

Visual Navigation for Flying Robots

Lecture Notes Summer Term 2012

Lecturer: Dr. Jürgen Sturm

Teaching Assistant: Nikolas Engelhard

<http://vision.in.tum.de/teaching/ss2012/visnav2012>

Acknowledgements

This slide set would not have been possible without the help and support of many other people. In particular, I would like to thank all my great colleagues who made their lecture slides available for everybody on the internet or sent them to me personally.

My thanks go to (in alphabetical order)

- Alexander Kleiner
- Andrew Davison
- Andrew Zisserman
- Antonio Torralba
- Chad Jenkins
- Cyrill Stachniss
- Daniel Cremers
- Georgio Grisetti
- Jan Peters
- Jana Kosecka
- Jörg Müller
- Jürgen Hess
- Kai Arras
- Kai Wurm
- Kurt Konolige
- Li Fei-Fei
- Maxim Likhachev
- Margaritha Chli
- Nicholas Roy
- Paul Newman
- Richard Newcombe
- Richard Szeliski
- Roland Siegwart
- Sebastian Thrun
- Steve Seitz
- Steven Lavalle
- Szymon Rusinkiewicz
- Volker Grabe
- Vijay Kumar
- Wolfram Burgard

Table of Contents

Introduction	1
3D Geometry and Sensors	17
Probabilistic Models and State Estimation	37
Robot Control	55
Visual Motion Estimation	69
Simultaneous Localization and Mapping (SLAM)	85
Bundle Adjustment and Stereo Correspondence	101
Place Recognition, ICP, and Dense Reconstruction	117
Motion Planning	133
Planning under Uncertainty, Exploration and Coordination	149
Experimentation, Evaluation and Benchmarking	165

Visual Navigation for Flying Robots

Welcome

Dr. Jürgen Sturm

Organization

- Tue 10:15-11:45
 - Lectures, discussions
 - Lecturer: Jürgen Sturm
- Thu 14:15-15:45
 - Lab course, homework & programming exercises
 - Teaching assistant: Nikolas Engelhard
- Course website
 - Dates, additional material
 - Exercises, deadlines
 - <http://cvpr.in.tum.de/teaching/ss2012/visnav2012>

Who are we?



- Computer Vision group:
1 Professor, 2 Postdocs, 7 PhD students
- Research topics:
Optical flow and motion estimation, 3D reconstruction, image segmentation, convex optimization
- My research goal:
Apply solutions from computer vision to real-world problems in robotics.

Goal of this Course

- Provide an overview on problems/approaches for autonomous quadrocopters
- Strong focus on vision as the main sensor
- Areas covered: Mobile Robotics and Computer Vision
- Hands-on experience in lab course

Course Material

- Probabilistic Robotics. Sebastian Thrun, Wolfram Burgard and Dieter Fox. MIT Press, 2005.
- Computer Vision: Algorithms and Applications. Richard Szeliski. Springer, 2010.
<http://szeliski.org/Book/>



Lecture Plan

1. Introduction
2. Robots, sensor and motion models
3. State estimation and control
4. Guest talks
5. Feature detection and matching
6. Motion estimation
7. Simultaneous localization and mapping
8. Stereo correspondence
9. 3D reconstruction
10. Navigation and path planning
11. Exploration
12. Evaluation and Benchmarking

Basics on mobile robotics

Camera-based localization and mapping

Advanced topics

Lab Course

- Thu 14:15 – 15:45, given by Nikolas Engelhard
 - Exercises: room 02.09.23
(6x, obliged, homework discussion)
 - Robot lab: room 02.09.34/36
(in weeks without exercises, in case you need help, recommended!)



Exercises Plan

- Exercise sheets contain both theoretical and programming problems
- 3 exercise sheets + 1 mini-project
- Deadline: before lecture (Tue 10:15)
- Hand in by email (visnav2012@cvpr.in.tum.de)

Group Assignment and Schedule

- 3 Ardrone (red/green/blue) + Joystick + 2x Batteries + Charger + PC
- 20 students in the course, 2-3 students per group → 7-8 groups
- Either use lab computers or bring own laptop (recommended)
- Will put up lists for groups and robot schedule in robot lab (room 02.09.36)

VISNAV2012: Team Assignment

Team Name				
Student Name				
Student Name				
Student Name				

Team Name				
Student Name				
Student Name				
Student Name				

VISNAV2012: Robot Schedule

- Each team gets one time slot with programming support
- The robots/PCs are also available during the rest of the week (but without programming support)

	Red	Green	Blue
Thu 2pm – 3pm			
Thu 3pm – 4pm			
Thu 4pm – 5pm			



Safety Warning



- Quadcopters are dangerous objects
- Read the manual carefully before you start
- Always use the protective hull
- If somebody gets injured, report to us so that we can improve safety guidelines
- If something gets damaged, report it to us so that we can fix it
- **NEVER TOUCH THE PROPELLORS**
- **DO NOT TRY TO CATCH THE QUADROPTER WHEN IT FAILS – LET IT FALL/CRASH!**

Agenda for Today

- History of mobile robotics
- Brief intro on quadcopters
- Paradigms in robotics
- Architectures and middleware


General background

- Autonomous, automaton
 - self-willed (Greek, auto+matos)
- Robot
 - Karel Capek in 1923 play R.U.R. (Rossum's Universal Robots)
 - labor (Czech or Polish, robota)
 - workman (Czech or Polish, robotnik)


History

In 1966, Marvin Minsky at MIT asked his undergraduate student Gerald Jay Sussman to “spend the summer linking a camera to a computer and getting the computer to describe what it saw”. We now know that the problem is slightly more difficult than that. (Szeliski 2009, Computer Vision)

Shakey the Robot (1966-1972)



Shakey the Robot (1966-1972)



Stanford Cart (1961-80)




Rhino and Minerva (1998-99)

- Museum tour guide robots
- University of Bonn and CMU
- Deutsches Museum, Smithsonian Museum




Roomba (2002)

- Sensor: one contact sensor
- Control: random movements
- Over 5 million units sold



Neato XV-11 (2010)

- Sensors:
 - 1D range sensor for mapping and localization
 - Improved coverage



Darpa Grand Challenge (2005)



Kiva Robotics (2007)

- Pick, pack and ship automation




Fork Lift Robots (2010)



Quadrocopters (2001-)




Aggressive Maneuvers (2010)



Autonomous Construction (2011)



Mapping with a Quadcopter (2011)



