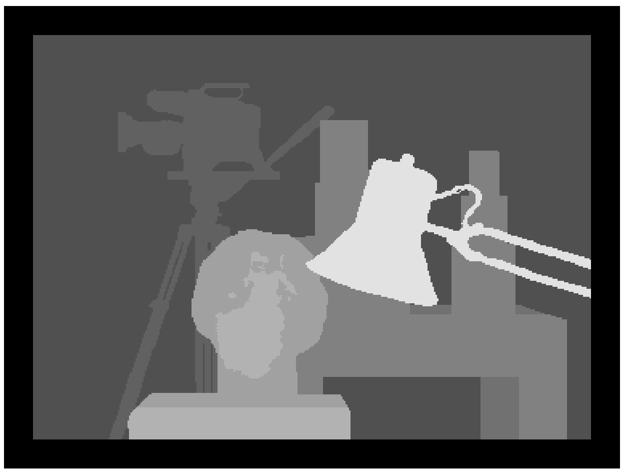

Disparity

$$-\log \psi_{st}(x_s, x_t)$$

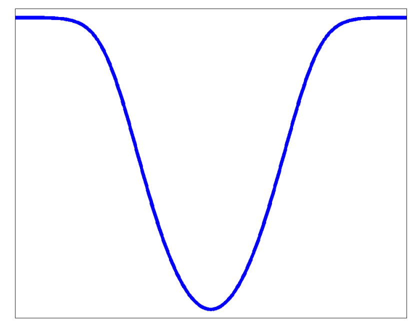

- State space is horizontal displacement (disparity) between corresponding pixels in aligned images (~50 options)
- Yanover, Meltzer, Weiss (JMLR 2006) show reweighted max-product finds global MAP in ~90% of test instances

Loopy Particle Max-Product

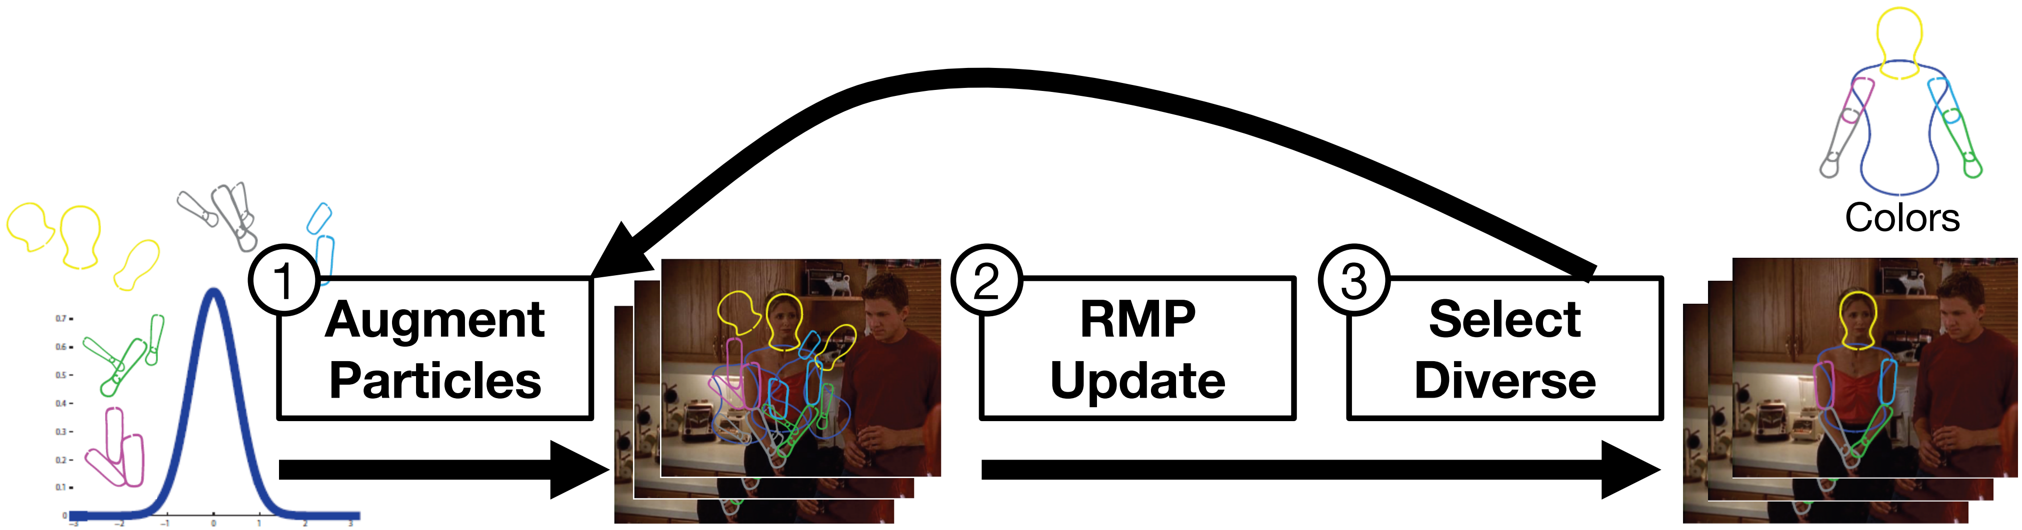

Select diverse subset and repeat…

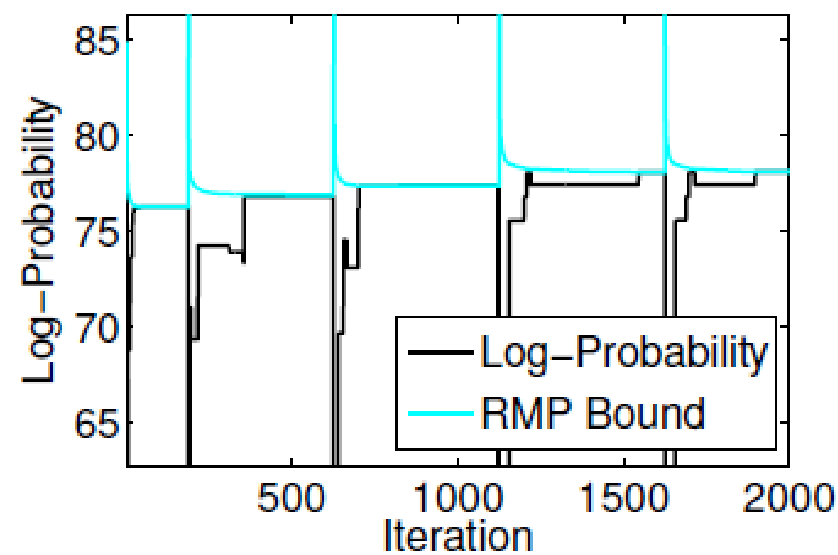

Diverse Particle Selection

Minimize reweighted message distortion:

$$\underset{z}{\text{minimize}} \sum_{s \in \Gamma(t)} \sum_{a=1}^{\alpha N} \left(m_{ts}(a)^{\rho_{st}} - \hat{m}_{ts}(a, z)^{\rho_{st}} \right)$$

$$\text{subject to } \|z\|_1 \leq N, z \in \{0, 1\}^{\alpha N}$$

- Accounts for spanning tree distribution
- Remains submodular
- Same greedy approximation

Pseudo-Max-Marginal Error

Selection IP objective upper bounds pseudo-max-marginal distortion.

$$\|\nu_s - \hat{\nu}_s\|_1 \leq \sum_{t \in \Gamma(s)} \sum_{a=1}^{\alpha N} \left(m_{ts}(a)^{\rho_{st}} - \hat{m}_{ts}(a)^{\rho_{st}} \right)$$

Recall pseudo-max-marginal definitions:

$$\nu_s(a) \propto \psi_s(a) \prod_{t \in \Gamma(s)} m_{ts}(a) \qquad \hat{\nu}_s(a) \propto \psi_s(a) \prod_{t \in \Gamma(s)} \hat{m}_{ts}(a)$$

