

Tracking by Regression


Tomáš Svoboda

Department of Cybernetics, Center for Machine Perception

March 31, 2014



Linear mapping for tracking



$$\Phi(\text{img}_1) = (0, 0)^T$$

$$\Phi(\text{img}_2) = (12, 7)^T$$

$$\Phi(\text{img}_3) = (-14, 2)^T$$

$$\Phi(\text{img}_4) = (-9, 18)^T$$

$$\Phi(\text{img}_5) = (14, -14)^T$$

$$\Phi(\text{img}_6) = (-16, -12)^T$$



Optimization vs. regression based tracking

KLT:

$$\Delta \mathbf{p}^* = \operatorname{argmin}_{\Delta \mathbf{p} \in \mathcal{S}} \sum_{\mathbf{x}} [I(\mathbf{W}(\mathbf{x}; \mathbf{p} + \Delta \mathbf{p})) - T(\mathbf{x})]^2$$

Optimization in general:

$$\mathbf{p}^* = \operatorname{argmin}_{\mathbf{p} \in \mathcal{S}} f(\mathbf{p}; I, \mathbf{p}_0),$$

Regression based tracking:

$$\mathbf{p}^* = \varphi(I, \mathbf{p}_0)$$



Optimization vs. regression based tracking

KLT:

$$\Delta \mathbf{p}^* = \operatorname{argmin}_{\Delta \mathbf{p} \in \mathcal{S}} \sum_{\mathbf{x}} [I(\mathbf{W}(\mathbf{x}; \mathbf{p} + \Delta \mathbf{p})) - T(\mathbf{x})]^2$$

Optimization in general:

$$\mathbf{p}^* = \operatorname{argmin}_{\mathbf{p} \in \mathcal{S}} f(\mathbf{p}; I, \mathbf{p}_0),$$

Regression based tracking:

$$\mathbf{p}^* = \varphi(I, \mathbf{p}_0)$$



Optimization vs. regression based tracking

KLT:

$$\Delta \mathbf{p}^* = \operatorname{argmin}_{\Delta \mathbf{p} \in \mathcal{S}} \sum_{\mathbf{x}} [l(\mathbf{W}(\mathbf{x}; \mathbf{p} + \Delta \mathbf{p})) - T(\mathbf{x})]^2$$

Optimization in general:


$$\mathbf{p}^* = \operatorname{argmin}_{\mathbf{p} \in \mathcal{S}} f(\mathbf{p}; l, \mathbf{p}_0),$$

Regression based tracking:

$$\mathbf{p}^* = \varphi(l, \mathbf{p}_0)$$



Regression learned from data - sample image



$$\Phi(\text{img}_1) = (0, 0)^T$$

$$\Phi(\text{img}_2) = (12, 7)^T$$

$$\Phi(\text{img}_3) = (-14, 2)^T$$

$$\Phi(\text{img}_4) = (-9, 18)^T$$

$$\Phi(\text{img}_5) = (14, -14)^T$$

$$\Phi(\text{img}_6) = (-16, -12)^T$$



Regression learning

$$\varphi^* = \operatorname{argmin}_{\varphi} \sum_{\mathbf{p}} \|\varphi(I(\mathbf{p} \circ \mathbf{x})) - \mathbf{p}\|.$$



Connection to KLT

$$\Delta \mathbf{p} = \mathbf{H}^{-1} \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^{\top} [T(\mathbf{x}) - I(\mathbf{W}(\mathbf{x}; \mathbf{p}))]$$

where:

$$\mathbf{H} = \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^{\top} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]$$

Reformulating *regression*:

$$\mathbf{p} = \varphi(I(\mathbf{x})) = \mathbf{H}^{-1} (I(\mathbf{x}) - T(\mathbf{x}))$$

How to get \mathbf{H} ?