Our Own Recent Work (2011-)

- RGB-D SLAM (Nikolas Engelhard)
- Visual odometry (Frank Steinbrücker)
- Camera-based navigation (Jakob Engel)



Current Trends in Robotics

Robots are entering novel domains

- Industrial automation
- Domestic service robots
- Medical, surgery
- Entertainment, toys
- Autonomous cars
- Aerial monitoring/inspection/construction

Flying Robots

- Recently increased interest in flying robots
 - Shift focus to different problems (control is much more difficult for flying robots, path planning is simpler, ...)
- Especially quadcopters because
 - Can keep position
 - Reliable and compact
 - Low maintenance costs
- Trend towards miniaturization

Application Domains of Flying Robots

- Stunts for action movies, photography, sportscasts
- Search and rescue missions
- Aerial photogrammetry
- Documentation
- Aerial inspection of bridges, buildings, ...
- Construction tasks
- Military
- Today, quadcopters are often still controlled by human pilots

Quadcopter Platforms

- Commercial platforms
 - Ascending Technologies
 - Height Tech
 - Parrot Ardrone ← Used in the lab course
 - ...
- Community/open-source projects
 - Mikrokopter
 - Paparazzi
 - ...


For more, see http://multicopter.org/wiki/Multicopter_Table

Flying Principles


- Fixed-wing airplanes
 - generate lift through forward airspeed and the shape of the wings
 - controlled by flaps
- Helicopters/rotorcrafts
 - main rotor for lift, tail rotor to compensate for torque
 - controlled by adjusting rotor pitch
- Quadcopter/quadrotor
 - four rotors generate lift
 - controlled by changing the speeds of rotation

Helicopter

- Swash plate adjusts pitch of propeller cyclically, controls pitch and roll
- Yaw is controlled by tail rotor




Quadcopter




Keep position:

- Torques of all four rotors sum to zero
- Thrust compensates for earth gravity

Quadrocopter: Basic Motions




Ascend




Descend

Quadrocopter: Basic Motions




Turn Left




Turn Right

Quadrocopter: Basic Motions




Accelerate Forward




Accelerate Backward

Quadrocopter: Basic Motions



Accelerate to the Right



Accelerate to the Left

Autonomous Flight

- Low level control (not covered in this course)
 - Maintain attitude, stabilize
 - Compensate for disturbances
- High level control
 - Compensate for drift
 - Avoid obstacles
 - Localization and Mapping
 - Navigate to point
 - Return to take-off position
 - Person following

Challenges

- Limited payload
 - Limited computational power
 - Limited sensors
- Limited battery life
- Fast dynamics, needs electronic stabilization
- Quadrocopter is always in motion
- Safety considerations

Robot Ethics

- Where does the responsibility for a robot lie?
- How are robots motivated?
- Where are humans in the control loop?
- How might society change with robotics?
- Should robots be programmed to follow a code of ethics, if this is even possible?

Robot Ethics

Three Laws of Robotics (Asimov, 1942):

- A robot may not injure a human being or, through inaction, allow a human being to come to harm.
- A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
- A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Robot Design

Imagine that we want to build a robot that has to perform navigation tasks...

How would you tackle this?

- What hardware would you choose?
- What software architecture would you choose?

Robot Hardware/Components


- Sensors
- Actuators
- Control Unit/Software



Evolution of Paradigms in Robotics

- Classical robotics (mid-70s)
 - Exact models
 - No sensing necessary
- Reactive paradigms (mid-80s)
 - No models
 - Relies heavily on good sensing
- Hybrid approaches (since 90s)
 - Model-based at higher levels
 - Reactive at lower levels

Classical / hierarchical paradigm




- Inspired by methods from Artificial Intelligence (70's)
- Focus on automated reasoning and knowledge representation
- STRIPS (Stanford Research Institute Problem Solver): Perfect world model, closed world assumption
- Shakey: Find boxes and move them to designated positions

Classical paradigm: Stanford Cart

- Take nine images of the environment, identify interesting points, estimate depth
- Integrate information into global world model
- Correlate images with previous image set to estimate robot motion
- On basis of desired motion, estimated motion, and current estimate of environment, determine direction in which to move
- Execute motion

Classical paradigm as horizontal/functional decomposition



Characteristics of hierarchical paradigm

Good old-fashioned Artificial Intelligence (GOFAI):

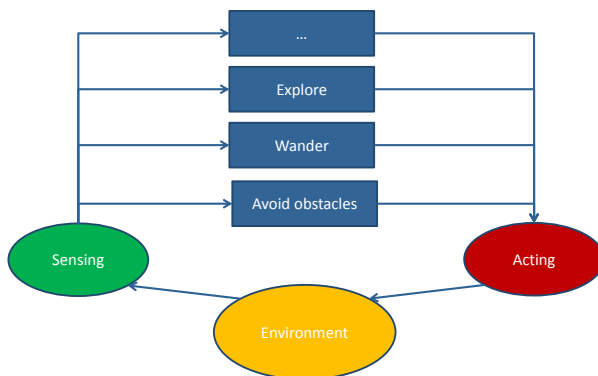
- Symbolic approaches
- Robot perceives the world, plans the next action, acts
- All data is inserted into a single, global world model
- Sequential data processing

Reactive Paradigm



- Sense-act type of organization
- Multiple instances of stimulus-response loops (called behaviors)
- Each behavior uses local sensing to generate the next action
- Combine several behaviors to solve complex tasks
- Run behaviors in parallel, behavior can override (subsume) output of other behaviors

Reactive Paradigm as Vertical Decomposition




Characteristics of Reactive Paradigm

- Situated agent, robot is integral part of the world
- No memory, controlled by what is happening in the world
- Tight coupling between perception and action via behaviors
- Only local, behavior-specific sensing is permitted (ego-centric representation)


Subsumption Architecture

- Introduced by Rodney Brooks in 1986
- Behaviors are networks of sensing and acting modules (augmented finite state machines)
- Modules are grouped into layers of competence
- Layers can subsume lower layers