VideoPose2 Experiments

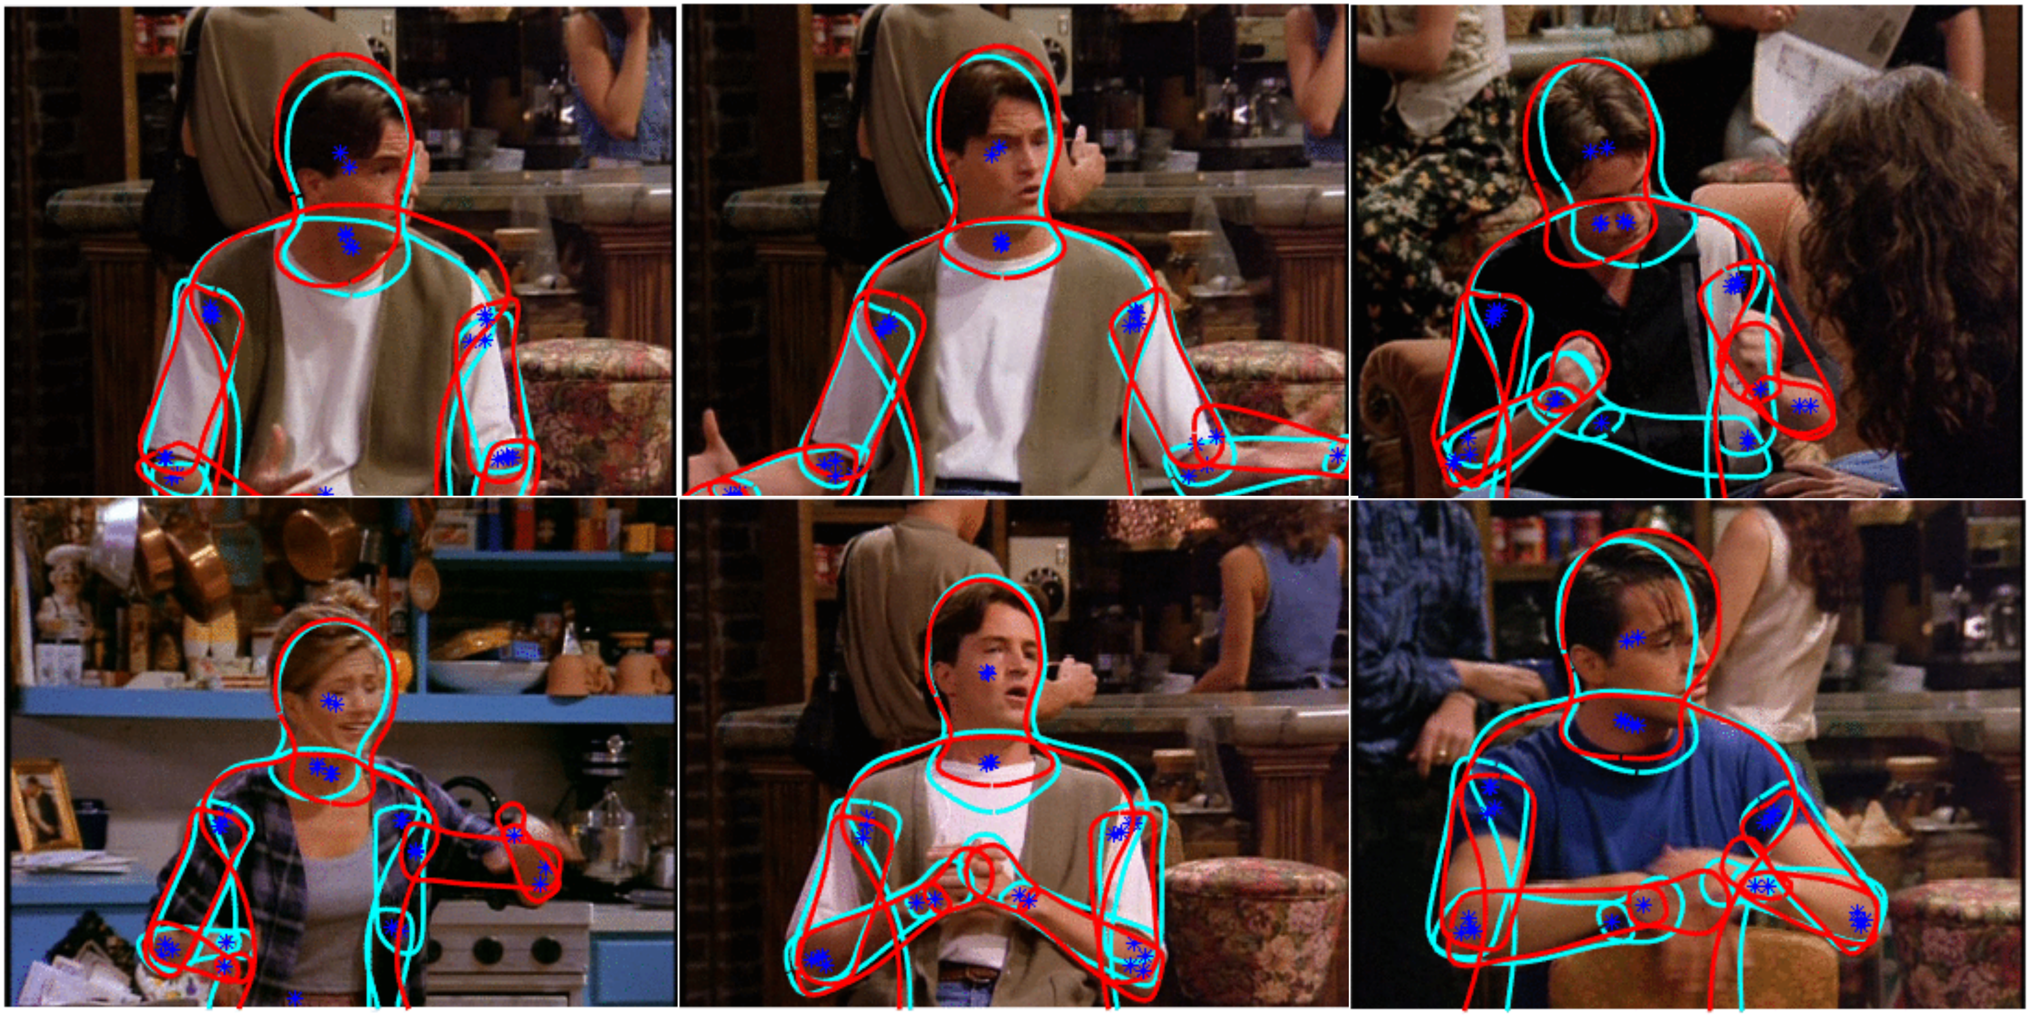

Comparison on VideoPose2 dataset of ~2,000 video frames from TV shows [Sapp et al., 2011]

T-PMP
D-PMP

Pose Tracking Particles

T-PMP

D-PMP **Both right arm hypotheses**

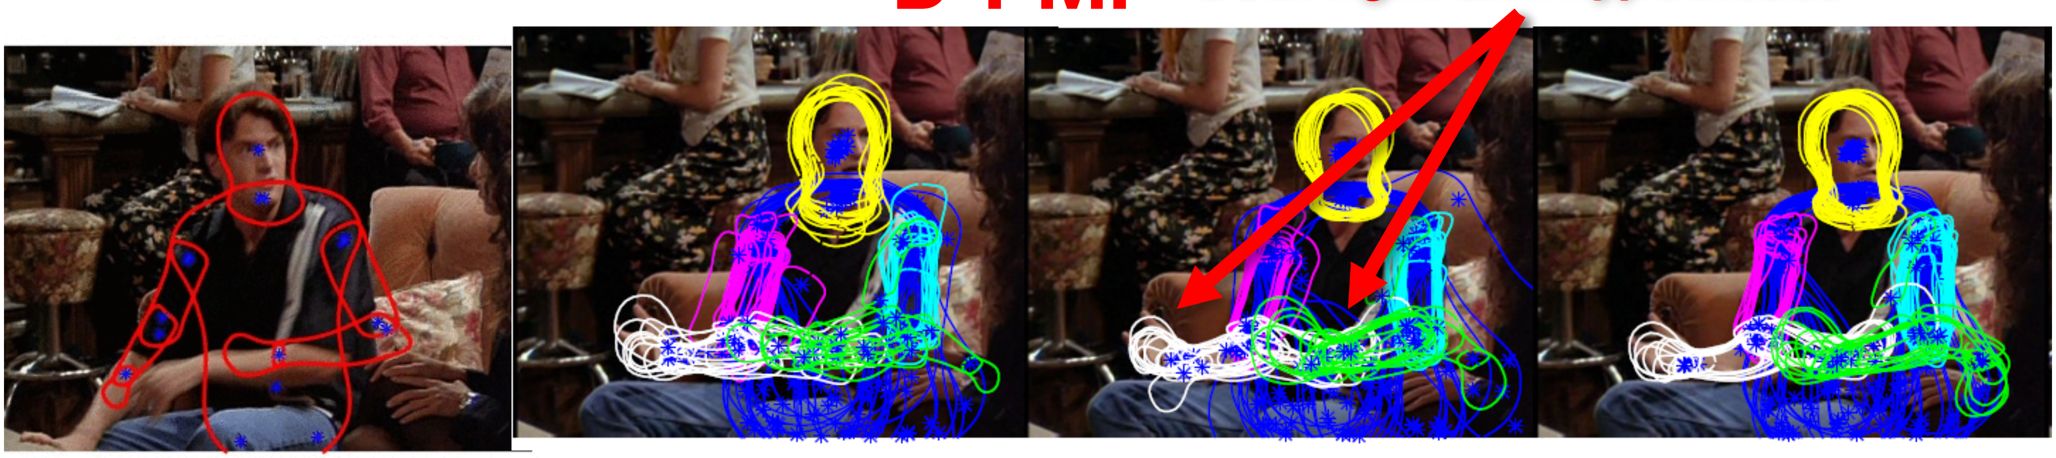

Greater diversity in particles allows D-PMP to reason more globally

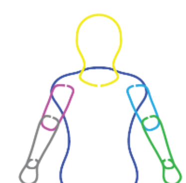

Colors

VideoPose2 Experiments [Sapp et al. 2011]