Connection to KLT

$$\Delta \mathbf{p} = \mathbf{H}^{-1} \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^{\top} [T(\mathbf{x}) - I(\mathbf{W}(\mathbf{x}; \mathbf{p}))]$$

where:

$$\mathbf{H} = \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^{\top} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]$$

Reformulating *regression*:

$$\mathbf{p} = \varphi(I(\mathbf{x})) = \mathbf{H}(I(\mathbf{x}) - T(\mathbf{x}))$$

How to get \mathbf{H} ?



Connection to KLT

$$\Delta \mathbf{p} = \mathbf{H}^{-1} \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^{\top} [T(\mathbf{x}) - I(\mathbf{W}(\mathbf{x}; \mathbf{p}))]$$

where:

$$\mathbf{H} = \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^{\top} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]$$

Reformulating *regression*:

$$\mathbf{p} = \varphi(I(\mathbf{x})) = \mathbf{H}(I(\mathbf{x}) - T(\mathbf{x}))$$

How to get H?



Collecting training data

Image template: $T = T(\mathbf{x})$

Collected training pairs (I^i, \mathbf{p}^i) ($i = 1 \dots d$) of observed intensities I^i and corresponding motion parameters \mathbf{p}^i , which align the object with current frame.

Training set is an ordered pair (\mathbf{I}, \mathbf{P}) , such that $\mathbf{I} = [I^1 - T, I^2 - T, \dots, I^d - T]$ and $\mathbf{P} = [\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^d]$.

The bolds \mathbf{I}, \mathbf{T} are vectorized I, T . Think about: $\mathbf{I} = I(:)$



Collecting training data

Image template: $T = T(\mathbf{x})$

Collected training pairs (I^i, \mathbf{p}^i) ($i = 1 \dots d$) of observed intensities I^i and corresponding motion parameters \mathbf{p}^i , which align the object with current frame.


Training set is an ordered pair (\mathbf{I}, \mathbf{P}) , such that

$\mathbf{I} = [I^1 - \mathbf{T}, I^2 - \mathbf{T}, \dots, I^d - \mathbf{T}]$ and $\mathbf{P} = [\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^d]$.

The bolds \mathbf{I}, \mathbf{T} are vectorized I, T . Think about: $\mathbf{I} = I(:)$



Learning alignment for one predictor



▶ $\varphi(\text{img}) = (0, 0)^T$

▶ $\varphi(\text{img}) = (-25, 0)^T$

▶ $\varphi(\text{img}) = (25, -15)^T$


▶ $\varphi(\text{img}) = (0, 0)^T$

▶ $\varphi(\text{img}) = (-25, 0)^T$

▶ $\varphi(\text{img}) = (25, -15)^T$



Learning alignment for one predictor



▶ $\varphi(\text{img}) = (0, 0)^T$

▶ $\varphi(\text{img}) = (-25, 0)^T$

▶ $\varphi(\text{img}) = (25, -15)^T$


▶ $\varphi(\text{img}) = (0, 0)^T$

▶ $\varphi(\text{img}) = (-25, 0)^T$

▶ $\varphi(\text{img}) = (25, -15)^T$



Learning alignment for one predictor



▶ $\varphi(\text{img}) = (0, 0)^T$

▶ $\varphi(\text{img}) = (-25, 0)^T$

▶ $\varphi(\text{img}) = (25, -15)^T$


▶ $\varphi(\text{img}) = (0, 0)^T$

▶ $\varphi(\text{img}) = (-25, 0)^T$


▶ $\varphi(\text{img}) = (25, -15)^T$



Learning alignment for one predictor




- ▶ $\varphi(\text{img}) = (0, 0)^T$
- ▶ $\varphi(\text{img}) = (-25, 0)^T$
- ▶ $\varphi(\text{img}) = (25, -15)^T$




- ▶ $\varphi(\text{img}) = (0, 0)^T$
- ▶ $\varphi(\text{img}) = (-25, 0)^T$
- ▶ $\varphi(\text{img}) = (25, -15)^T$



Learning alignment for one predictor




- ▶ $\varphi(\text{img}) = (0, 0)^T$
- ▶ $\varphi(\text{img}) = (-25, 0)^T$
- ▶ $\varphi(\text{img}) = (25, -15)^T$




