

RWTH AACHEN
UNIVERSITY

Computer Vision II - Lecture 8

Tracking with Linear Dynamic Models

20.05.2014

Bastian Leibe
RWTH Aachen
<http://www.vision.rwth-aachen.de>
leibe@vision.rwth-aachen.de

Computer Vision II, Summer'14

RWTH AACHEN
UNIVERSITY

Course Outline

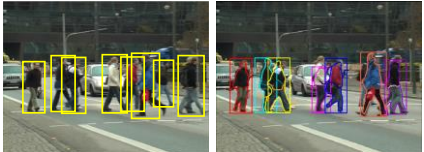
- Single-Object Tracking
 - Background modeling
 - Template based tracking
 - Color based tracking
 - Contour based tracking
 - Tracking by online classification
 - Tracking-by-detection
- Bayesian Filtering
 - Kalman filter
 - Particle filter
- Multi-Object Tracking
- Articulated Tracking

2
Image source: Helmut Grabner, Disney/Pixar

Computer Vision II, Summer'14

RWTH AACHEN
UNIVERSITY

Recap: Tracking-by-Detection



- Main ideas
 - Apply a generic object detector to find objects of a certain class
 - Based on the detections, extract object appearance models
 - Link detections into trajectories

3
B. Leibe

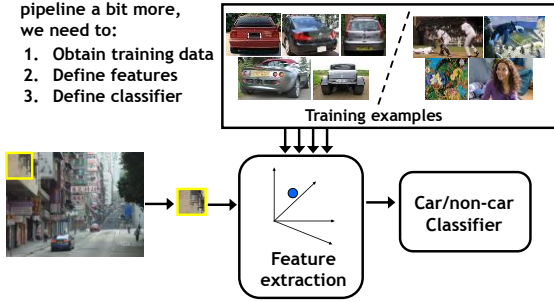
Computer Vision II, Summer'14

RWTH AACHEN
UNIVERSITY

Recap: Sliding-Window Object Detection

Fleshing out this pipeline a bit more, we need to:

1. Obtain training data
2. Define features
3. Define classifier



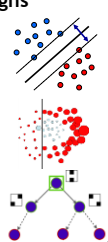
4
Slide credit: Kristen Grauman
B. Leibe

Computer Vision II, Summer'14

RWTH AACHEN
UNIVERSITY

Recap: Object Detector Design

- In practice, the classifier often determines the design.
 - Types of features
 - Speedup strategies
- We'll look at 3 state-of-the-art detector designs
 - Based on SVMs → Last lecture
 - Based on Boosting → Last lecture
 - Based on Random Forests → Postponed to a later slot...



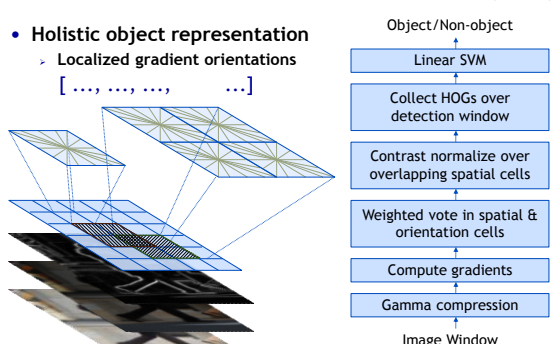
5
B. Leibe

Computer Vision II, Summer'14

RWTH AACHEN
UNIVERSITY

Recap: Histograms of Oriented Gradients (HOG)

- Holistic object representation
 - Localized gradient orientations [..., ..., ..., ...]



6
Slide adapted from Navneet Dalal

Computer Vision II, Summer'14

RWTH AACHEN UNIVERSITY

Recap: Deformable Part-based Model (DPM)

Score of filter: dot product of filter with HOG features underneath it

Score of object hypothesis is sum of filter scores minus deformation costs

- Multiscale model captures features at two resolutions

Computer Vision II, Summer'14 | Slide credit: Pedro Felzenszwalb | B. Leibe | [Felzenszwalb, McAllister, Ramanan, CVPR'08]

RWTH AACHEN UNIVERSITY

Recap: DPM Hypothesis Score

$$\text{score}(p_0, \dots, p_n) = \sum_{i=0}^n F_i \cdot \phi(H, p_i) - \sum_{i=1}^n d_i \cdot (dx_i^2, dy_i^2)$$