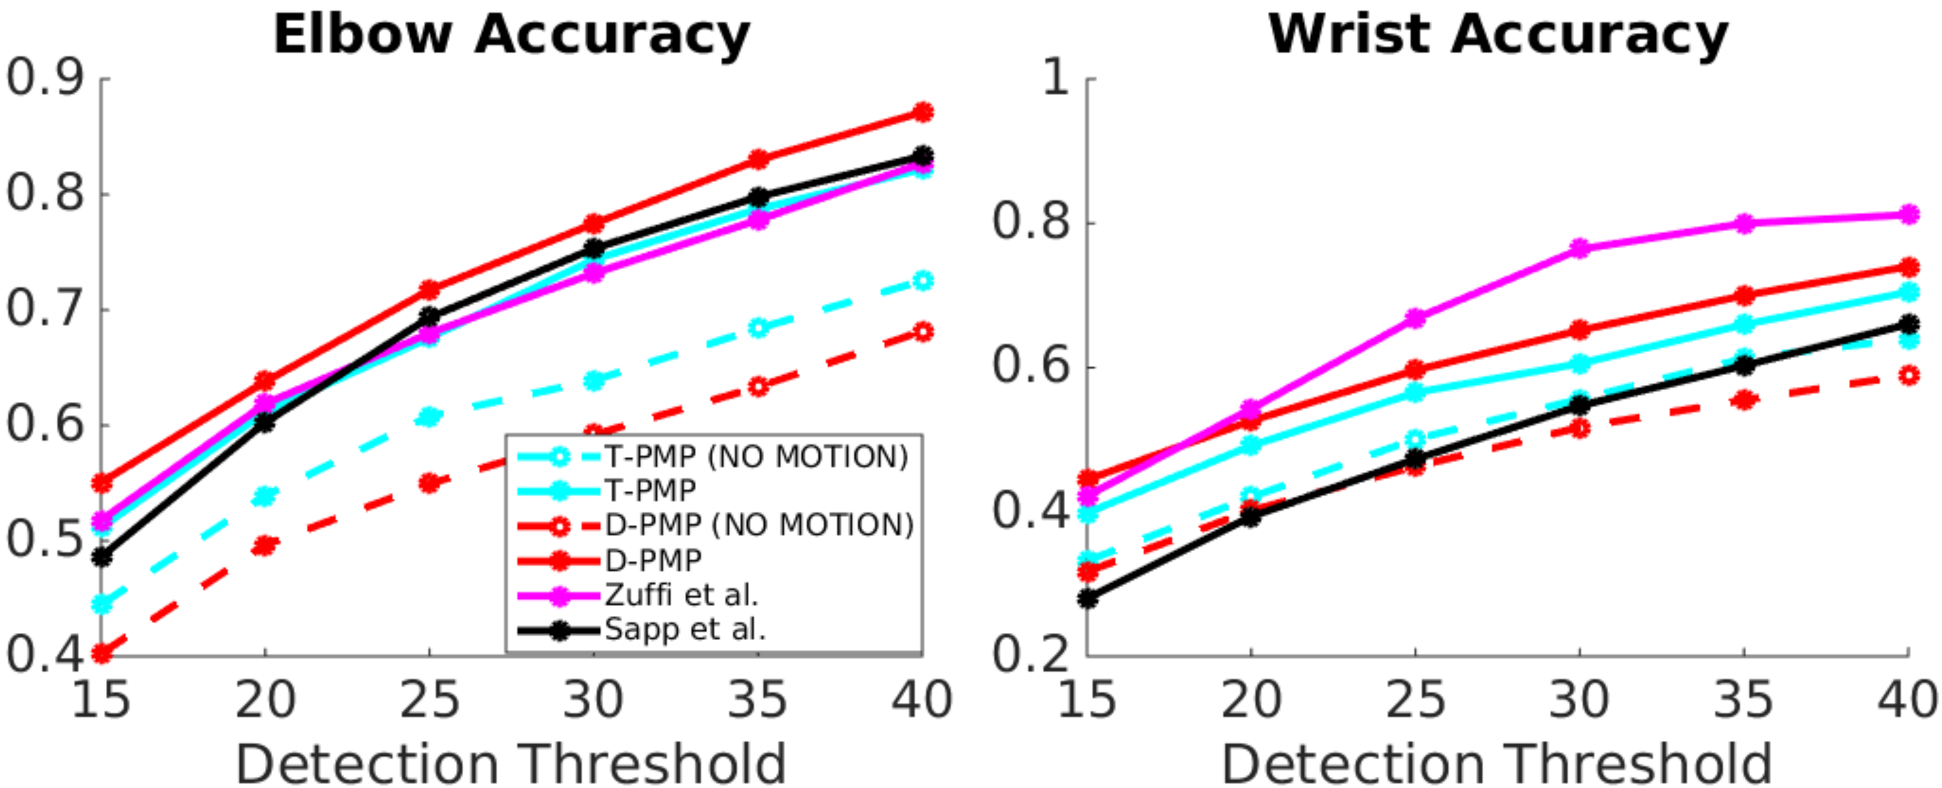

- Superior to static image estimates (--,--)
- Clear improvement over Sapp et al. baseline
- D-PMP superior to Flowing Puppets in close detection ranges. Looking at failure cases.

Protein Structure Prediction

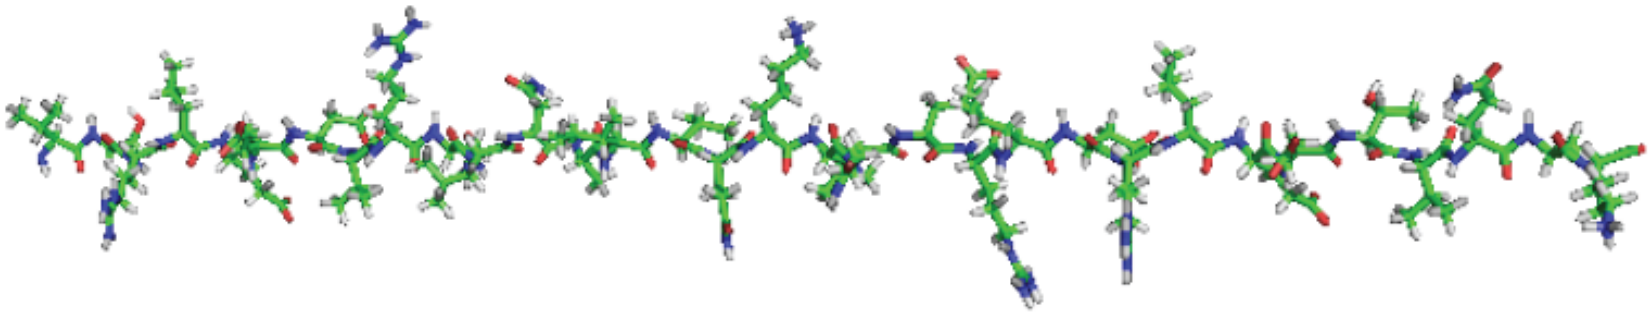

V-S-R-L-E-E-D-V-R-N-L-N-A-I-V-Q-K-L-Q-E-R-L-D-R-L-E-E-T-V-Q-A-K

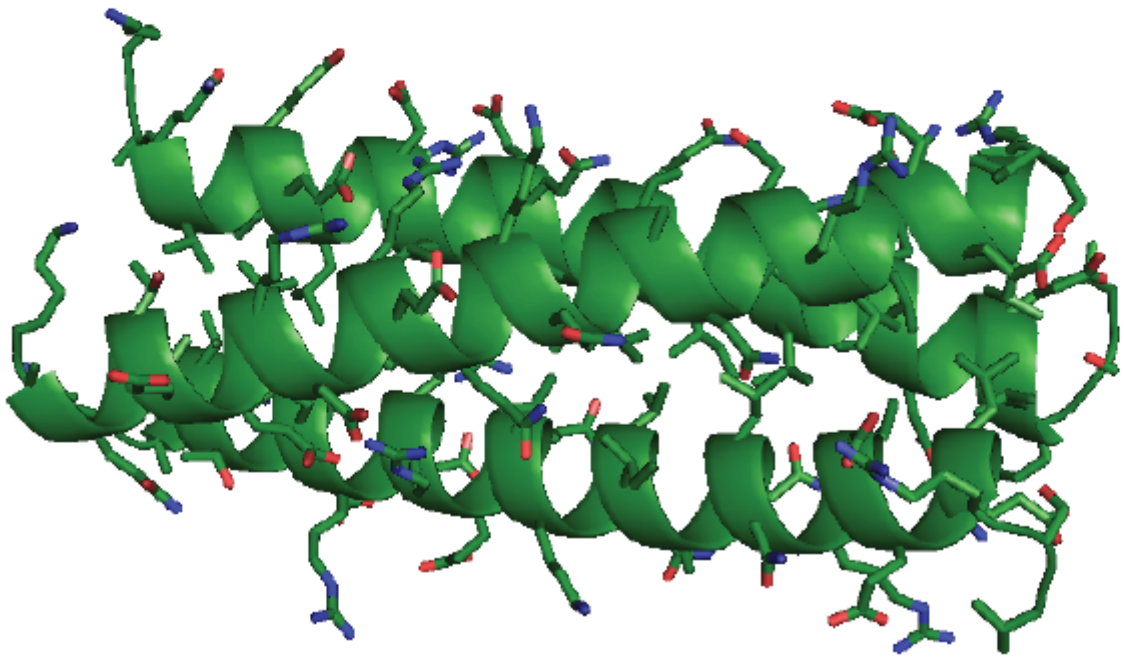

All information for predicting 3D structure encoded in amino acid sequence and physics

Protein Side Chains

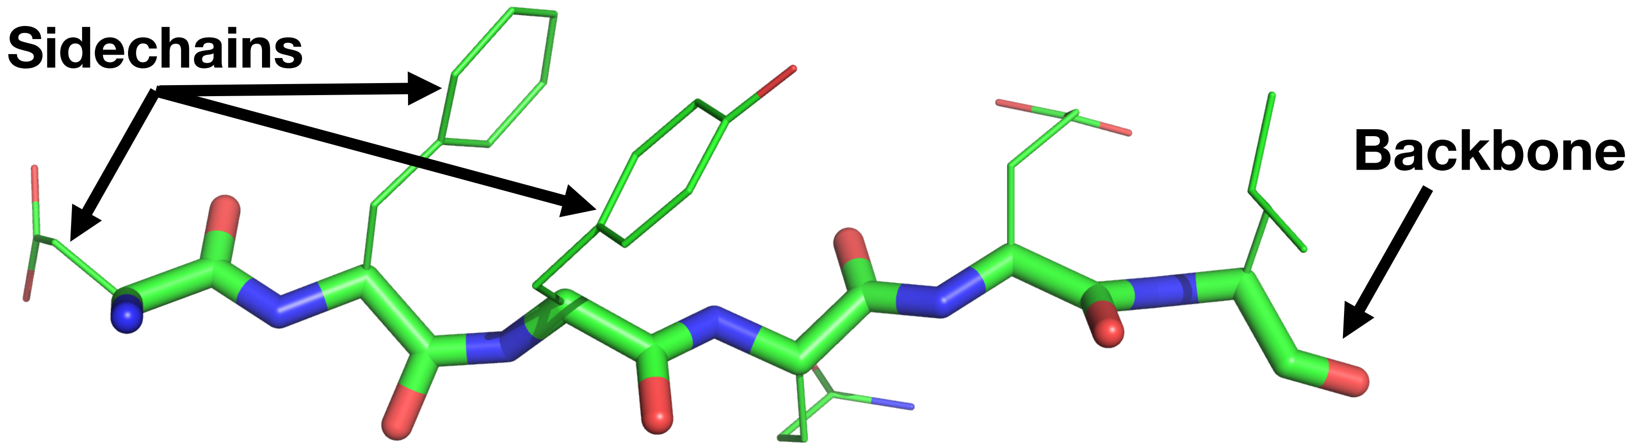

Side chain prediction: Estimate side chains given <u>fixed</u> backbone.

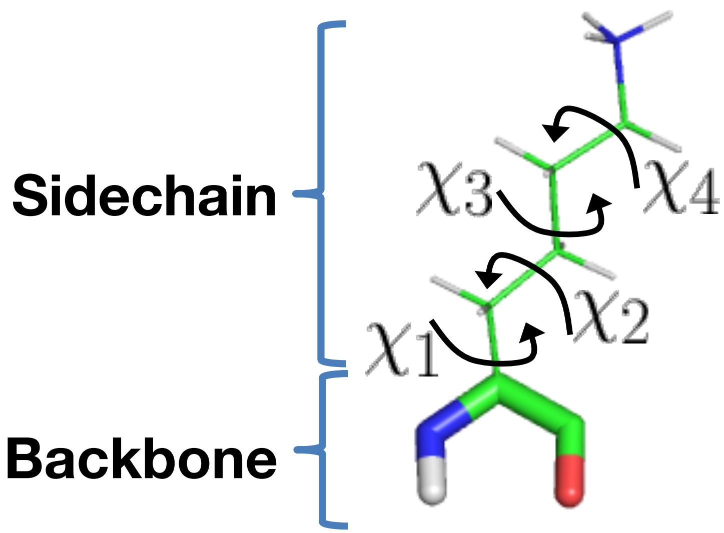

20 Amino Acid Types

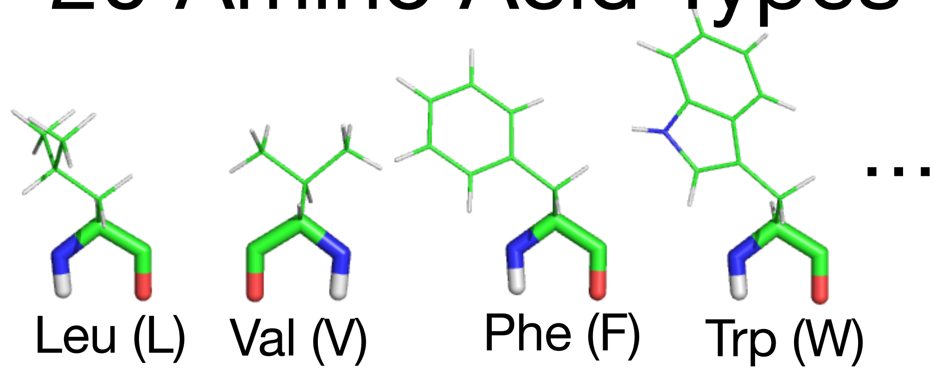

Dihedrals and Rotamers

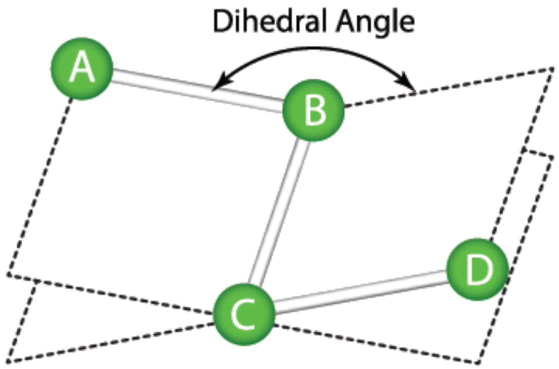

Dihedral Angles:

- ➢ Compact angular encoding
- ➢ 1D-4D continuous state

Rotamer discretization based on marginal statistics fails to capture fine details…

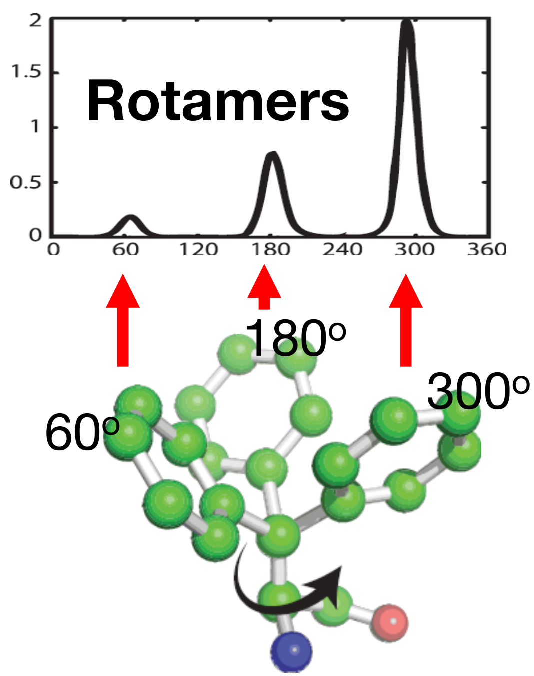

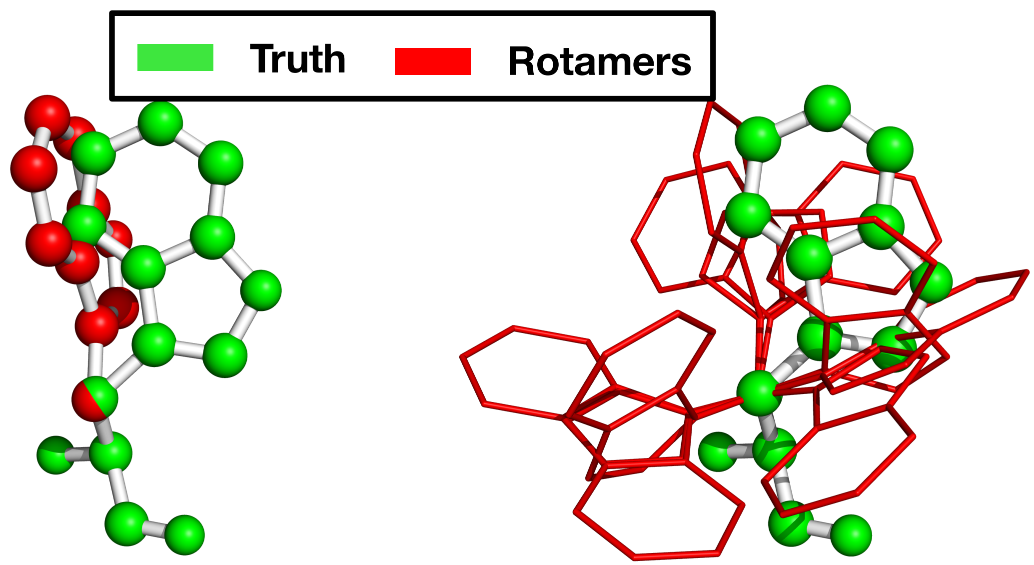

[Shapovalov & Dunbrack 2007]

Side Chain Prediction

$$p(x) \propto \prod_{s \in \mathcal{V}} \psi_s(x_s) \prod_{(s,t) \in \mathcal{E}} \psi_{st}(x_s, x_t)$$

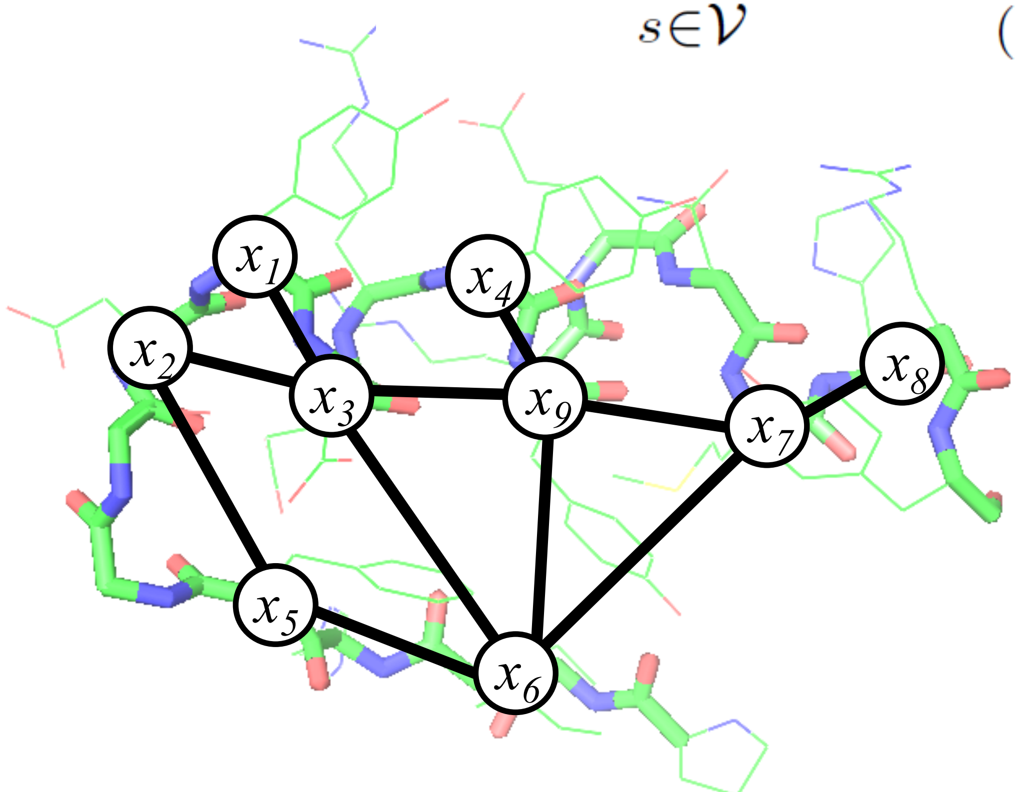

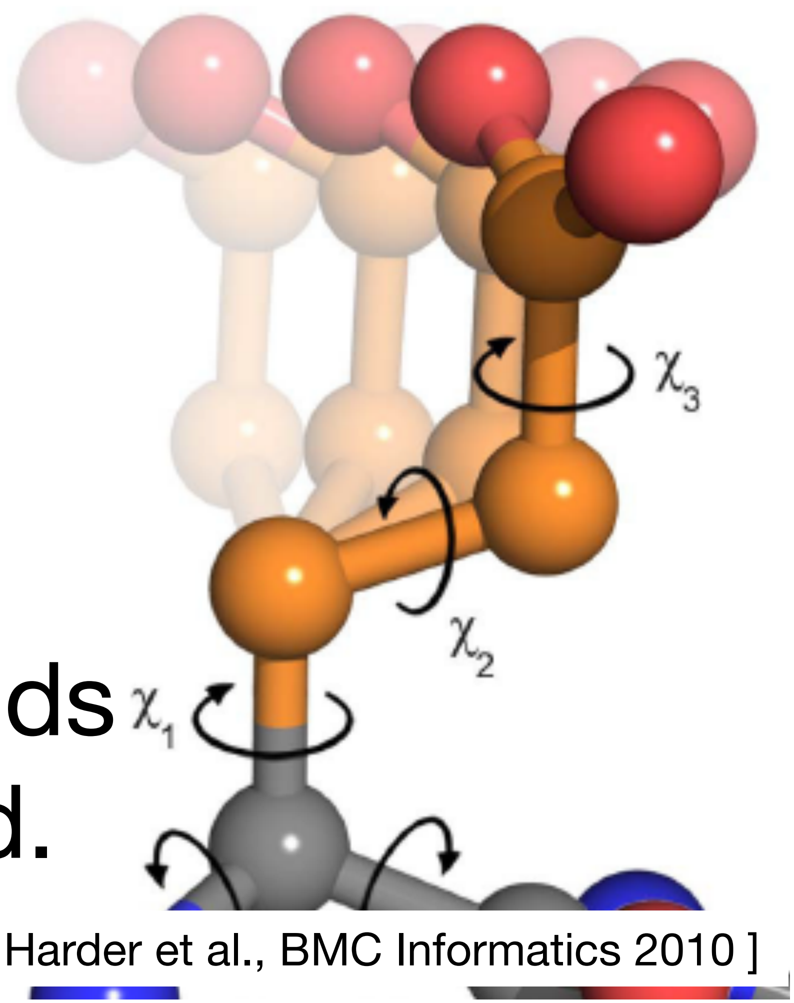

Edges between amino acids within distance threshold.