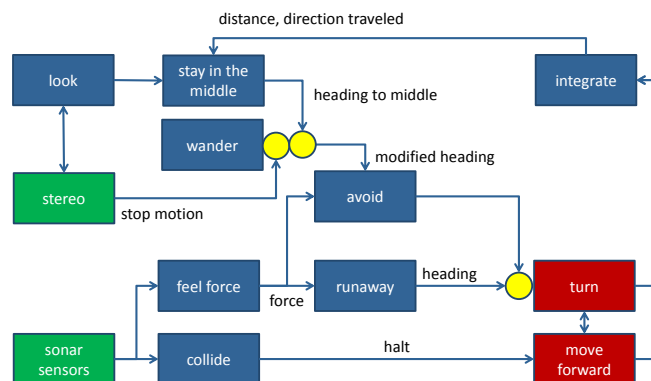
Level 1: Avoid



Level 2: Wander



Level 3: Follow Corridor



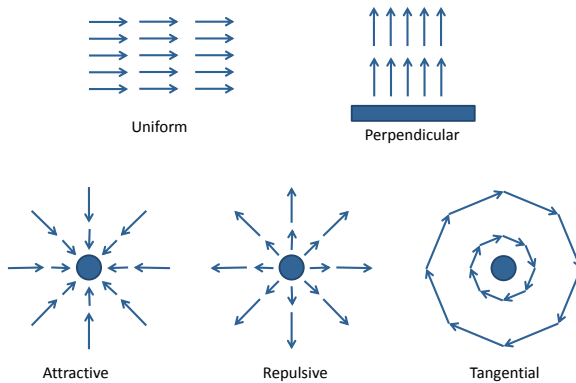
Roomba Robot

- Exercise: Model the behavior of a Roomba robot.


Navigation with Potential Fields

- Treat robot as a particle under the influence of a potential field
- Robot travels along the derivative of the potential
- Field depends on obstacles, desired travel directions and targets
- Resulting field (vector) is given by the summation of primitive fields
- Strength of field may change with distance to obstacle/target

Primitive Potential Fields



Example: reach goal and avoid obstacles




Corridor Following Robot


- Level 1 (collision avoidance)
add repulsive fields for the detected obstacles
- Level 2 (wander)
add a uniform field into a (random) direction
- Level 3 (corridor following)
replaces the wander field by three fields (two perpendicular, one parallel to the walls)

Characteristics of Potential Fields

- Simple method which is often used
- Easy to visualize
- Easy to combine different fields (with parameter tuning)
- But: Suffer from local minima
 - Random motion to escape local minimum
 - Backtracking
 - Increase potential of visited regions
 - High-level planner



Hybrid deliberative/reactive Paradigm



- Combines advantages of previous paradigms
 - World model used in high-level planning
 - Closed-loop, reactive low-level control

Modern Robot Architectures

- Robots became rather complex systems
- Often, a large set of individual capabilities is needed
- Flexible composition of different capabilities for different tasks

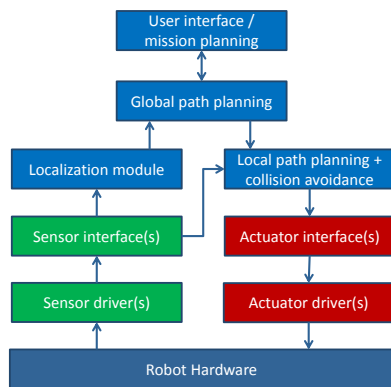
Best Practices for Robot Architectures

- Modular
- Robust
- De-centralized
- Facilitate software re-use
- Hardware and software abstraction
- Provide introspection
- Data logging and playback
- Easy to learn and to extend

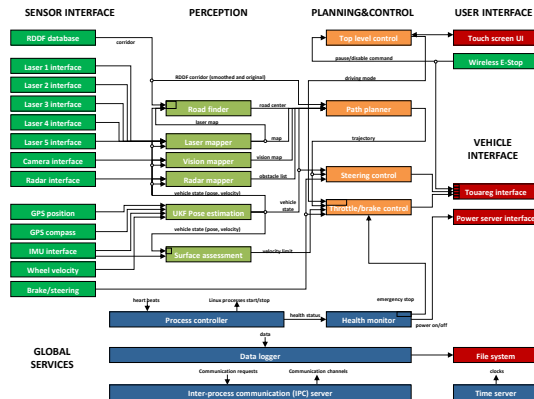
Robotic Middleware

- Provides infrastructure
- Communication between modules
- Data logging facilities
- Tools for visualization
- Several systems available
 - Open-source: ROS (Robot Operating System), Player/Stage, CARMEN, YARP, OROCOS
 - Closed-source: Microsoft Robotics Studio

Example Architecture for Navigation




Stanley's Software Architecture




PR2 Software Architecture

- Two 7-DOF arms, grippers, torso, 2-DOF head
- 7 cameras, 2 laser scanners
- Two 8-core CPUs, 3 network switches
- 73 nodes, 328 message topics, 174 services




Communication Paradigms

- Message-based communication



- Direct (shared) memory access




Forms of Communication

- Push
- Pull
- Publisher/subscriber
- Publish to blackboard
- Remote procedure calls / service calls
- Preemptive tasks / actions


Push

- Broadcast
- One-way communication
- Send as the information is generated by the producer P



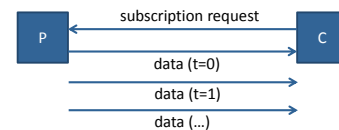
Pull

- Data is delivered upon request by the consumer C (e.g., a map of the building)
- Useful if the consumer C controls the process and the data is not required (or available) at high frequency




Publisher/Subscriber

- The consumer C requests a subscription for the data by the producer P (e.g., a camera or GPS)
- The producer P sends the subscribed data as it is generated to C
- Data generated according to a trigger (e.g., sensor data, computations, other messages, ...)




Publish to Blackboard

- The producer P sends data to the blackboard B (e.g., parameter server)
- A consumer C pull data from the blackboard B
- Only the last instance of data is stored in the blackboard B



Service Calls

- The client C sends a request to the server S
- The server returns the result
- The client waits for the result (synchronous communication)
- Also called: Remote Procedure Call

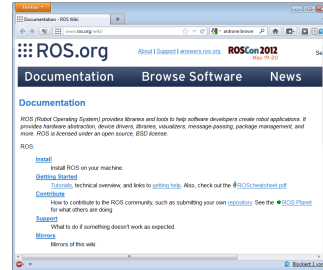


Actions (Preemptive Tasks)

- The client requests the execution of an enduring action (e.g., navigate to a goal location)
- The server executes this action and sends continuously status updates
- Task execution may be canceled from both sides (e.g., timeout, new navigation goal,...)