"data term" "spatial prior"
↑ filters ↑ displacements
↑ deformation parameters

$$\text{score}(z) = \beta \cdot \Psi(H, z)$$

↑ concatenation filters and deformation parameters ↑ concatenation of HOG features and part displacement features

Computer Vision II, Summer'14 | Slide credit: Pedro Felzenszwalb | B. Leibe | [Felzenszwalb, McAllister, Ramanan, CVPR'08]

RWTH AACHEN UNIVERSITY

Recap: Integral Channel Features

- Generalization of Haar Wavelet idea from Viola-Jones
 - Instead of only considering intensities, also take into account other feature channels (gradient orientations, color, texture).
 - Still efficiently represented as integral images.

P. Dollár, Z. Tu, P. Perona, S. Belongie. [Integral Channel Features](#), BMVC'09.

Computer Vision II, Summer'14 | B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Integral Channel Features

- Generalize also block computation
 - 1st order features:
 - Sum of pixels in rectangular region.
 - 2nd-order features:
 - Haar-like difference of sum-over-blocks
 - Generalized Haar:
 - More complex combinations of weighted rectangles
 - Histograms
 - Computed by evaluating local sums on quantized images.

Computer Vision II, Summer'14 | B. Leibe

RWTH AACHEN UNIVERSITY

Recap: VeryFast Detector

- Idea 1: Invert the relation

1 model, 50 image scales \rightarrow 50 models, 1 image scale

R. Benenson, M. Mathias, R. Timofte, L. Van Gool. [Pedestrian Detection at 100 Frames per Second](#), CVPR'12.

Computer Vision II, Summer'14 | Slide credit: Rodrigo Benenson | B. Leibe

RWTH AACHEN UNIVERSITY

Recap: VeryFast Detector

- Idea 2: Reduce training time by feature interpolation

5 models, 1 image scale \approx 50 models, 1 image scale

- Shown to be possible for Integral Channel features
 - P. Dollár, S. Belongie, Perona. [The Fastest Pedestrian Detector in the West](#), BMVC 2010.

Computer Vision II, Summer'14 | Slide adapted from Rodrigo Benenson | B. Leibe

Computer Vision II, Summer'14

RWTH AACHEN UNIVERSITY

Recap: VeryFast Classifier Construction

6 Orientation bins Gradient magnitude LUV color channels

score = $w_1 \cdot h_1 + w_2 \cdot h_2 + \dots + w_N \cdot h_N$

- Ensemble of short trees, learned by AdaBoost

Slide credit: Rodrigo Benenson B. Leibe 13

Computer Vision II, Summer'14

RWTH AACHEN UNIVERSITY

Elements of Tracking

Detection Data association Prediction

- Detection
 - Where are candidate objects?
- Data association
 - Which detection corresponds to which object?
- Prediction
 - Where will the tracked object be in the next time step?

Last lecture Today's topic

Computer Vision II, Summer'14 B. Leibe 14

Computer Vision II, Summer'14

RWTH AACHEN UNIVERSITY

Today: Tracking with Linear Dynamic Models

Figure from Forsyth & Ponce

B. Leibe 15

Computer Vision II, Summer'14

RWTH AACHEN UNIVERSITY

Topics of This Lecture

- Tracking with Dynamics
 - Detection vs. Tracking
 - Tracking as probabilistic inference
 - Prediction and Correction
- Linear Dynamic Models
 - Zero velocity model
 - Constant velocity model
 - Constant acceleration model
- The Kalman Filter
 - Kalman filter for 1D state
 - General Kalman filter
 - Limitations

Computer Vision II, Summer'14 B. Leibe 16

Computer Vision II, Summer'14

RWTH AACHEN UNIVERSITY

Tracking with Dynamics

- Key idea
 - Given a model of expected motion, predict where objects will occur in next frame, even before seeing the image.
- Goals
 - Restrict search for the object
 - Improved estimates since measurement noise is reduced by trajectory smoothness.
- Assumption: continuous motion patterns
 - Camera is not moving instantly to new viewpoint.
 - Objects do not disappear and reappear in different places.
 - Gradual change in pose between camera and scene.

Slide adapted from S. Lazebnik, K. Grauman B. Leibe 17

Computer Vision II, Summer'14

RWTH AACHEN UNIVERSITY

General Model for Tracking

- Representation
 - The moving object of interest is characterized by an underlying state X .
 - State X gives rise to measurements or observations Y .
 - At each time t , the state changes to X_t and we get a new observation Y_t .