- ▶ $\varphi(\text{img}) = (0, 0)^T$
- ▶ $\varphi(\text{img}) = (-25, 0)^T$
- ▶ $\varphi(\text{img}) = (25, -15)^T$



Learning alignment for one predictor



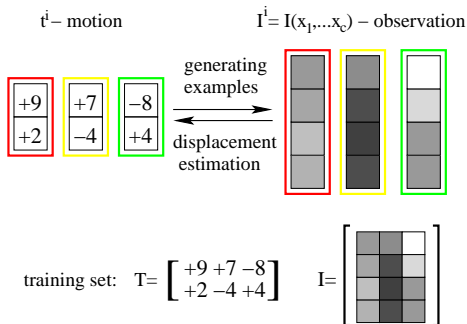

- ▶ $\varphi(\text{img}) = (0, 0)^T$
- ▶ $\varphi(\text{img}) = (-25, 0)^T$
- ▶ $\varphi(\text{img}) = (25, -15)^T$



- ▶ $\varphi(\text{img}) = (0, 0)^T$
- ▶ $\varphi(\text{img}) = (-25, 0)^T$
- ▶ $\varphi(\text{img}) = (25, -15)^T$



Generating training examples



Training set: (I, P)

$I = [I^1 - T, I^2 - T, \dots, I^d - T]$ and $P = [p^1, p^2, \dots, p^d]$.



Learning H from data

$$\begin{aligned} H^* &= \operatorname{argmin}_H \sum_{i=1}^d \|H(\mathbf{I}^i - \mathbf{T}) - \mathbf{p}^i\|_2^2 = \operatorname{argmin}_H \|H\mathbf{I} - \mathbf{P}\|_F^2 = \\ &= \operatorname{argmin}_H \operatorname{trace}(H\mathbf{I} - \mathbf{P})(H\mathbf{I} - \mathbf{P})^\top = \\ &= \operatorname{argmin}_H \operatorname{trace}(H\mathbf{I}\mathbf{I}^\top H^\top - 2H\mathbf{I}\mathbf{P}^\top + \mathbf{P}\mathbf{P}^\top). \end{aligned}$$

Next its derivative is set equal to zero:

$$\begin{aligned} 2H^*\mathbf{I}\mathbf{I}^\top - 2\mathbf{P}\mathbf{I}^\top &= 0 \\ H^*\mathbf{I}\mathbf{I}^\top &= \mathbf{P}\mathbf{I}^\top \\ H^* &= \underbrace{\mathbf{P}\mathbf{I}^\top(\mathbf{I}\mathbf{I}^\top)^{-1}}_{\mathbf{I}^+} = \mathbf{P}\mathbf{I}^+. \end{aligned}$$



Learning H from data

$$\begin{aligned} H^* &= \operatorname{argmin}_H \sum_{i=1}^d \|H(\mathbf{I}^i - \mathbf{T}) - \mathbf{p}^i\|_2^2 = \operatorname{argmin}_H \|HI - P\|_F^2 = \\ &= \operatorname{argmin}_H \operatorname{trace}(HI - P)(HI - P)^\top = \\ &= \operatorname{argmin}_H \operatorname{trace}(HII^\top H^\top - 2HIP^\top + PP^\top). \end{aligned}$$

Next its derivative is set equal to zero:

$$\begin{aligned} 2H^*II^\top - 2PI^\top &= 0 \\ H^*II^\top &= PI^\top \\ H^* &= \underbrace{PI^\top(II^\top)^{-1}}_{I^+} = PI^+. \end{aligned}$$



LS learning - summary

$$\varphi^* = \operatorname{argmin}_{\varphi} \sum_{\mathbf{p}} \|\varphi(I(\mathbf{p} \circ \mathbf{x})) - \mathbf{p}\|^2.$$

Minimizes sum of square errors over all training set. Leads to matrix pseudoinverse computation.