Robot Operating System (ROS)

- We will use ROS in the lab course
- <http://www.ros.org/>
- Installation instructions, tutorials, docs



Concepts in ROS

- Nodes: programs that communicate with each other
- Messages: data structure (e.g., "Image")
- Topics: typed message channels to which nodes can publish/subscribe (e.g., "/camera1/image_color")
- Parameters: stored in a blackboard



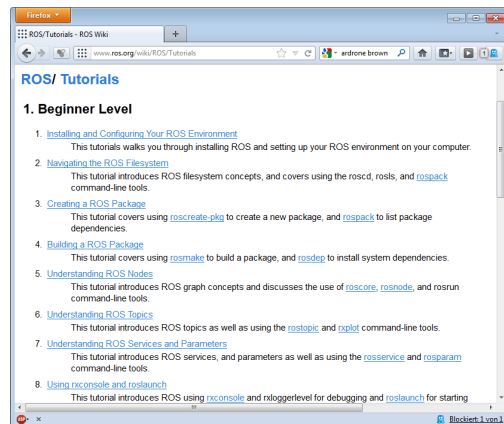
Software Management

- Package: atomic unit of building, contains one or more nodes and/or message definitions
- Stack: atomic unit of releasing, contains several packages with a common theme
- Repository: contains several stacks, typically one repository per institution

Useful Tools

- roscat-pkg
- rosmake
- roscore
- rosnodet list/info
- rostopic list/echo
- rosbag record/play
- rosrund

Tutorials in ROS



Exercise Sheet 1

- On the course website
- Solutions are due in 2 weeks (May 1st)
- Theory part:
Define the motion model of a quadcopter
(will be covered next week)
- Practical part:
Playback a bag file with data from
quadcopter & plot trajectory

Summary

- History of mobile robotics
- Brief intro on quadcopters
- Paradigms in robotics
- Architectures and middleware

Questions?

- See you next week!

Visual Navigation for Flying Robots

3D Geometry and Sensors

Dr. Jürgen Sturm

Organization: Lecture

- Student request to change lecture time to Tuesday afternoon due to time conflicts with other course
- Problem: At least 3 students who are enrolled for this lecture have time Tuesday morning but not on Tuesday afternoon
- Therefore: No change
- Lectures are important, please choose which course to follow
- Note: Still students on the waiting list

Organization: Lab Course

- Robot lab: room 02.09.38 (around the corner)
- Exercises: room 02.09.23 (here)
- You have to sign up for a team before May 1st (team list in student lab)
- After May 1st, remaining places will be given to students on waiting list
- This Thursday: Visual navigation demo at 2pm in the student lab (in conjunction with TUM Girls' Day)

Today's Agenda


- Linear algebra
- 2D and 3D geometry
- Sensors

Vectors

- Vector and its coordinates


$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \in \mathbb{R}^n$$

- Vectors represent points in an n-dimensional space




Vector Operations

- **Scalar multiplication**
- Addition/subtraction
- Length
- Normalized vector
- Dot product
- Cross product




Vector Operations

- Scalar multiplication
- **Addition/subtraction**
- Length
- Normalized vector
- Dot product
- Cross product



Vector Operations


- Scalar multiplication
- Addition/subtraction
- **Length**
- Normalized vector
- Dot product
- Cross product



$$\|x\|_2 = \|x\| = \sqrt{x_1^2 + x_2^2 + \dots}$$

Vector Operations


- Scalar multiplication
- Addition/subtraction
- Length
- **Normalized vector**
- Dot product
- Cross product



$$\hat{x} = \frac{x}{\|x\|}$$

Vector Operations

- Scalar multiplication
- Addition/subtraction
- Length
- Normalized vector
- **Dot product**
- Cross product



$$x \cdot y = \|x\| \|y\| \cos \theta$$


x, y are orthogonal if $x \cdot y = 0$

y is linearly dependent from $\{x_1, x_2, \dots\}$ if

$$y = \sum_i k_i x_i$$

Vector Operations

- Scalar multiplication
- Addition/subtraction
- Length
- Normalized vector
- Dot product
- **Cross product**



$$x \times y = \|x\| \|y\| \sin(\theta) \mathbf{n}$$

Cross Product

- Definition

$$x \times y = \begin{pmatrix} x_2 y_3 - x_3 y_2 \\ x_3 y_1 - x_1 y_3 \\ x_1 y_2 - x_2 y_1 \end{pmatrix}$$

- Matrix notation for the cross product

$$[x]_{\times} = \begin{pmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{pmatrix}$$

- Verify that $x \times y = [x]_{\times} y$

Matrices

- Rectangular array of numbers

$$X = \begin{pmatrix} x_{11} & x_{12} & \dots & x_{1m} \\ x_{21} & x_{22} & \dots & x_{2m} \\ \vdots & & & \\ x_{n1} & x_{n2} & \dots & x_{nm} \end{pmatrix} \in \mathbb{R}^{n \times m}$$

rows ↓ columns ↓

- First index refers to row
- Second index refers to column

Matrices

- Column vectors of a matrix

$$X = \begin{pmatrix} x_{11} & x_{12} & \dots & x_{1m} \\ x_{21} & x_{22} & \dots & x_{2m} \\ \vdots & & & \\ x_{n1} & x_{n2} & \dots & x_{nm} \end{pmatrix}$$

$$= (\mathbf{x}_{*1} \quad \mathbf{x}_{*2} \quad \dots \quad \mathbf{x}_{*m})$$

- Geometric interpretation: for example, column vectors can form basis of a coordinate system

Matrices

- Row vectors of a matrix

$$X = \begin{pmatrix} x_{11} & x_{12} & \dots & x_{1m} \\ x_{21} & x_{22} & \dots & x_{2m} \\ \vdots & & & \\ x_{n1} & x_{n2} & \dots & x_{nm} \end{pmatrix} = \begin{pmatrix} \mathbf{x}_{1*}^\top \\ \mathbf{x}_{2*}^\top \\ \vdots \\ \mathbf{x}_{n*}^\top \end{pmatrix}$$

Matrices

- Square matrix
- Diagonal matrix
- Upper and lower triangular matrix
- Symmetric matrix
- Skew-symmetric matrix
- (Semi-)positive definite matrix
- Invertible matrix
- Orthonormal matrix
- Matrix rank

Matrices

- Square matrix
- Diagonal matrix
- Upper and lower triangular matrix
- Symmetric matrix $X = X^\top$
- Skew-symmetric matrix $X = -X^\top (= \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix})$
- (Semi-)positive definite matrix $\mathbf{a}^\top X \mathbf{a} \geq 0$
- Invertible matrix
- Orthonormal matrix
- Matrix rank