[Image: Harder et al., BMC Informatics 2010]

Side Chain Prediction

$$p(x) \propto \prod_{s \in \mathcal{V}} \psi_s(x_s) \prod_{(s,t) \in \mathcal{E}} \psi_{st}(x_s, x_t)$$

Rotamer Likelihood

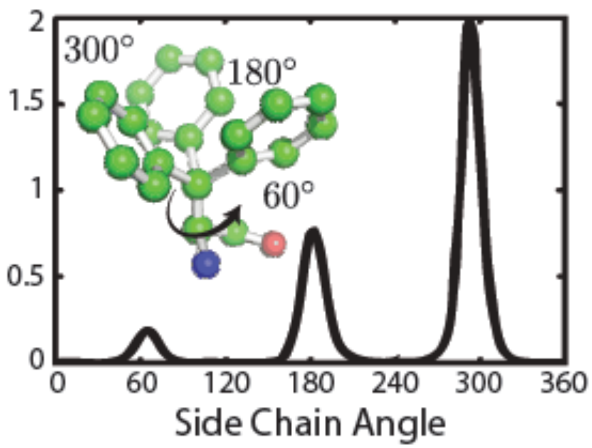

Atomic Interaction

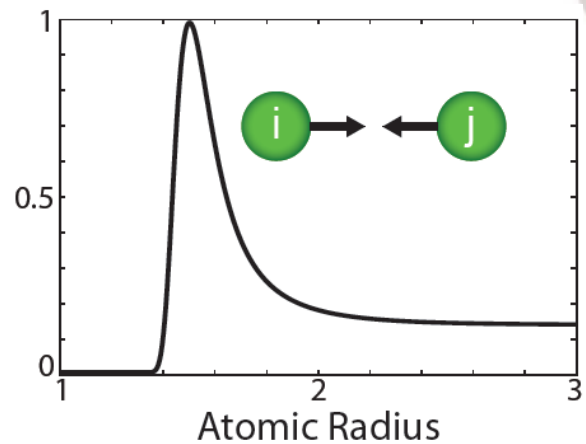

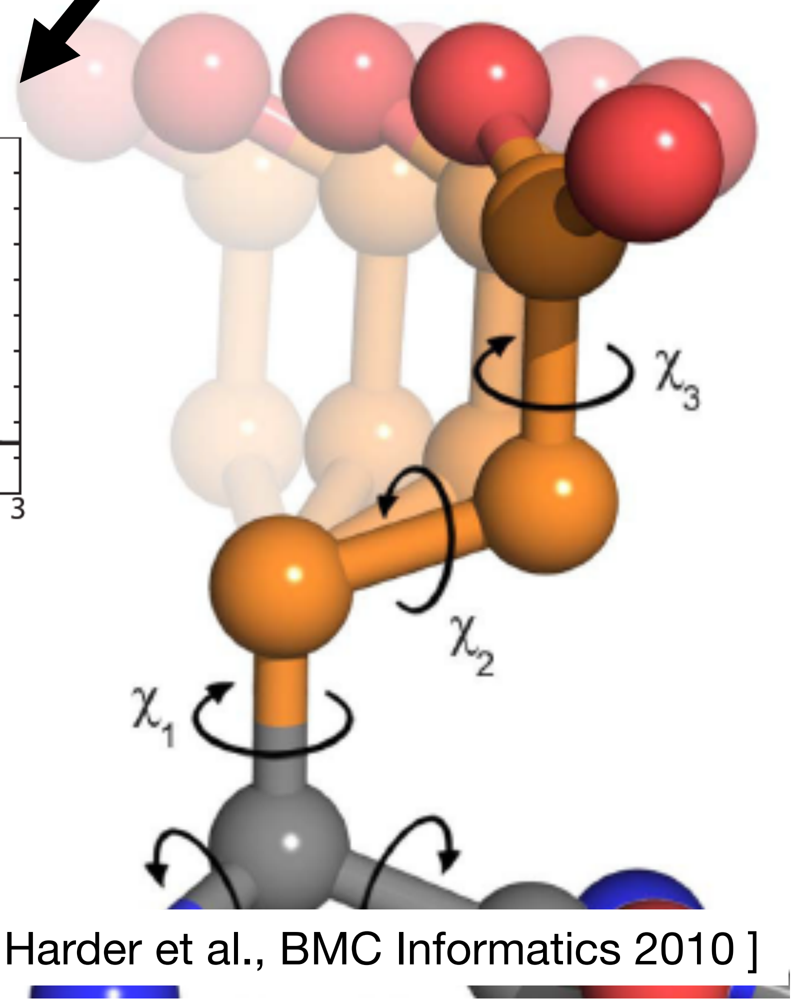

Statistical and physical potential functions.

[Image: Harder et al., BMC Informatics 2010]

D-PMP for Side Chains

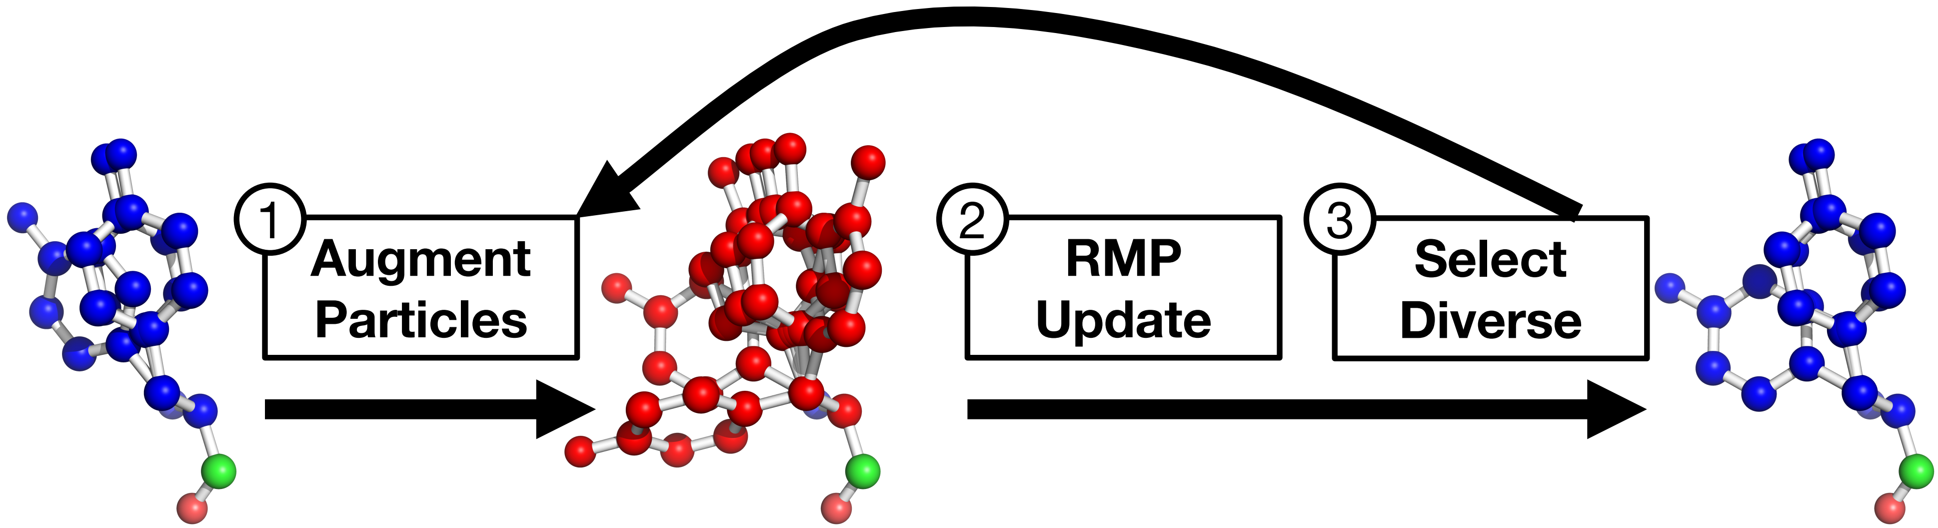

Continuous optimization of side chains:

- ➢Captures non-rotameric side chains
- ➢Conformational diversity
- ➢Likelihood-based proposals

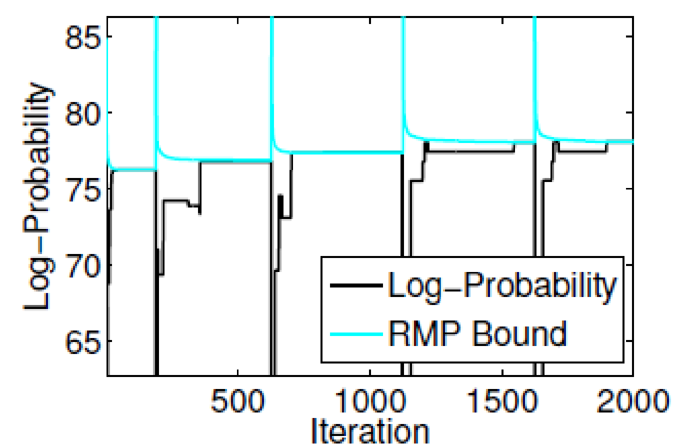

Rosetta

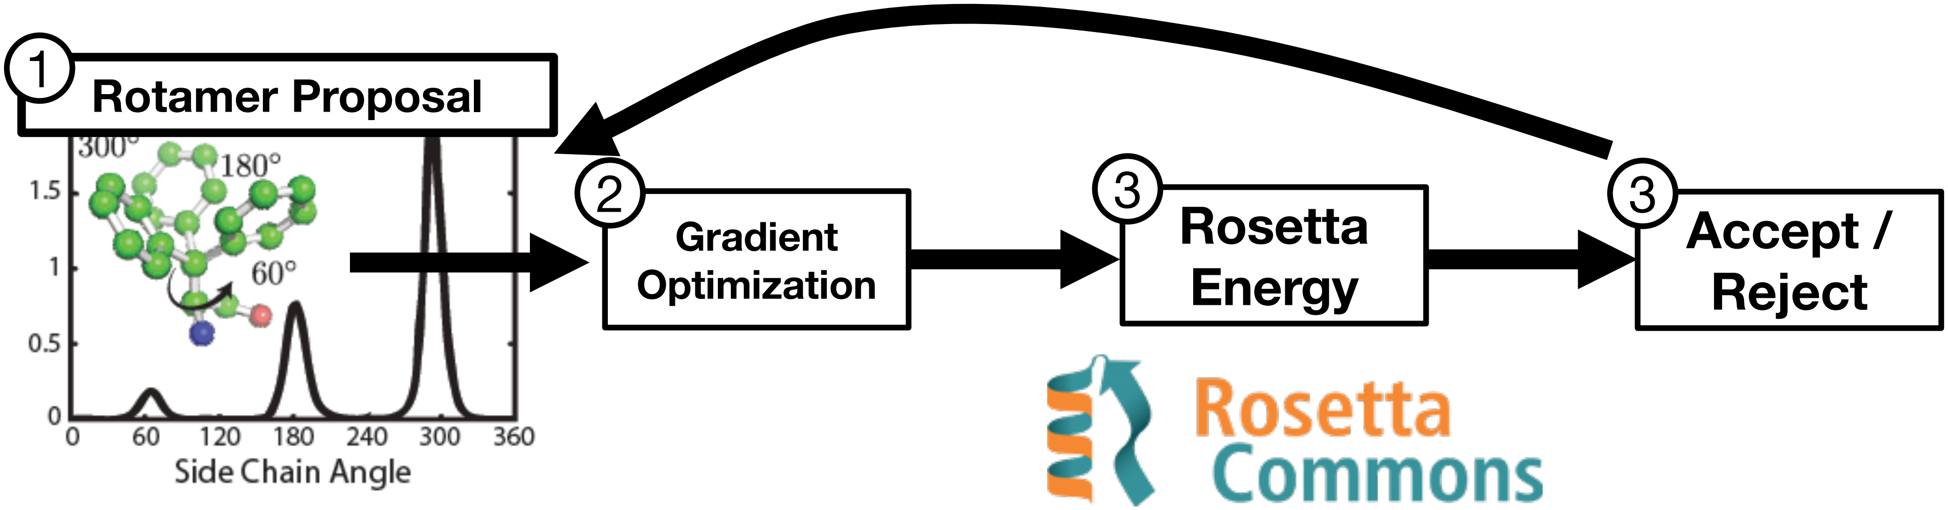

- ➢ Energy model used in FoldIt game
- ➢ Simulated annealing (SA) Monte Carlo
- ➢ Independent chains for multiple optima

Replace SA with D-PMP. Use Rosetta as black-box energy method.