Example for *linear mapping*:

$$\mathbf{H}^* = \operatorname{argmin}_{\mathbf{H}} \sum_{i=1}^d \|\mathbf{H}(\mathbf{I}^i - \mathbf{T}) - \mathbf{t}^i\|_2^2 = \operatorname{argmin}_{\mathbf{H}} \|\mathbf{H}\mathbf{I} - \mathbf{P}\|_F^2$$

after some derivation

$$\mathbf{H}^* = \mathbf{P} \underbrace{\mathbf{I}^T (\mathbf{I}\mathbf{I}^T)^{-1}}_{\mathbf{I}^+} = \mathbf{P}\mathbf{I}^+.$$




Min-max learning

$$\varphi^* = \operatorname{argmin}_{\varphi} \max_{\mathbf{p}} \|\varphi(l(\mathbf{p} \circ \mathbf{x})) - \mathbf{p}\|_{\infty}.$$

Minimizes the *worst case* (the biggest estimation error) in the training set. Leads to linear programming.



Tracking of a single point by a *sequence* of predictors




$$t_1 = \hat{\Phi}_1 \left(\begin{array}{c} \text{gray bar} \\ \text{gray bar} \\ \text{gray bar} \\ \text{gray bar} \end{array} \right)$$


Old position



Tracking of a single point by a *sequence* of predictors



Tracking of a single point by a *sequence* of predictors




$$t_1 = \hat{\Phi}_1 \left(\begin{array}{c} \text{gray bar} \\ \text{gray bar} \\ \text{gray bar} \end{array} \right)$$


Old position



Tracking of a single point by a *sequence* of predictors



Tracking of a single point by a *sequence* of predictors




$$t_1 = \hat{\Phi}_1 \left(\begin{array}{c} \text{gray bar} \\ \text{gray bar} \\ \text{gray bar} \end{array} \right)$$

$$t_2 = \hat{\Phi}_2 \left(\begin{array}{c} \text{gray bar} \\ \text{gray bar} \end{array} \right)$$

Old position



Tracking of a single point by a *sequence* of predictors




$$t_1 = \hat{\Phi}_1 \left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right)$$

$$t_2 = \hat{\Phi}_2 \left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right)$$

Old position



Tracking of a single point by a *sequence* of predictors




$$t_1 = \hat{\Phi}_1 \left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right)$$

$$t_2 = \hat{\Phi}_2 \left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right)$$


Old position



Tracking of a single point by a *sequence* of predictors



Tracking of a single point by a *sequence* of predictors



$$t_1 = \hat{\Phi}_1 \left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right)$$


$$t_2 = \hat{\Phi}_2 \left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \right)$$

$$t_3 = \hat{\Phi}_3 \left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right)$$

Old position



Tracking of a single point by a *sequence* of predictors



$$t_1 = \hat{\Phi}_1 \left(\begin{array}{c} \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \end{array} \right)$$

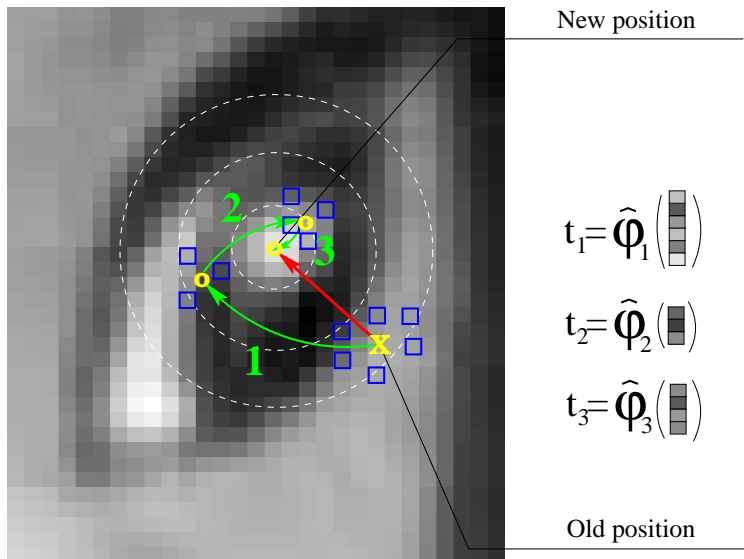
$$t_2 = \hat{\Phi}_2 \left(\begin{array}{c} \blacksquare \\ \blacksquare \\ \blacksquare \end{array} \right)$$

$$t_3 = \hat{\Phi}_3 \left(\begin{array}{c} \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \end{array} \right)$$