Matrix Operations

- Scalar multiplication
- Addition/subtraction
- Transposition
- Matrix-vector multiplication
- Matrix-matrix multiplication
- Inversion

Matrix Operations

- Scalar multiplication
- Addition/subtraction
- Transposition
- **Matrix-vector multiplication** $X\mathbf{b}$
- Matrix-matrix multiplication
- Inversion

Matrix-Vector Multiplication

- Definition

$$X \cdot \mathbf{b} = \begin{pmatrix} x_{11} & x_{12} & \dots & x_{1m} \\ x_{21} & x_{22} & \dots & x_{2m} \\ \vdots & & & \\ x_{n1} & x_{n2} & \dots & x_{nm} \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{pmatrix} = \sum_{k=1}^m \mathbf{x}_{*k} \cdot b_k$$

↑
column vectors

- Geometric interpretation:
a linear combination of the columns of X scaled by the coefficients of \mathbf{b}

Matrix-Vector Multiplication

$$X \cdot \mathbf{b} = \sum_{k=1}^n \mathbf{x}_{*k} \cdot b_k$$

↑
column vectors

- Geometric interpretation:
A linear combination of the columns of A scaled by the coefficients of \mathbf{b}
→ coordinate transformation

Matrix Operations

- Scalar multiplication
- Addition/subtraction
- Transposition
- Matrix-vector multiplication
- **Matrix-matrix multiplication**
- Inversion

Matrix-Matrix Multiplication

- Operator $\mathbb{R}^{n \times m} \times \mathbb{R}^{m \times p} \rightarrow \mathbb{R}^{n \times p}$
- Definition $C = AB$
 $= A(\mathbf{b}_{*1} \ \mathbf{b}_{*2} \ \dots \ \mathbf{b}_{*p})$
- Interpretation: transformation of coordinate systems
- Can be used to concatenate transforms

Matrix-Matrix Multiplication

- Not commutative (in general)
 $AB \neq BA$
- Associative
 $A(BC) = (AB)C$
- Transpose
 $(AB)^T = B^T A^T$

Matrix Operations

- Scalar multiplication
- Addition/subtraction
- Transposition
- Matrix-vector multiplication
- Matrix-matrix multiplication
- **Inversion**

Matrix Inversion

- If A is a square matrix of full rank, then there is a unique matrix $B = A^{-1}$ such that $AB = I$.
- Different ways to compute, e.g., Gauss-Jordan elimination, LU decomposition, ...
- When A is orthonormal, then

$$A^{-1} = A^T$$

Recap: Linear Algebra

- Vectors
- Matrices
- Operators
- Now let's apply these concepts to 2D+3D geometry

Geometric Primitives in 2D

- 2D point $\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2$
- Augmented vector $\bar{\mathbf{x}} = \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \in \mathbb{R}^3$
- Homogeneous coordinates $\tilde{\mathbf{x}} = \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{w} \end{pmatrix} \in \mathbb{P}^2$

Geometric Primitives in 2D


- Homogeneous vectors that differ only by scale represent the same 2D point
- Convert back to inhomogeneous coordinates by dividing through last element

$$\tilde{\mathbf{x}} = \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{w} \end{pmatrix} = \begin{pmatrix} \tilde{x}/\tilde{w} \\ \tilde{y}/\tilde{w} \\ 1 \end{pmatrix} = \tilde{w} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \tilde{w} \bar{\mathbf{x}}$$

- Points with $\tilde{w} = 0$ are called points at infinity or ideal points

Geometric Primitives in 2D

- 2D line $\tilde{\mathbf{l}} = (a, b, c)^T$
- 2D line equation $\bar{\mathbf{x}} \cdot \tilde{\mathbf{l}} = ax + by + c = 0$




Geometric Primitives in 2D

- Normalized line equation vector

$$\tilde{\mathbf{I}} = (\hat{n}_x, \hat{n}_y, d)^\top = (\hat{\mathbf{n}}, d)^\top \quad \text{with} \quad \|\hat{\mathbf{n}}\| = 1$$


where d is the distance of the line to the origin



Geometric Primitives in 2D

- Polar coordinates of a line: $(\theta, d)^\top$
(e.g., used in Hough transform for finding lines)

$$\hat{\mathbf{n}} = (\cos \theta, \sin \theta)^\top$$



Geometric Primitives in 2D

- Line joining two points $\tilde{\mathbf{I}} = \tilde{\mathbf{x}}_1 \times \tilde{\mathbf{x}}_2$
- Intersection point of two lines $\tilde{\mathbf{x}} = \tilde{\mathbf{I}}_1 \times \tilde{\mathbf{I}}_2$


Geometric Primitives in 3D

- 3D point (same as before) $\mathbf{x} = \begin{pmatrix} x \\ y \\ z \end{pmatrix} \in \mathbb{R}^3$
- Augmented vector $\bar{\mathbf{x}} = \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \in \mathbb{R}^4$
- Homogeneous coordinates $\tilde{\mathbf{x}} = \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \\ \tilde{w} \end{pmatrix} \in \mathbb{P}^3$

Geometric Primitives in 3D

- 3D plane $\tilde{\mathbf{m}} = (a, b, c, d)^\top$
- 3D plane equation $\tilde{\mathbf{x}} \cdot \tilde{\mathbf{m}} = ax + by + cz + d = 0$


- Normalized plane with unit normal vector $\mathbf{m} = (\hat{n}_x, \hat{n}_y, \hat{n}_z, d)^\top = (\hat{\mathbf{n}}, d)$ ($\|\hat{\mathbf{n}}\| = 1$) and distance d




Geometric Primitives in 3D

- 3D line $\mathbf{r} = (1 - \lambda)\mathbf{p} + \lambda\mathbf{q}$ through points \mathbf{p}, \mathbf{q}

- Infinite line: $\lambda \in \mathbb{R}$
- Line segment joining \mathbf{p}, \mathbf{q} : $0 \leq \lambda \leq 1$



2D Planar Transformations



2D Transformations

- Translation $\mathbf{x}' = \mathbf{x} + \mathbf{t}$

$$\mathbf{x}' = \underbrace{\begin{pmatrix} \mathbf{I} & \mathbf{t} \end{pmatrix}}_{2 \times 3} \bar{\mathbf{x}}$$

$$\bar{\mathbf{x}}' = \underbrace{\begin{pmatrix} \mathbf{I} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{pmatrix}}_{3 \times 3} \bar{\mathbf{x}}$$

where \mathbf{I} is the identity matrix (2x2)
and $\mathbf{0}$ is the zero vector