Slide credit: Svetlana Lazebnik B. Leibe 18

RWTH AACHEN UNIVERSITY

State vs. Observation

The diagram shows three parts: a 3D plot of a trajectory, a photograph of a bowling ball hitting pins, and a diagram of a water level measurement. The water level is labeled as the 'State variable' and the measurement as 'd'.

- Hidden state : parameters of interest
- Measurement: what we get to directly observe

19

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Tracking as Inference

- Inference problem
 - The hidden state consists of the true parameters we care about, denoted X .
 - The measurement is our noisy observation that results from the underlying state, denoted Y .
 - At each time step, state changes (from X_{t-1} to X_t) and we get a new observation Y_t .
- Our goal: recover most likely state X_t given
 - All observations seen so far.
 - Knowledge about dynamics of state transitions.

20

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Steps of Tracking

- Prediction: What is the next state of the object given past measurements?

$$P(X_t | Y_0 = y_0, \dots, Y_{t-1} = y_{t-1})$$
- Correction: Compute an updated estimate of the state from prediction and measurements.

$$P(X_t | Y_0 = y_0, \dots, Y_{t-1} = y_{t-1}, Y_t = y_t)$$
- Tracking can be seen as the process of propagating the posterior distribution of state given measurements across time.

21

Slide credit: Svetlana Lazebnik B. Leibe

RWTH AACHEN UNIVERSITY

Simplifying Assumptions

- Only the immediate past matters

$$P(X_t | X_0, \dots, X_{t-1}) = P(X_t | X_{t-1})$$

Dynamics model
- Measurements depend only on the current state

$$P(Y_t | X_0, Y_0, \dots, X_{t-1}, Y_{t-1}, X_t) = P(Y_t | X_t)$$

Observation model

22

Slide credit: Svetlana Lazebnik B. Leibe

RWTH AACHEN UNIVERSITY

Tracking as Induction

- Base case:
 - Assume we have initial prior that predicts state in absence of any evidence: $P(X_0)$
 - At the first frame, correct this given the value of $Y_0 = y_0$.

$$P(X_0 | Y_0 = y_0) = \frac{P(y_0 | X_0)P(X_0)}{P(y_0)} \propto P(y_0 | X_0)P(X_0)$$

Posterior prob. of state given measurement
Likelihood of measurement
Prior of the state

23

Slide credit: Svetlana Lazebnik B. Leibe

RWTH AACHEN UNIVERSITY

Tracking as Induction

- Base case:
 - Assume we have initial prior that predicts state in absence of any evidence: $P(X_0)$
 - At the first frame, correct this given the value of $Y_0 = y_0$.
- Given corrected estimate for frame t :
 - Predict for frame $t+1$
 - Correct for frame $t+1$

24

Slide credit: Svetlana Lazebnik B. Leibe

RWTH AACHEN UNIVERSITY

Induction Step: Prediction

- Prediction involves representing $P(X_t | y_0, \dots, y_{t-1})$ given $P(X_{t-1} | y_0, \dots, y_{t-1})$

$$P(X_t | y_0, \dots, y_{t-1}) = \int P(X_t, X_{t-1} | y_0, \dots, y_{t-1}) dX_{t-1}$$

Law of total probability

$$P(A) = \int P(A, B) dB$$

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 25

RWTH AACHEN UNIVERSITY

Induction Step: Prediction

- Prediction involves representing $P(X_t | y_0, \dots, y_{t-1})$ given $P(X_{t-1} | y_0, \dots, y_{t-1})$

$$P(X_t | y_0, \dots, y_{t-1}) = \int P(X_t, X_{t-1} | y_0, \dots, y_{t-1}) dX_{t-1} = \int P(X_t | X_{t-1}, y_0, \dots, y_{t-1}) P(X_{t-1} | y_0, \dots, y_{t-1}) dX_{t-1}$$

Conditioning on X_{t-1}

$$P(A, B) = P(A | B) P(B)$$

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 26

RWTH AACHEN UNIVERSITY

Induction Step: Prediction

- Prediction involves representing $P(X_t | y_0, \dots, y_{t-1})$ given $P(X_{t-1} | y_0, \dots, y_{t-1})$

$$P(X_t | y_0, \dots, y_{t-1}) = \int P(X_t, X_{t-1} | y_0, \dots, y_{t-1}) dX_{t-1} = \int P(X_t | X_{t-1}, y_0, \dots, y_{t-1}) P(X_{t-1} | y_0, \dots, y_{t-1}) dX_{t-1} = \int P(X_t | X_{t-1}) P(X_{t-1} | y_0, \dots, y_{t-1}) dX_{t-1}$$