Protein Side Chain Prediction

Log-probability of MAP estimate for…

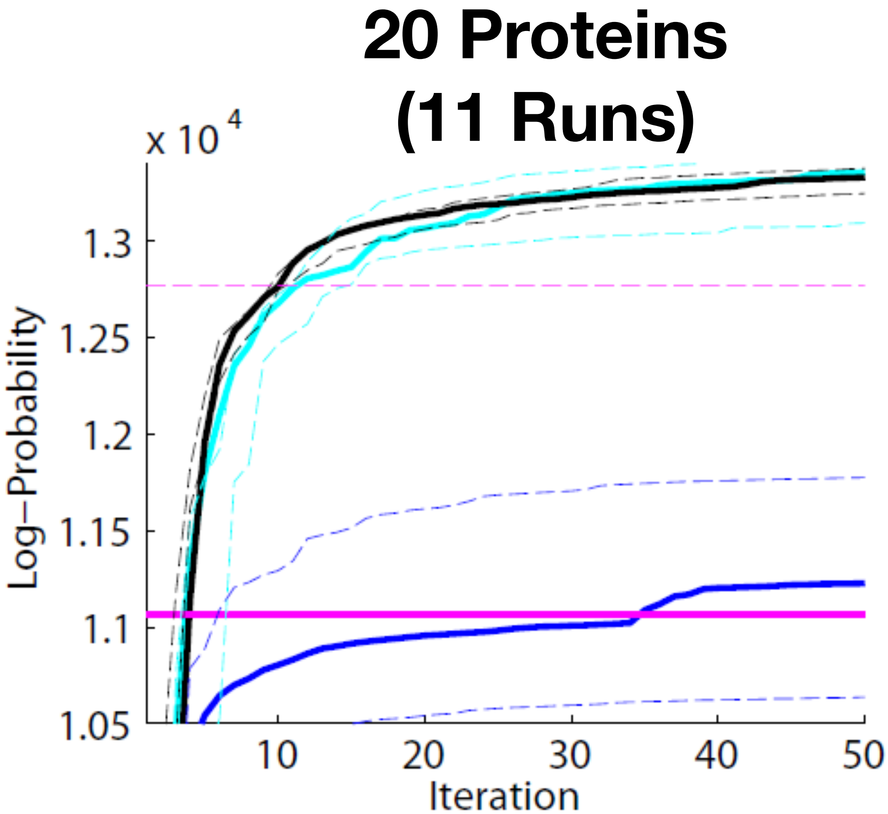

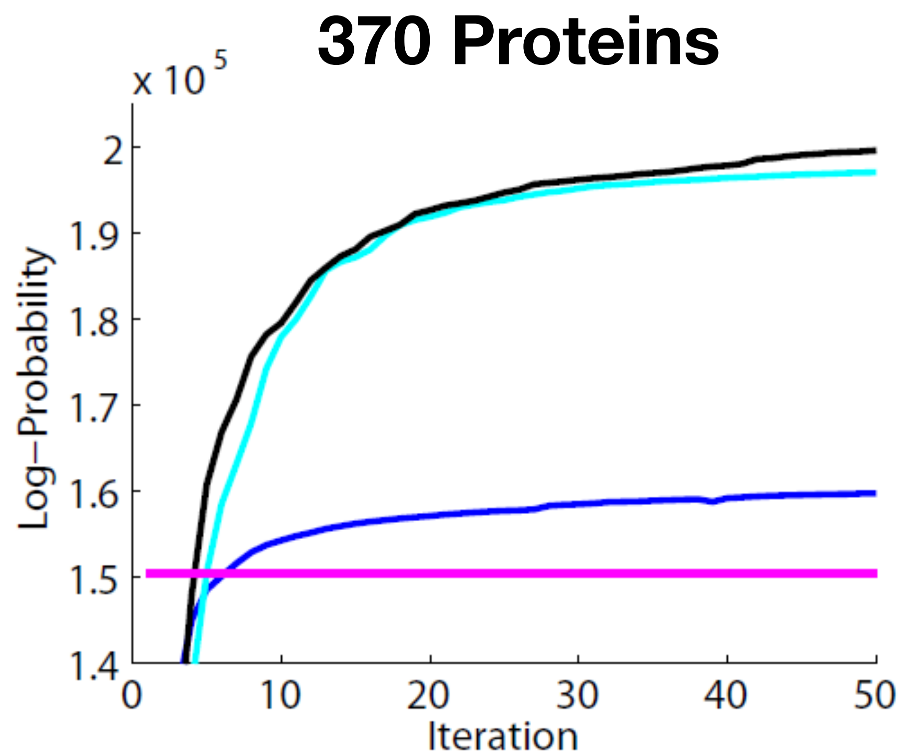

G-PMP, T-PMP, D-PMP, Rosetta simulated annealing [Rohl et al., 2004]

[Pacheco et al., ICML 2015]

Protein Side Chain Prediction

Root mean square deviation (RMSD) from x-ray structure.

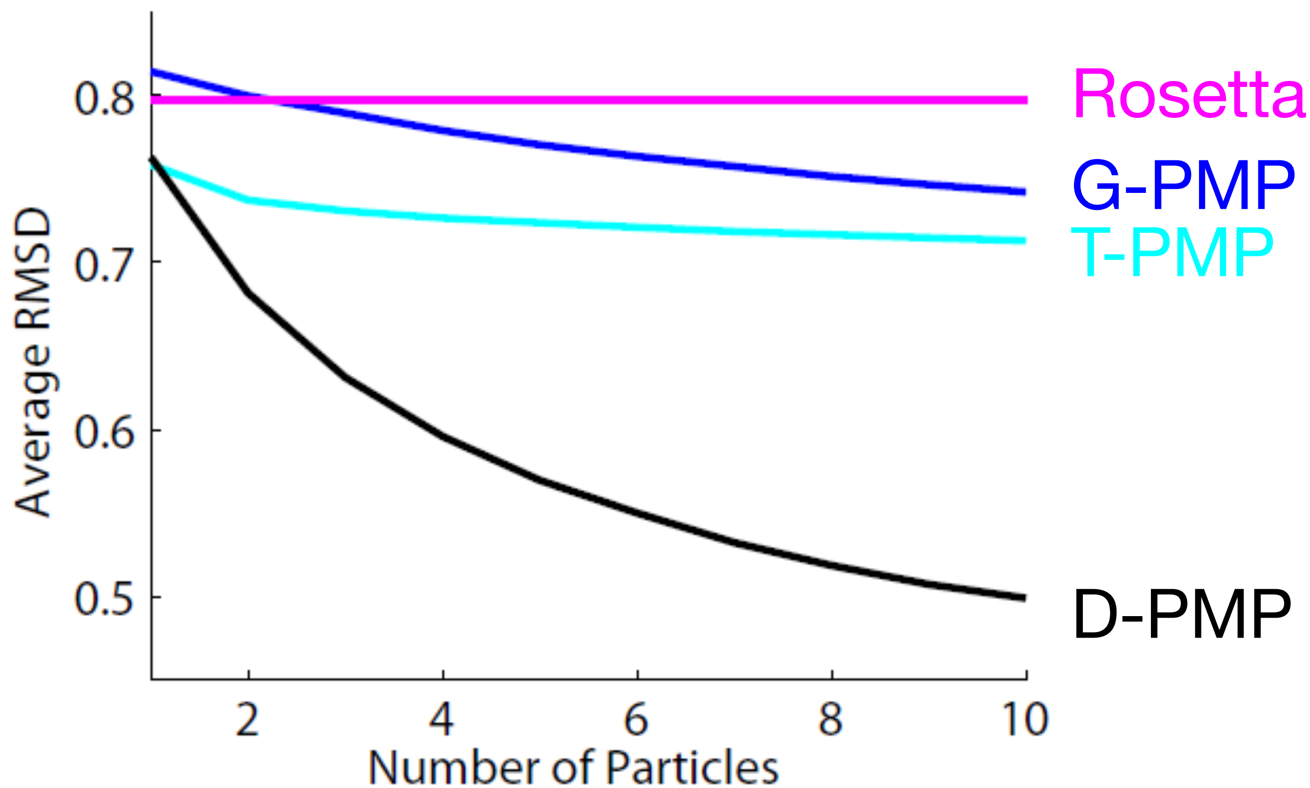

Oracle selects best configuration in current particle set.