2D Transformations

- Rotation + translation (2D rigid body motion, or 2D Euclidean transformation)

$$\mathbf{x}' = \mathbf{R}\mathbf{x} + \mathbf{t} \quad \text{or} \quad \bar{\mathbf{x}}' = \begin{pmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{pmatrix} \bar{\mathbf{x}}$$

where $\mathbf{R} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$

is an orthonormal rotation matrix, i.e., $\mathbf{R}\mathbf{R}^\top = \mathbf{I}$

- Distances (and angles) are preserved

2D Transformations

- Scaled rotation/similarity transform

$$\mathbf{x}' = s\mathbf{R}\mathbf{x} + \mathbf{t} \quad \text{or} \quad \bar{\mathbf{x}}' = \begin{pmatrix} s\mathbf{R} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{pmatrix} \bar{\mathbf{x}}$$

- Preserves angles between lines

2D Transformations

- Affine transform

$$\bar{\mathbf{x}}' = A\bar{\mathbf{x}} = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{pmatrix} \bar{\mathbf{x}}$$

- Parallel lines remain parallel




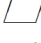

2D Transformations

- Projective/perspective transform

$$\tilde{\mathbf{x}}' = \tilde{H} = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \tilde{\mathbf{x}}$$

- Note that \tilde{H} is homogeneous (only defined up to scale)
- Resulting coordinates are homogeneous
- Parallel lines remain parallel

2D Transformations

Transformation	Matrix	# DoF	Preserves	Icon
translation	$\begin{bmatrix} I & t \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	2	orientation	
rigid (Euclidean)	$\begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	3	lengths	
similarity	$\begin{bmatrix} sR & t \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	4	angles	
affine	$\begin{bmatrix} A \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	6	parallelism	
projective	$\begin{bmatrix} \tilde{H} \\ 0 & 1 \end{bmatrix}_{3 \times 3}$	8	straight lines	

3D Transformations

- Translation






$$\bar{x}' = \underbrace{\begin{pmatrix} I & t \\ 0^T & 1 \end{pmatrix}}_{4 \times 4} \bar{x}$$

- Euclidean transform (translation + rotation), (also called the Special Euclidean group SE(3))

$$\bar{x}' = \begin{pmatrix} R & t \\ 0^T & 1 \end{pmatrix} \bar{x}$$

- Scaled rotation, affine transform, projective transform...

3D Transformations

Transformation	Matrix	# DoF	Preserves	Icon
translation	$\begin{bmatrix} I & t \\ 0 & 1 \end{bmatrix}_{3 \times 4}$	3	orientation	
rigid (Euclidean)	$\begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix}_{3 \times 4}$	6	lengths	
similarity	$\begin{bmatrix} sR & t \\ 0 & 1 \end{bmatrix}_{3 \times 4}$	7	angles	
affine	$\begin{bmatrix} A \\ 0 & 1 \end{bmatrix}_{3 \times 4}$	12	parallelism	
projective	$\begin{bmatrix} \tilde{H} \\ 0 & 1 \end{bmatrix}_{4 \times 4}$	15	straight lines	

3D Rotations

- Rotation matrix (also called the special orientation group SO(3))

- Euler angles
- Axis/angle
- Unit quaternion

Rotation Matrix


- Orthonormal 3x3 matrix

$$R = \begin{pmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{pmatrix}$$

- Column vectors correspond to coordinate axes
- Special orientation group $R \in SO(3)$
- Main disadvantage: Over-parameterized (9 parameters instead of 3)

Euler Angles

- Product of 3 consecutive rotations
- Roll-pitch-yaw convention is very common in aerial navigation (DIN 9300)



Euler Angles

- Yaw Ψ , Pitch Θ , Roll Φ to rotation matrix

$$R = R_Z(\Psi)R_Y(\Theta)R_X(\Phi)$$

$$= \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \Phi & \sin \Phi \\ 0 & -\sin \Phi & \cos \Phi \end{pmatrix} \begin{pmatrix} \cos \Theta & 0 & -\sin \Theta \\ 0 & 1 & 0 \\ \sin \Theta & 0 & \cos \Theta \end{pmatrix} \begin{pmatrix} \cos \Psi & \sin \Psi & 0 \\ -\sin \Psi & \cos \Psi & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$= \begin{pmatrix} \cos \Theta \cos \Psi & \cos \Theta \sin \Psi & -\sin \Theta \\ \sin \Phi \sin \Theta \cos \Psi - \cos \Phi \sin \Psi & \sin \Phi \sin \Theta \sin \Psi + \cos \Phi \cos \Psi & \sin \Phi \cos \Theta \\ \cos \Phi \sin \Theta \cos \Psi + \sin \Phi \sin \Psi & \cos \Phi \sin \Theta \sin \Psi - \sin \Phi \cos \Psi & \cos \Phi \cos \Theta \end{pmatrix}$$

- Rotation matrix to Yaw-Pitch-Roll

$$\phi = \text{Atan2} \left(-r_{31}, \sqrt{r_{11}^2 + r_{21}^2} \right)$$

$$\psi = -\text{Atan2} \left(\frac{r_{21}}{\cos(\phi)}, \frac{r_{11}}{\cos(\phi)} \right)$$


$$\theta = \text{Atan2} \left(\frac{r_{32}}{\cos(\phi)}, \frac{r_{33}}{\cos(\phi)} \right)$$

Euler Angles

- Advantage:
 - Minimal representation (3 parameters)
 - Easy interpretation
- Disadvantages:
 - Many "alternative" Euler representations exist (XYZ, ZXZ, ZYX, ...)
 - Singularities (gimbal lock)


Gimbal Lock

- When the axes align, one degree-of-freedom (DOF) is lost...



Axis/Angle

- Represent rotation by
 - rotation axis \hat{n} and
 - rotation angle θ
- 4 parameters (\hat{n}, θ)
- 3 parameters $\omega = \theta \hat{n}$
 - length is rotation angle
 - also called the angular velocity
 - minimal but not unique (why?)