Independence assumption

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 27

RWTH AACHEN UNIVERSITY

Induction Step: Correction

- Correction involves computing $P(X_t | y_0, \dots, y_t)$ given predicted value $P(X_t | y_0, \dots, y_{t-1})$

$$P(X_t | y_0, \dots, y_t) = \frac{P(y_t | X_t, y_0, \dots, y_{t-1}) P(X_t | y_0, \dots, y_{t-1})}{P(y_t | y_0, \dots, y_{t-1})}$$

Bayes rule

$$P(A | B) = \frac{P(B | A) P(A)}{P(B)}$$

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 28

RWTH AACHEN UNIVERSITY

Induction Step: Correction

- Correction involves computing $P(X_t | y_0, \dots, y_t)$ given predicted value $P(X_t | y_0, \dots, y_{t-1})$

$$P(X_t | y_0, \dots, y_t) = \frac{P(y_t | X_t, y_0, \dots, y_{t-1}) P(X_t | y_0, \dots, y_{t-1})}{P(y_t | y_0, \dots, y_{t-1})} = \frac{P(y_t | X_t) P(X_t | y_0, \dots, y_{t-1})}{P(y_t | y_0, \dots, y_{t-1})}$$

Independence assumption (observation y_t depends only on state X_t)

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 29

RWTH AACHEN UNIVERSITY

Induction Step: Correction

- Correction involves computing $P(X_t | y_0, \dots, y_t)$ given predicted value $P(X_t | y_0, \dots, y_{t-1})$

$$P(X_t | y_0, \dots, y_t) = \frac{P(y_t | X_t, y_0, \dots, y_{t-1}) P(X_t | y_0, \dots, y_{t-1})}{P(y_t | y_0, \dots, y_{t-1})} = \frac{P(y_t | X_t) P(X_t | y_0, \dots, y_{t-1})}{P(y_t | y_0, \dots, y_{t-1})} = \int P(y_t | X_t) P(X_t | y_0, \dots, y_{t-1}) dX_t$$

Conditioning on X_t

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 30

RWTH AACHEN UNIVERSITY

Summary: Prediction and Correction

- Prediction:**

$$P(X_t | y_0, \dots, y_{t-1}) = \int \underbrace{P(X_t | X_{t-1})}_{\text{Dynamics model}} \underbrace{P(X_{t-1} | y_0, \dots, y_{t-1})}_{\text{Corrected estimate from previous step}} dX_{t-1}$$

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 31

RWTH AACHEN UNIVERSITY

Summary: Prediction and Correction

- Prediction:**

$$P(X_t | y_0, \dots, y_{t-1}) = \int \underbrace{P(X_t | X_{t-1})}_{\text{Dynamics model}} \underbrace{P(X_{t-1} | y_0, \dots, y_{t-1})}_{\text{Corrected estimate from previous step}} dX_{t-1}$$

- Correction:**

$$P(X_t | y_0, \dots, y_t) = \frac{\underbrace{P(y_t | X_t)}_{\text{Observation model}} \underbrace{P(X_t | y_0, \dots, y_{t-1})}_{\text{Predicted estimate}}}{\int P(y_t | X_t) P(X_t | y_0, \dots, y_{t-1}) dX_t}$$

Computer Vision II, Summer'14

Slide credit: Svetlana Lazebnik B. Leibe 32

RWTH AACHEN UNIVERSITY

Topics of This Lecture

- Tracking with Dynamics
 - Detection vs. Tracking
 - Tracking as probabilistic inference
 - Prediction and Correction
- Linear Dynamic Models
 - Zero velocity model
 - Constant velocity model
 - Constant acceleration model
- The Kalman Filter
 - Kalman filter for 1D state
 - General Kalman filter
 - Limitations

Computer Vision II, Summer'14

B. Leibe 33

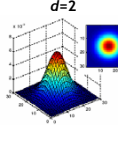
RWTH AACHEN UNIVERSITY

Notation Reminder

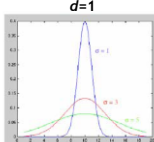
$$\mathbf{x} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

- Random variable with Gaussian probability distribution that has the mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.
- \mathbf{x} and $\boldsymbol{\mu}$ are d -dimensional, $\boldsymbol{\Sigma}$ is $d \times d$.