Non-Rotameric Side Chains

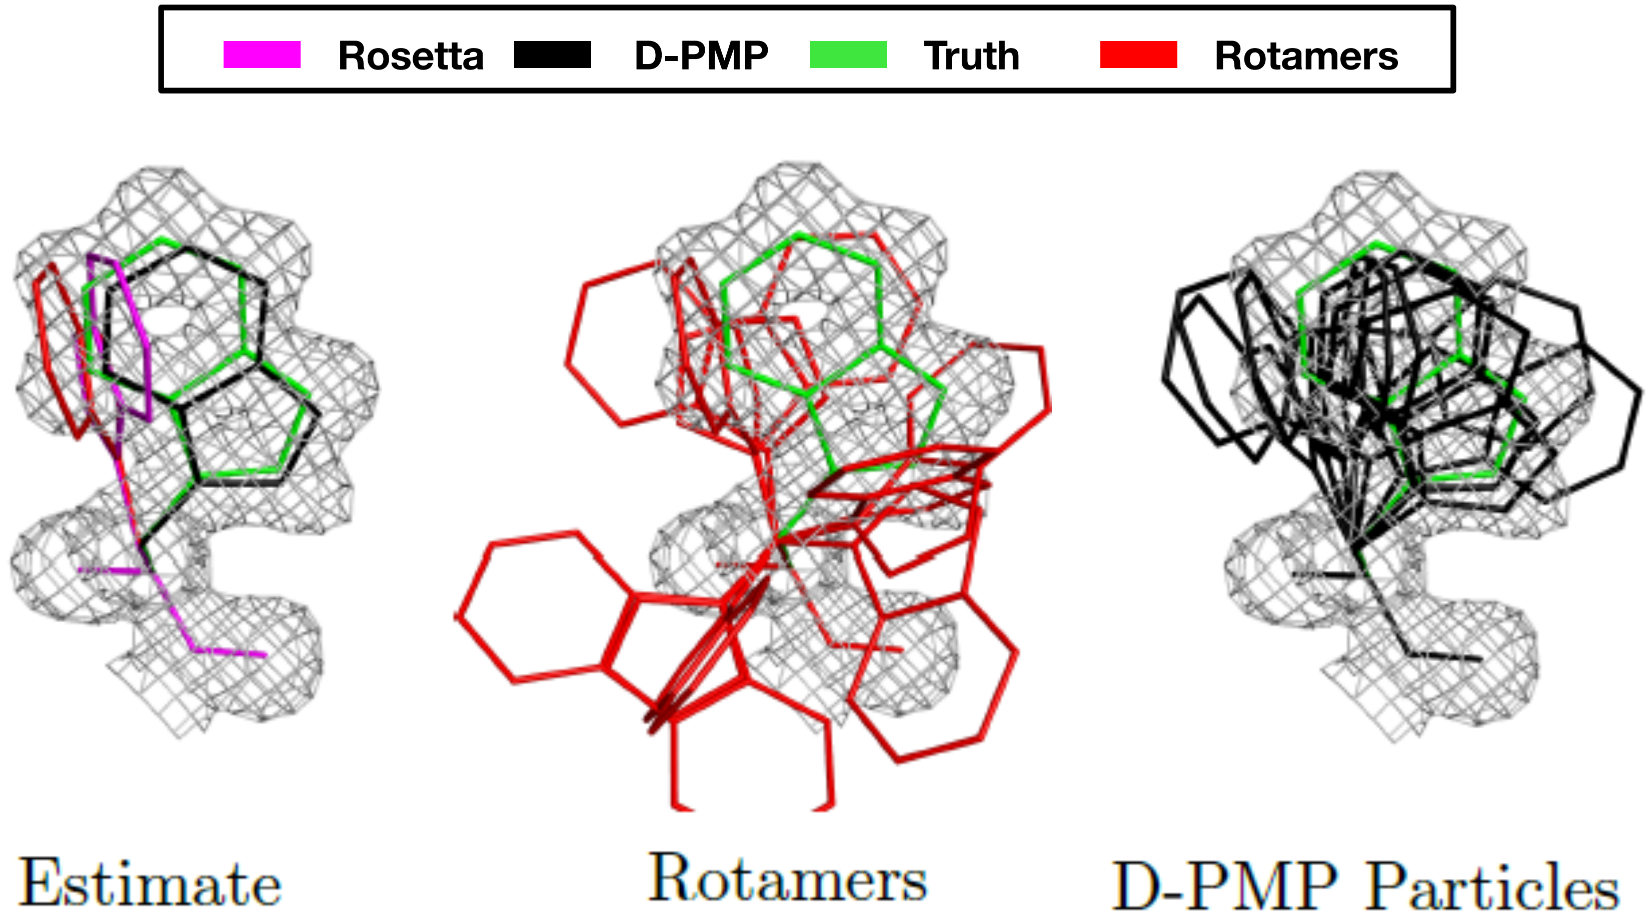

Not all side chains obey standard rotamer discretization.

Penicillin Acylase Complex, Trp154 [Shapovalov & Dunbrack 2007]