Derivation of Angular Velocities

- Assume we have a rotational motion in $SO(3)$

$$R(t) \in SO(3) \quad t \in \mathbb{R}$$
- As this rotations are orthonormal matrices, we have

$$R(t)R^T(t) = I$$
- Now take the derivative on both sides (w.r.t. t)

$$\dot{R}(t)R^T(t) + R(t)\dot{R}^T(t) = 0$$

$$\dot{R}(t)R^T(t) = -(\dot{R}(t)R^T(t))^T$$
- Thus, $\dot{R}(t)R^T(t)$ must be skew-symmetric, i.e.,

$$[\omega(t)]_{\times} = \dot{R}(t)R^T(t)$$

Derivation of Angular Velocities

→ Linear ordinary differential equation (ODE)

$$\dot{R}(t) = [\omega]_{\times} R(t) = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix} R(t)$$

- Solution of this ODE

$$R(t) = \exp([\omega]_{\times})R(0)$$
- Conversions

$$R = \exp([\omega]_{\times}) \quad [\omega]_{\times} = \log R$$

Derivation of Angular Velocities

→ Linear ordinary differential equation (ODE)

$$\dot{R}(t) = [\boldsymbol{\omega}]_{\times} R(t) = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix} R(t)$$

- The space of all skew-symmetric matrices is called the *tangent space*

$$\text{so}(3) = \{[\boldsymbol{\omega}]_{\times} \in \mathbb{R}^{3 \times 3} \mid \boldsymbol{\omega} \in \mathbb{R}^3\}$$

- Space of all rotations in 3D (Special orientation group)

$$\text{SO}(3) = \{R \in \mathbb{R}^{3 \times 3} \mid R^{\top} R = I, \det R = 1\}$$

Conversion

- Rodriguez' formula

$$R(\hat{\mathbf{n}}, \theta) = I + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2$$

- Inverse

$$\theta = \cos^{-1} \left(\frac{\text{trace}(R) - 1}{2} \right), \hat{\mathbf{n}} = \frac{1}{2 \sin \theta} \begin{pmatrix} r_{32} - r_{23} \\ r_{13} - r_{31} \\ r_{21} - r_{12} \end{pmatrix}$$

see: An Invitation to 3D Vision, Y. Ma, S. Soatto, J. Kosecka, S. Sastry, Chapter 2 (available online)

Exponential Twist

- The exponential map can be generalized to Euclidean transformations (incl. translations)

- Tangent space $\text{se}(3) = \text{so}(3) \times \mathbb{R}^3$

- (Special) Euclidean group $\text{SE}(3) = \text{SO}(3) \times \mathbb{R}^3$ (group of all Euclidean transforms)

- Rigid body velocity

$$\xi = \left(\underbrace{\omega_x, \omega_y, \omega_z}_{\text{angular vel.}}, \underbrace{v_x, v_y, v_z}_{\text{linear vel.}} \right) \in \mathbb{R}^6$$

Exponential Twist

- Convert to homogeneous coordinates

$$\hat{\xi} = \begin{pmatrix} 0 & -\omega_z & \omega_y & v_x \\ \omega_z & 0 & -\omega_x & v_y \\ -\omega_y & \omega_x & 0 & v_z \\ 0 & 0 & 0 & 0 \end{pmatrix} \in \text{se}(3)$$

- Exponential map between $\text{se}(3)$ and $\text{SE}(3)$

$$M = \exp \hat{\xi} \quad \hat{\xi} = \log M$$

- There are also direct formulas (similar to Rodriguez)


Unit Quaternions

- Quaternion $\mathbf{q} = (q_x, q_y, q_z, q_w)^{\top} \in \mathbb{R}^4$

- Unit quaternions have $\|\mathbf{q}\| = 1$

- Opposite sign quaternions represent the same rotation $\mathbf{q} = -\mathbf{q}$

- Otherwise unique



Unit Quaternions

- Advantage: multiplication and inversion operations are really fast

- Quaternion-Quaternion Multiplication

$$\begin{aligned} \mathbf{q}_0 \mathbf{q}_1 &= (\mathbf{v}_0, w_0)(\mathbf{v}_1, w_1) \\ &= (\mathbf{v}_0 \times \mathbf{v}_1 + w_0 \mathbf{v}_1 + w_1 \mathbf{v}_0, w_0 w_1 - \mathbf{v}_0 \mathbf{v}_1) \end{aligned}$$

- Inverse (flip sign of \mathbf{v} or w)

$$\begin{aligned} \mathbf{q}_0 / \mathbf{q}_1 &= (\mathbf{v}_0, w_0) / (\mathbf{v}_1, w_1) \\ &= (\mathbf{v}_0, w_0)(\mathbf{v}_1, -w_1) \\ &= (\mathbf{v}_0 \times \mathbf{v}_1 + w_0 \mathbf{v}_1 - w_1 \mathbf{v}_0, -w_0 w_1 - \mathbf{v}_0 \mathbf{v}_1) \end{aligned}$$

Unit Quaternions

- Quaternion-Vector multiplication (rotate point \mathbf{p} with rotation \mathbf{q})

$$\mathbf{p}' = \mathbf{v}\bar{\mathbf{p}}/\mathbf{q}$$

with $\bar{\mathbf{p}} = (x, y, z, 0)^\top$

- Relation to Axis/Angle representation

$$\mathbf{q} = (\mathbf{v}, w) = \left(\sin \frac{\theta}{2} \hat{\mathbf{n}}, \cos \frac{\theta}{2} \right)$$

Spherical Linear Interpolation (SLERP)

- Useful for interpolating between two rotations


procedure $slerp(\mathbf{q}_0, \mathbf{q}_1, \alpha)$:

1. $\mathbf{q}_r = \mathbf{q}_1/\mathbf{q}_0 = (v_r, w_r)$
2. if $w_r < 0$ then $\mathbf{q}_r \leftarrow -\mathbf{q}_r$
3. $\theta_r = 2 \tan^{-1}(\|\mathbf{v}_r\|/w_r)$
4. $\hat{\mathbf{n}}_r = \mathcal{N}(\mathbf{v}_r) = \mathbf{v}_r/\|\mathbf{v}_r\|$
5. $\theta_\alpha = \alpha \theta_r$
6. $\mathbf{q}_\alpha = (\sin \frac{\theta_\alpha}{2} \hat{\mathbf{n}}_r, \cos \frac{\theta_\alpha}{2})$


3D to 2D Projections

- Orthographic projections
- Perspective projections

3D to 2D Perspective Projection



3D to 2D Perspective Projection



3D to 2D Perspective Projection

- 3D point \mathbf{p} (in the camera frame)
- 2D point \mathbf{x} (on the image plane)
- Pin-hole camera model


$$\tilde{\mathbf{x}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \bar{\mathbf{p}}$$