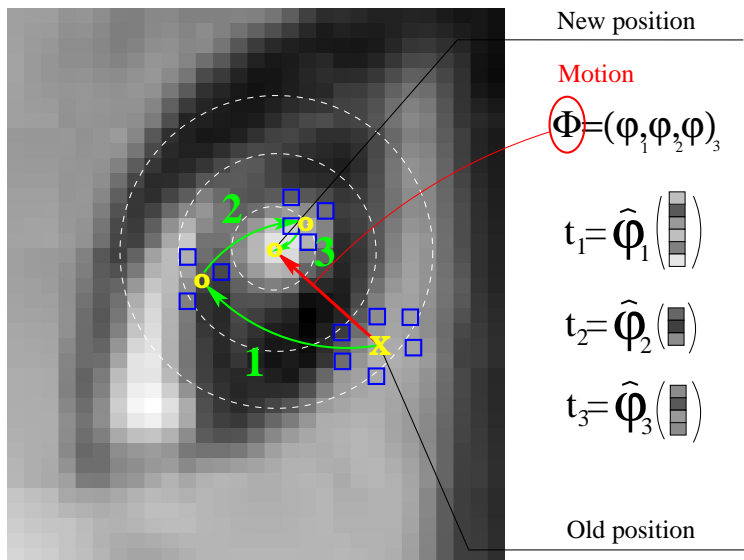
Old position



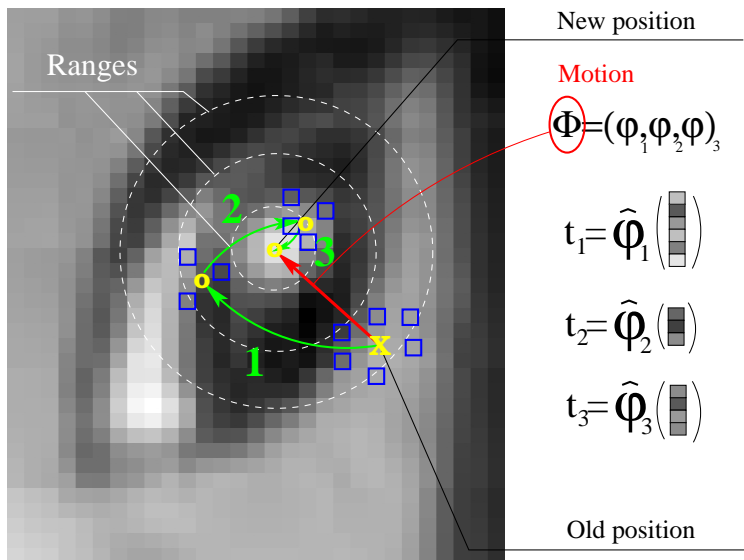
Tracking of a single point by a *sequence* of predictors



Tracking of a single point by a *sequence* of predictors



Tracking of a single point by a *sequence* of predictors



Learning of sequential predictor

- ▶ **Learning** - searching for the sequence with predefined *range*, *accuracy* and minimal *computational cost*.
 - ▶ [Zimmermann-PAMI-2009] - Dynamic programming estimates the optimal sequence of linear predictors.
 - ▶ [Zimmermann-IVC-2009] - Branch & bound search allows for time constrained learning (demo in MATLAB).

[Zimmermann-PAMI-2009] K.Zimmermann, J.Matas, T.Svoboda. Tracking by an Optimal Sequence of Linear Predictors, in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, IEEE computer society, 2009, vol. 31, No 4, pp 677–692.