Protein Side Chain Prediction

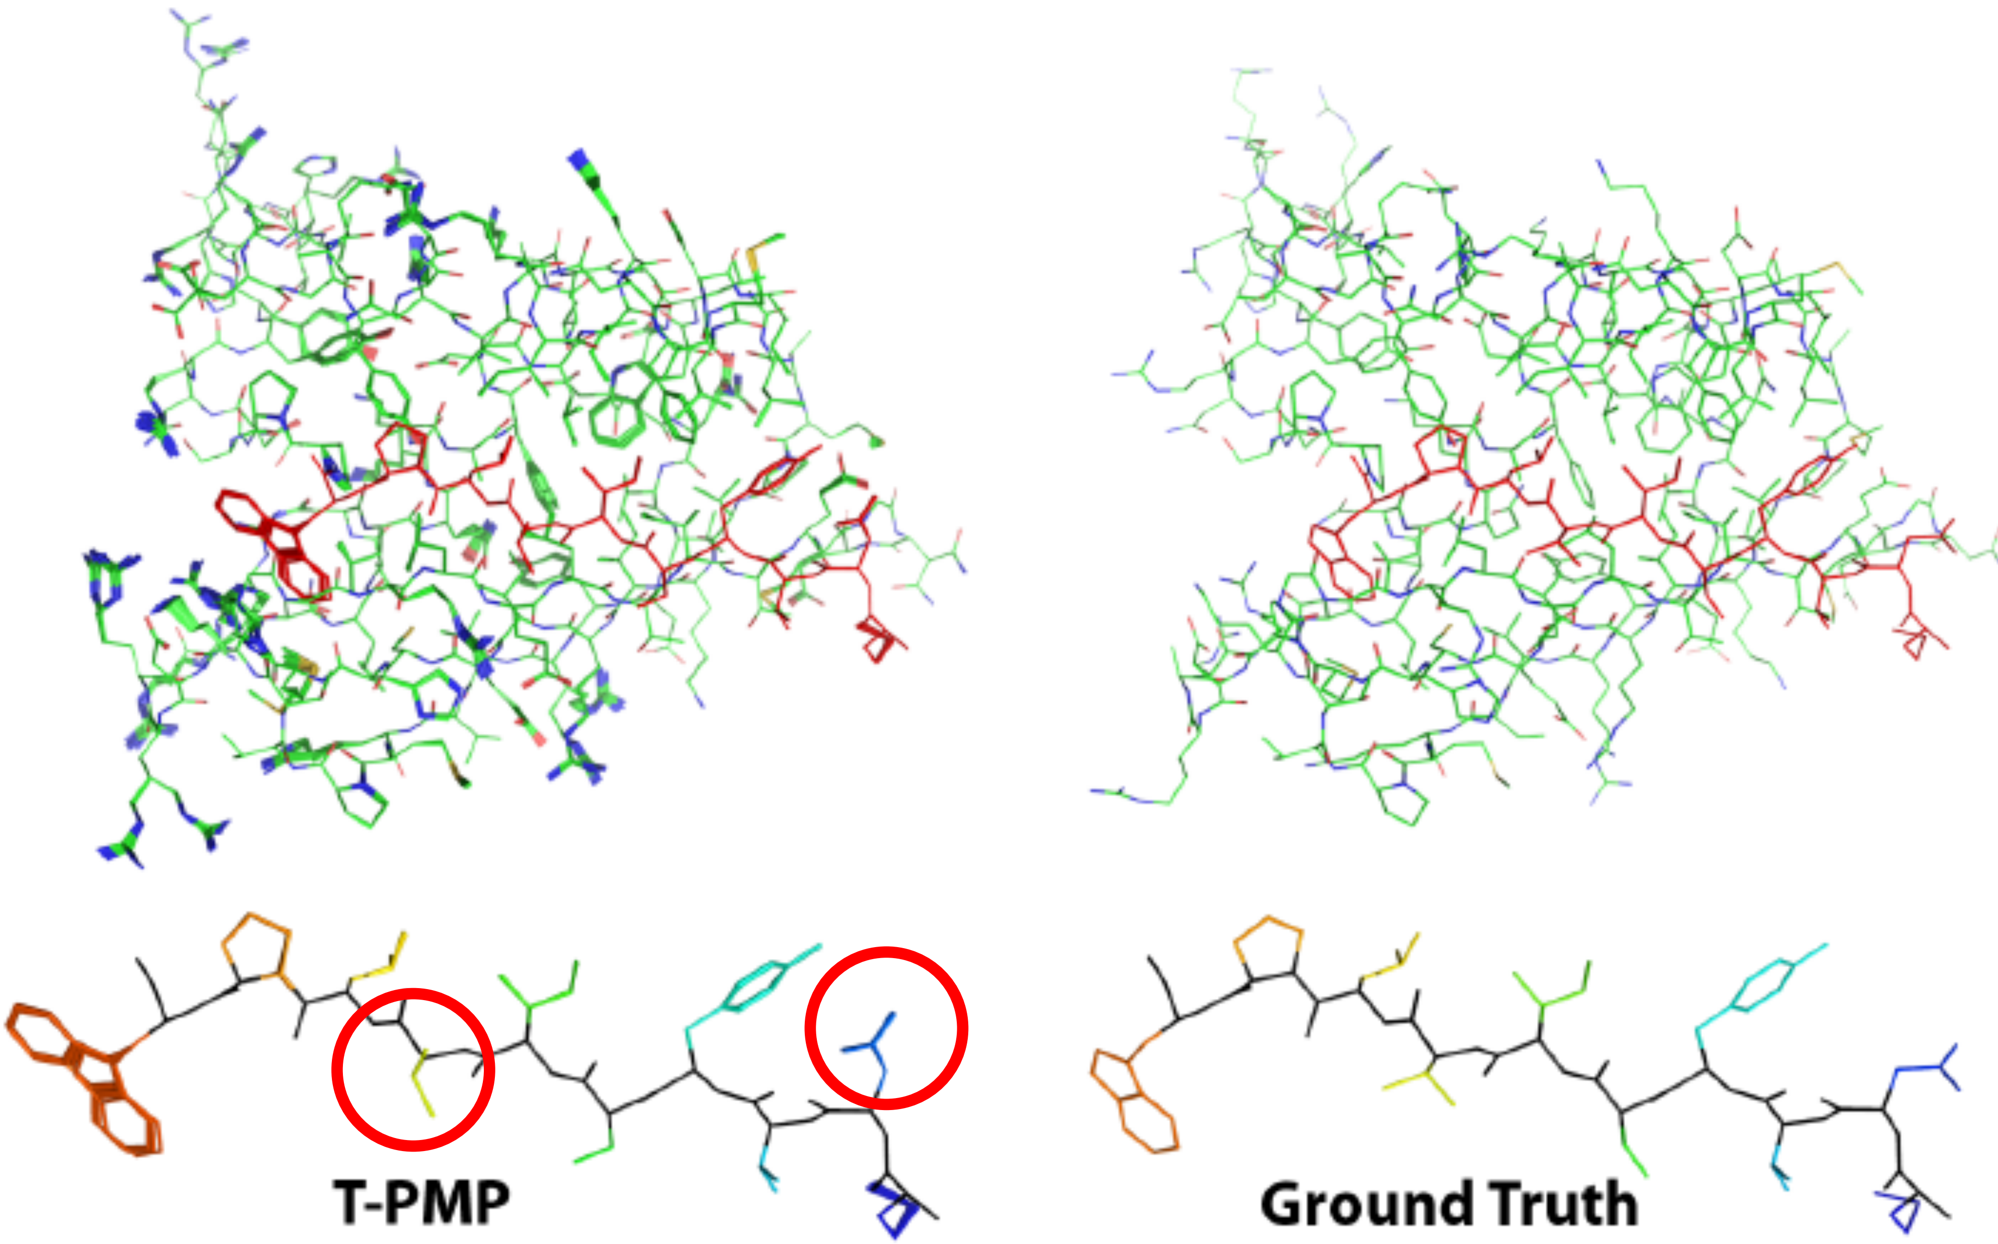

Protein Side Chain Prediction

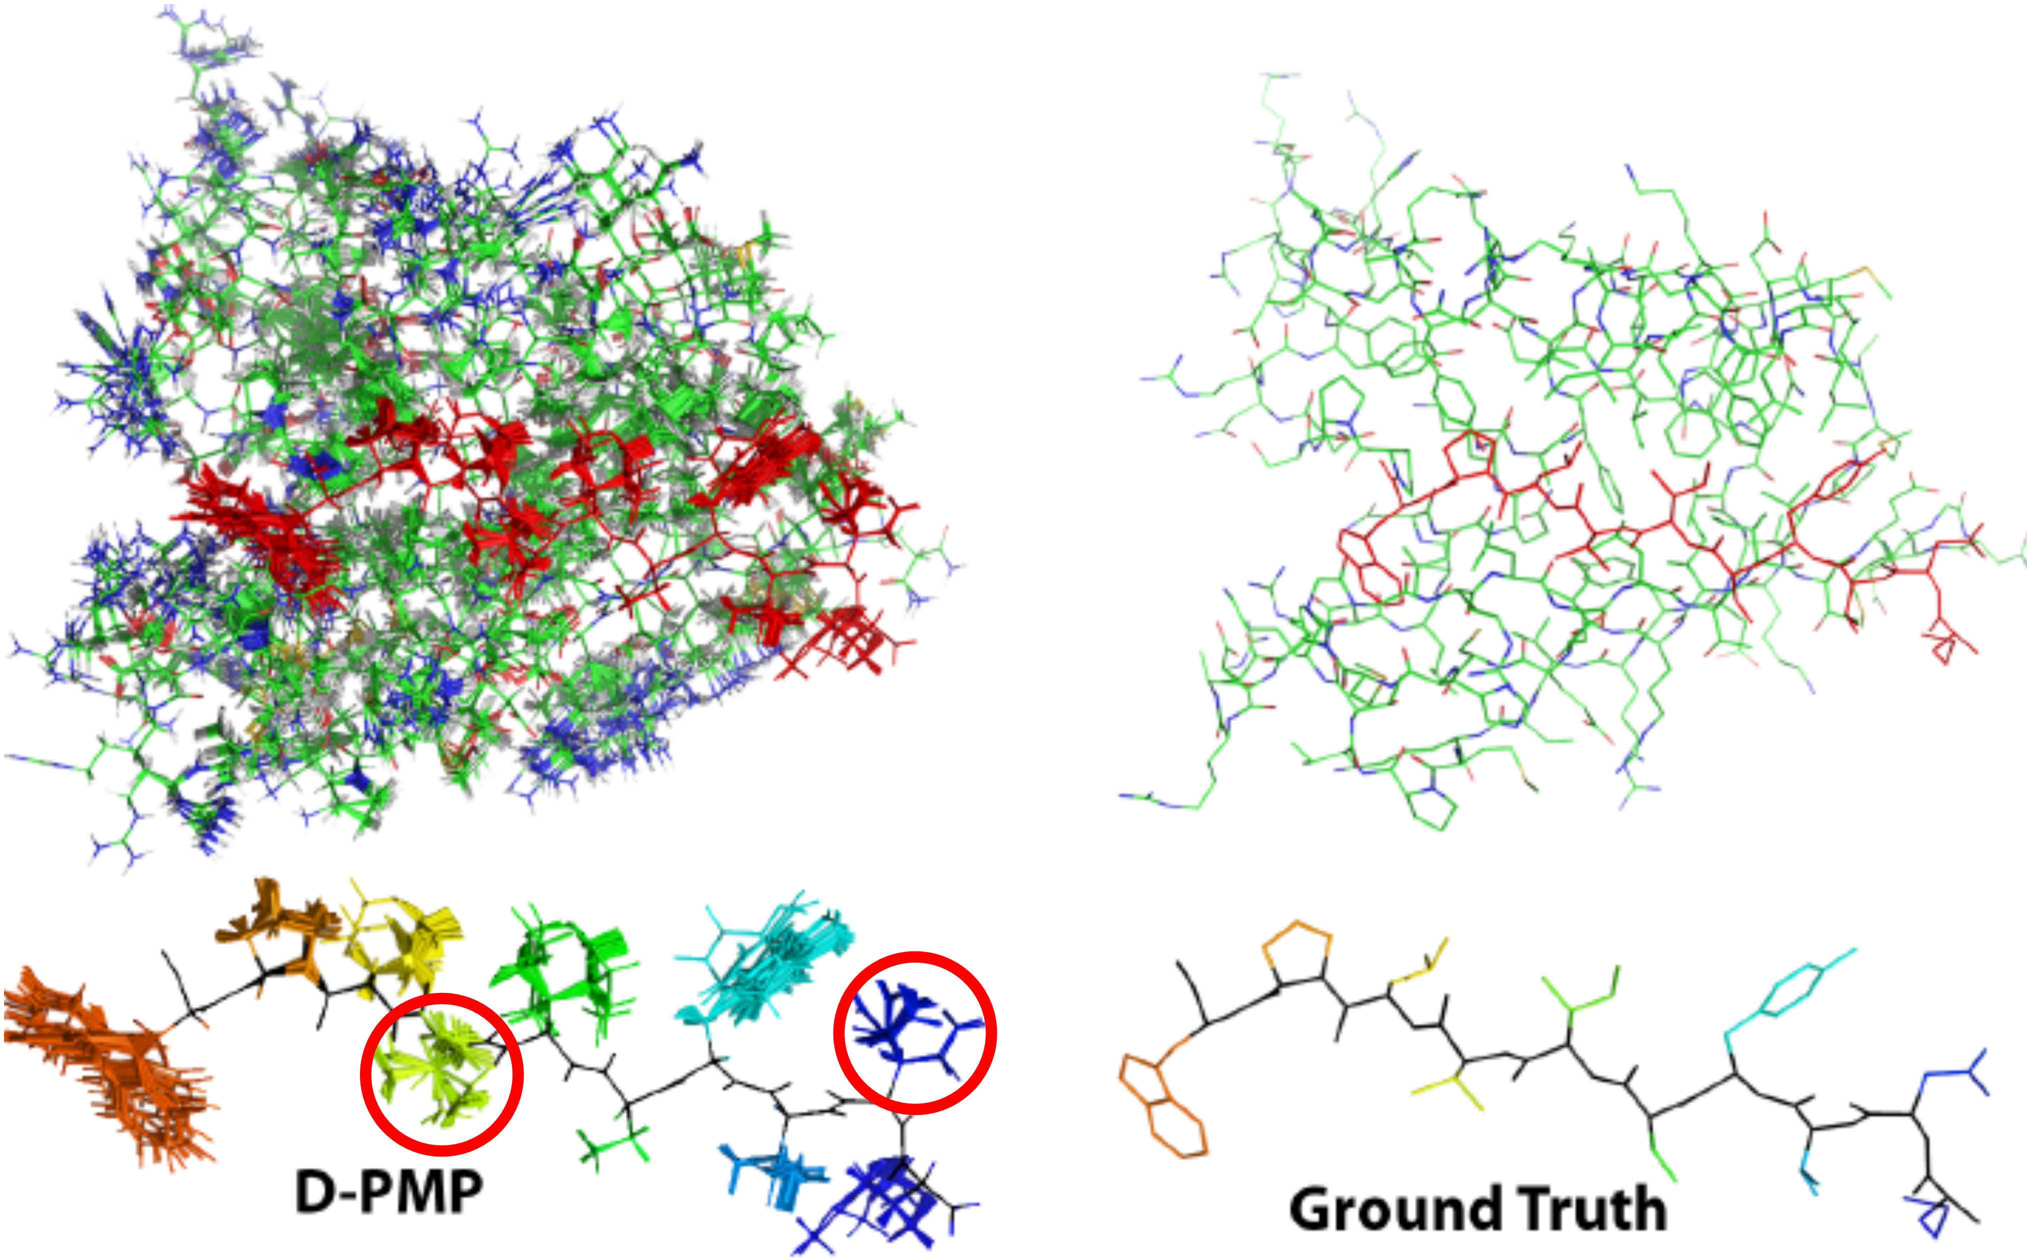

Contributions

Reliable particle-based MAP inference

for graphical models with continuous variables:

object shape, articulation, position, motion, …

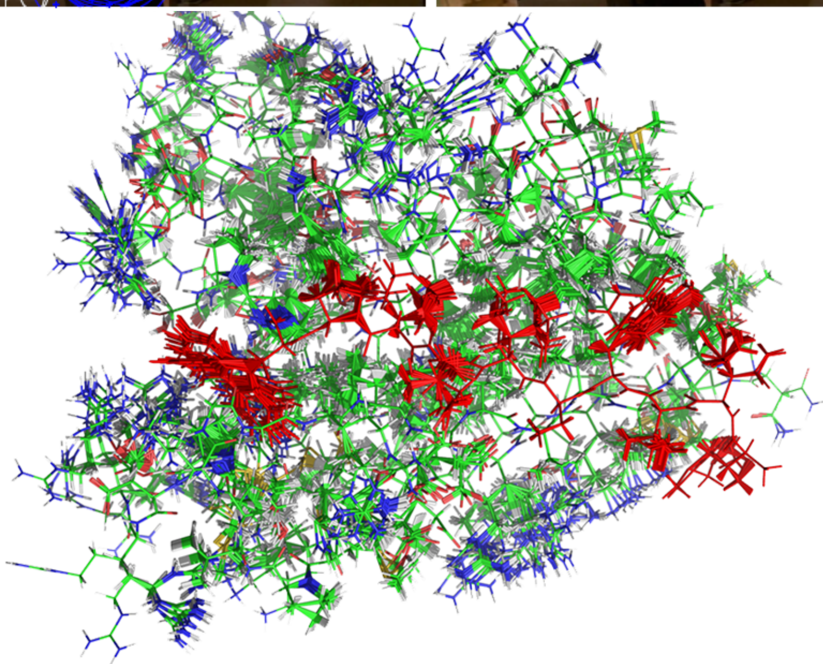

Validation: *Inference of multiple poses, motions, protein conformations, …*

Guarantees of Reliabilty: *Rigorous, non-asymptotic bounds on accuracy of diverse particle selection*

Code: *General-purpose, black-box inference for continuous graphical models*