$d=2$



$d=1$



If \mathbf{x} is 1D, we just have one $\boldsymbol{\Sigma}$ parameter: the variance σ^2

Computer Vision II, Summer'14

Slide credit: Kristen Grauman B. Leibe 34

RWTH AACHEN UNIVERSITY

Linear Dynamic Models

- Dynamics model**
 - State undergoes linear transformation D_t plus Gaussian noise
$$\mathbf{x}_t \sim N(\underbrace{D_t}_{n \times n} \mathbf{x}_{t-1}, \underbrace{\Sigma_{d_t}}_{n \times n})$$
- Observation model**
 - Measurement is linearly transformed state plus Gaussian noise
$$\mathbf{y}_t \sim N(\underbrace{M_t}_{m \times n} \mathbf{x}_t, \underbrace{\Sigma_{m_t}}_{m \times m})$$

Computer Vision II, Summer'14

Slide credit: S. Lazebnik, K. Grauman B. Leibe 35



RWTH AACHEN UNIVERSITY

Example: Randomly Drifting Points

- Consider a stationary object, with state as position.
 - Position is constant, only motion due to random noise term.
$$x_t = p_t \quad p_t = p_{t-1} + \epsilon$$

\Rightarrow State evolution is described by identity matrix $D=I$

$$x_t = D_t x_{t-1} + \text{noise} = I p_{t-1} + \text{noise}$$

Computer Vision II, Summer'14

Slide credit: Kristen Grauman B. Leibe 36

Computer Vision II, Summer'14

Example: Constant Velocity (1D Points)

Slide credit: Kristen Grauman B. Leibe Figure from Forsyth & Ponce

Computer Vision II, Summer'14

Example: Constant Velocity (1D Points)

- State vector: position p and velocity v

$$x_t = \begin{bmatrix} p_t \\ v_t \end{bmatrix} \quad p_t =$$

(greek letters denote noise terms)

$$x_t = D_t x_{t-1} + noise =$$
- Measurement is position only
$$y_t = M x_t + noise =$$

Slide credit: S. Lazebnik, K. Grauman B. Leibe

Computer Vision II, Summer'14

Example: Constant Acceleration (1D Points)

Slide credit: Kristen Grauman B. Leibe Figure from Forsyth & Ponce

Computer Vision II, Summer'14

Example: Constant Acceleration (1D Points)

- State vector: position p , velocity v , and acceleration a .
$$x_t = \begin{bmatrix} p_t \\ v_t \\ a_t \end{bmatrix} \quad \begin{aligned} p_t &= p_{t-1} + (\Delta t)v_{t-1} + \varepsilon \\ v_t &= \\ a_t &= \end{aligned}$$

(greek letters denote noise terms)

$$x_t = D_t x_{t-1} + noise =$$
- Measurement is position only
$$y_t = M x_t + noise =$$

Slide credit: S. Lazebnik, K. Grauman

Computer Vision II, Summer'14

Example: General Motion Models

- Assuming we have differential equations for the motion
 - E.g. for (undamped) periodic motion of a pendulum
$$\frac{d^2 p}{dt^2} = -p$$
- Substitute variables to transform this into linear system
$$p_1 = p \quad p_2 = \frac{dp}{dt} \quad p_3 = \frac{d^2 p}{dt^2}$$
- Then we have
$$x_t = \begin{bmatrix} p_{1,t} \\ p_{2,t} \\ p_{3,t} \end{bmatrix} \quad \begin{aligned} p_{1,t} &= p_{1,t-1} + (\Delta t)p_{2,t-1} + \varepsilon \\ p_{2,t} &= p_{2,t-1} + (\Delta t)p_{3,t-1} + \zeta \\ p_{3,t} &= -p_{1,t-1} + \zeta \end{aligned} \quad D_t = \begin{bmatrix} 1 & \Delta t & 0 \\ 0 & 1 & \Delta t \\ -1 & 0 & 0 \end{bmatrix}$$

B. Leibe

Computer Vision II, Summer'14

Topics of This Lecture

- Tracking with Dynamics
 - Detection vs. Tracking
 - Tracking as probabilistic inference
 - Prediction and Correction
- Linear Dynamic Models
 - Zero velocity model
 - Constant velocity model
 - Constant acceleration model
- The Kalman Filter
 - Kalman filter for 1D state
 - General Kalman filter
 - Limitations

B. Leibe

RWTH AACHEN UNIVERSITY

The Kalman Filter

- Kalman filter
 - Method for tracking linear dynamical models in Gaussian noise
- The predicted/corrected state distributions are Gaussian
 - You only need to maintain the mean and covariance.
 - The calculations are easy (all the integrals can be done in closed form).