[Zimmermann-IVC-2009] K.Zimmermann, T.Svoboda, J.Matas. Anytime learning for the NoSLLiP tracker. *Image and Vision Computing*, vol. 27, No 11



Learning of sequential predictor

- ▶ **Learning** - searching for the sequence with predefined *range*, *accuracy* and minimal *computational cost*.
 - ▶ [\[Zimmermann-PAMI-2009\]](#) - Dynamic programming estimates the optimal sequence of linear predictors.
 - ▶ [\[Zimmermann-IVC-2009\]](#) - Branch & bound search allows for time constrained learning (demo in MATLAB).

[\[Zimmermann-PAMI-2009\]](#) K.Zimmermann, J.Matas, T.Svoboda. Tracking by an Optimal Sequence of Linear Predictors, in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, IEEE computer society, 2009, vol. 31, No 4, pp 677–692.

[\[Zimmermann-IVC-2009\]](#) K.Zimmermann, T.Svoboda, J.Matas. Anytime learning for the NoSLLiP tracker. *Image and Vision Computing*, vol. 27, No 11



Learning of sequential predictor


- ▶ **Learning** - searching for the sequence with predefined *range*, *accuracy* and minimal *computational cost*.
 - ▶ [\[Zimmermann-PAMI-2009\]](#) - Dynamic programming estimates the optimal sequence of linear predictors.
 - ▶ [\[Zimmermann-IVC-2009\]](#) - Branch & bound search allows for time constrained learning (demo in MATLAB).

[\[Zimmermann-PAMI-2009\]](#) K.Zimmermann, J.Matas, T.Svoboda. Tracking by an Optimal Sequence of Linear Predictors, in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, IEEE computer society, 2009, vol. 31, No 4, pp 677–692.

[\[Zimmermann-IVC-2009\]](#) K.Zimmermann, T.Svoboda, J.Matas. Anytime learning for the NoSLLiP tracker. *Image and Vision Computing*, vol. 27, No 11



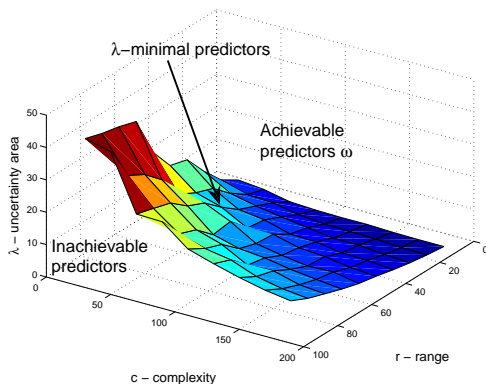
Learning of the optimal sequence of linear predictors



- ▶ **Range:** the set of admissible motions, r .
- ▶ **Complexity:** cardinality of support set, c .
- ▶ **Uncertainty region:** the region within which all predictions lie, λ . Small red circles show acceptable uncertainty.




Learning of the optimal sequence of linear predictors



- ▶ **Range:** the set of admissible motions, r .
- ▶ **Complexity:** cardinality of support set, c .
- ▶ **Uncertainty region:** the region within which all predictions lie, λ . Small red circles show acceptable uncertainty.




Learning of the optimal sequence of linear predictors



- ▶ **Range:** the set of admissible motions, r .
- ▶ **Complexity:** cardinality of support set, c .
- ▶ **Uncertainty region:** the region within which all predictions lie, λ . Small red circles show acceptable uncertainty.





Learning of the optimal sequence of linear predictors



- ▶ **Range:** the set of admissible motions, r .
- ▶ **Complexity:** cardinality of support set, c .
- ▶ **Uncertainty region:** the region within which all predictions lie, λ . Small red circles show acceptable uncertainty.





Branch and Bound



Don't forget to show the live demo!




Branch and Bound




Don't forget to show the live demo!



Tracking with one linear predictor.



Modeling motion by number of linear predictors.



Motion blur, fast motion, views from acute angles and other image distortions.