Computer Vision II, Summer'14

45

Slide credit: Svetlana Lazebnik B. Leibe

RWTH AACHEN UNIVERSITY

The Kalman Filter

Know corrected state from previous time step, and all measurements up to the current one
 → Predict distribution over next state.

Receive measurement

Know prediction of state and next measurement
 → Update distribution over current state.

Time update
("Predict")

Measurement update
("Correct")

$P(X_t | y_0, \dots, y_{t-1})$
 Mean and std. dev. of predicted state:
 μ_t^-, σ_t^-

Time advances: $t++$

$P(X_t | y_0, \dots, y_t)$
 Mean and std. dev. of corrected state:
 μ_t^+, σ_t^+

Computer Vision II, Summer'14

46

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Kalman Filter for 1D State

Want to represent and update

$$P(x_t | y_0, \dots, y_{t-1}) = N(\mu_t^-, (\sigma_t^-)^2)$$

$$P(x_t | y_0, \dots, y_t) = N(\mu_t^+, (\sigma_t^+)^2)$$

Computer Vision II, Summer'14

47

B. Leibe

RWTH AACHEN UNIVERSITY

Propagation of Gaussian densities

Computer Vision II, Summer'14

48

Slide credit: Svetlana Lazebnik

RWTH AACHEN UNIVERSITY

1D Kalman Filter: Prediction

- Have linear dynamic model defining predicted state evolution, with noise

$$X_t \sim N(dx_{t-1}, \sigma_d^2)$$
- Want to estimate predicted distribution for next state

$$P(X_t | y_0, \dots, y_{t-1}) = N(\mu_t^-, (\sigma_t^-)^2)$$
- Update the mean:

$$\mu_t^- = d\mu_{t-1}^+$$

for derivations, see F&P Chapter 17.3
- Update the variance:

$$(\sigma_t^-)^2 = \sigma_d^2 + (d\sigma_{t-1}^+)^2$$

Computer Vision II, Summer'14

49

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

1D Kalman Filter: Correction

- Have linear model defining the mapping of state to measurements:

$$Y_t \sim N(mx_t, \sigma_m^2)$$
- Want to estimate corrected distribution given latest measurement:

$$P(X_t | y_0, \dots, y_t) = N(\mu_t^+, (\sigma_t^+)^2)$$
- Update the mean:

$$\mu_t^+ = \frac{\mu_t^- \sigma_m^2 + m y_t (\sigma_t^-)^2}{\sigma_m^2 + m^2 (\sigma_t^-)^2}$$
- Update the variance:

$$(\sigma_t^+)^2 = \frac{\sigma_m^2 (\sigma_t^-)^2}{\sigma_m^2 + m^2 (\sigma_t^-)^2}$$

Computer Vision II, Summer'14

50

Slide credit: Kristen Grauman B. Leibe Derivations: F&P Chapter 17.3

RWTH AACHEN UNIVERSITY

Prediction vs. Correction

$$\mu_t^+ = \frac{\mu_t^- \sigma_m^2 + m y_t (\sigma_t^-)^2}{\sigma_m^2 + m^2 (\sigma_t^-)^2} \quad (\sigma_t^+)^2 = \frac{\sigma_m^2 (\sigma_t^-)^2}{\sigma_m^2 + m^2 (\sigma_t^-)^2}$$

- What if there is no prediction uncertainty ($\sigma_t^- = 0$)?

$$\mu_t^+ = \mu_t^- \quad (\sigma_t^+)^2 = 0$$

The measurement is ignored!
- What if there is no measurement uncertainty ($\sigma_m = 0$)?

$$\mu_t^+ = \frac{y_t}{m} \quad (\sigma_t^+)^2 = 0$$

The prediction is ignored!

Computer Vision II, Summer'14
Slide credit: Kristen Grauman
B. Leibe
51

RWTH AACHEN UNIVERSITY

Recall: Constant Velocity Example

position

time

measurements

state

State is 2D: position + velocity
Measurement is 1D: position

Computer Vision II, Summer'14
Slide credit: Kristen Grauman
B. Leibe
Figure from Forsyth & Ponce
52

RWTH AACHEN UNIVERSITY

Constant Velocity Model

o state

x measurement

* predicted mean estimate

+ corrected mean estimate

bars: variance estimates before and after measurements

Computer Vision II, Summer'14
Slide credit: Kristen Grauman
B. Leibe
Figure from Forsyth & Ponce
53

RWTH AACHEN UNIVERSITY

Constant Velocity Model

o state

x measurement

* predicted mean estimate

+ corrected mean estimate

bars: variance estimates before and after measurements

Computer Vision II, Summer'14
Slide credit: Kristen Grauman
B. Leibe
Figure from Forsyth & Ponce
54

RWTH AACHEN UNIVERSITY

Constant Velocity Model

o state

x measurement

* predicted mean estimate

+ corrected mean estimate

bars: variance estimates before and after measurements

Computer Vision II, Summer'14
Slide credit: Kristen Grauman
B. Leibe
Figure from Forsyth & Ponce
55

RWTH AACHEN UNIVERSITY

Constant Velocity Model

o state

x measurement

* predicted mean estimate

+ corrected mean estimate

bars: variance estimates before and after measurements

Computer Vision II, Summer'14
Slide credit: Kristen Grauman
B. Leibe
Figure from Forsyth & Ponce
56

RWTH AACHEN UNIVERSITY

Kalman Filter: General Case (>1dim)

- What if state vectors have more than one dimension?

PREDICT

$$x_t^- = D_t x_{t-1}^+$$

$$\Sigma_t^- = D_t \Sigma_{t-1}^+ D_t^T + \Sigma_{d_t}$$

CORRECT

$$K_t = \Sigma_t^- M_t^T (M_t \Sigma_t^- M_t^T + \Sigma_{m_t})^{-1}$$

$$x_t^+ = x_t^- + K_t (y_t - M_t x_t^-)$$

"residual"

$$\Sigma_t^+ = (I - K_t M_t) \Sigma_t^-$$

↻

More weight on residual when measurement error covariance approaches 0.
Less weight on residual as a priori estimate error covariance approaches 0.

for derivations, see F&P Chapter 17.3

Computer Vision II, Summer'14 | Slide credit: Kristen Grauman | B. Leibe | 57

RWTH AACHEN UNIVERSITY

Summary: Kalman Filter


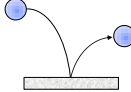
- Pros:**
 - Gaussian densities everywhere
 - Simple updates, compact and efficient
 - Very established method, very well understood
- Cons:**
 - Unimodal distribution, only single hypothesis
 - Restricted class of motions defined by linear model

Computer Vision II, Summer'14 | Slide adapted from Svetlana Lazebnik | B. Leibe | 58

RWTH AACHEN UNIVERSITY

Why Is This A Restriction?

- Many interesting cases don't have linear dynamics
 - E.g. pedestrians walking
- E.g. a ball bouncing

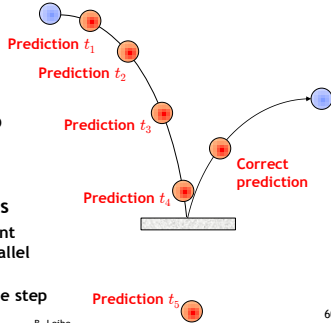



Computer Vision II, Summer'14 | B. Leibe | 59

RWTH AACHEN UNIVERSITY

Ball Example: What Goes Wrong Here?

- Assuming constant acceleration model
- Prediction is too far from true position to compensate...
- Possible solution: Keep multiple models
 - Keep multiple different motion models in parallel
 - I.e. would check for bouncing at each time step




Computer Vision II, Summer'14 | B. Leibe | 60

RWTH AACHEN UNIVERSITY

References and Further Reading

- A very good introduction to tracking with linear dynamic models and Kalman filters can be found in Chapter 17 of
 - D. Forsyth, J. Ponce, *Computer Vision - A Modern Approach*. Prentice Hall, 2003



Computer Vision II, Summer'14 | B. Leibe | 61