- Remember, $\tilde{\mathbf{x}}$ is homogeneous, need to normalize

$$\tilde{\mathbf{x}} = \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \end{pmatrix} \Rightarrow \mathbf{x} = \begin{pmatrix} \tilde{x}/\tilde{z} \\ \tilde{y}/\tilde{z} \end{pmatrix}$$

Camera Intrinsics

- So far, 2D point is given in meters on image plane
- But: we want 2D point be measured in pixels (as the sensor does)



Camera Intrinsics

- Need to apply some scaling/offset

$$\tilde{\mathbf{x}} = \underbrace{\begin{pmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix}}_{\text{intrinsics } K} \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}}_{\text{projection}} \tilde{\mathbf{p}}$$

- Focal length f_x, f_y
- Camera center c_x, c_y
- Skew s

Camera Extrinsics

- Assume $\tilde{\mathbf{p}}_w$ is given in world coordinates
- Transform from world to camera (also called the camera extrinsics)

$$\tilde{\mathbf{p}} = \begin{pmatrix} R & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{pmatrix} \tilde{\mathbf{p}}_w$$

- Full camera matrix

$$\tilde{\mathbf{x}} = \begin{pmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} R & \mathbf{t} \end{pmatrix} \tilde{\mathbf{p}}_w$$

Recap: 2D/3D Geometry


- points, lines, planes
- 2D and 3D transformations
- Different representations for 3D orientations
 - Choice depends on application
 - Which representations do you remember?
- 3D to 2D perspective projections
- You **really** have to know 2D/3D transformations by heart (read Szeliski, Chapter 2)

C++ Libraries for Lin. Alg./Geometry


- Many C++ libraries exist for linear algebra and 3D geometry
- Typically conversion necessary
- Examples:
 - C arrays, std::vector (no linear alg. functions)
 - gsl (gnu scientific library, many functions, plain C)
 - boost::array (used by ROS messages)
 - Bullet library (3D geometry, used by ROS tf)
 - Eigen (both linear algebra and geometry, my recommendation)

Example: Transform Trees in ROS

- TF package represents 3D transforms between rigid bodies in the scene as a tree



Example: Video from PR2



Sensors

Classification of Sensors

- **What:**
 - **Proprioceptive sensors**
 - Measure values internally to the system (robot)
 - Examples: battery status, motor speed, accelerations, ...
 - **Exteroceptive sensors**
 - Provide information about the environment
 - Examples: compass, distance to objects, ...
- **How:**
 - **Passive sensors**
 - Measure energy coming from the environment
 - **Active sensors**
 - Emit their proper energy and measure the reaction
 - Better performance, but influence on environment

Classification of Sensors

- **Tactile sensors**
Contact switches, bumpers, proximity sensors, pressure
- **Wheel/motor sensors**
Potentiometers, brush/optical/magnetic/inductive/capacitive encoders, current sensors
- **Heading sensors**
Compass, infrared, inclinometers, gyroscopes, accelerometers
- **Ground-based beacons**
GPS, optical or RF beacons, reflective beacons
- **Active ranging**
Ultrasonic sensor, laser rangefinder, optical triangulation, structured light
- **Motion/speed sensors**
Doppler radar, Doppler sound
- **Vision-based sensors**
CCD/CMOS cameras, visual servoing packages, object tracking packages

Example: Ardrone Sensors

- **Tactile sensors**
Contact switches, bumpers, proximity sensors, pressure
- **Wheel/motor sensors**
Potentiometers, brush/optical/magnetic/inductive/capacitive encoders, **current sensors**
- **Heading sensors**
Compass, infrared, inclinometers, **gyroscopes, accelerometers**
- **Ground-based beacons**
GPS, **optical** or **RF beacons**, reflective beacons
- **Active ranging**
Ultrasonic sensor, laser rangefinder, optical triangulation, structured light
- **Motion/speed sensors**
Doppler radar, Doppler sound
- **Vision-based sensors**
CCD/CMOS cameras, **visual servoing packages**, object tracking packages

Characterization of Sensor Performance


- **Bandwidth or Frequency**
- **Delay**
- **Sensitivity**
- **Cross-sensitivity (cross-talk)**
- **Error (accuracy)**
 - Deterministic errors (modeling/calibration possible)
 - Random errors
- **Weight, power consumption, ...**

Sensors

- Motor/wheel encoders
- Compass
- Gyroscope
- Accelerometers
- GPS
- Range sensors
- Cameras


Motor/wheel encoders

- Device for measuring angular motion
- Often used in (wheeled) robots
- Output: position, speed (possibly integrate speed to get odometry)




Motor/wheel encoders

- Working principle:
 - Regular: counts the number of transitions but cannot tell direction
 - Quadrature: uses two sensors in quadrature phase-shift, ordering of rising edge tells direction
 - Sometimes: Reference pulse (or zero switch)




Magnetic Compass

- Measures earth's magnetic field
- Inclination angle approx. 60deg (Germany)
- Does not work indoor/affected by metal
- Alternative: gyro compass (spinning wheel, aligns with earth's rotational poles, for ships)




Magnetic Declination

- Angle between magnetic north and true north
- Varies over time
- Good news ;-): by 2050, magnetic declination in central Europe will be zero




Magnetic Compass

- Sensing principle: Hall sensor
- Construction: 3 orthogonal sensors




Mechanical Gyroscope

- Measures orientation (standard gyro) or angular velocity (rate gyro, needs integration for angle)
- Spinning wheel mounted in a gimbal device (can move freely in 3 dimensions)
- Wheel keeps orientation due to angular momentum (standard gyro)



Modern Gyroscopes

- Vibrating structure gyroscope (MEMS)
 - Based on Coriolis effect
 - “Vibration keeps its direction under rotation”
 - Implementations: Tuning fork, vibrating wheels, ...
- Ring laser / fibre optic gyro
 - Interference between counter-propagating beams in response to rotation



Accelerometer

- Measures all external forces acting upon them (including gravity)
- Acts like a spring-damper system
- To obtain inertial acceleration (due to motion alone), gravity must be subtracted





Fig. 1. Schematic structure of an accelerometer

MEMS Accelerometers

- Micro Electro-Mechanical Systems (MEMS)
- Spring-like structure with a proof mass
- Damping results from residual gas
- Implementations: capacitive, piezoelectric, ...




Inertial Measurement Unit

- 3-axes MEMS gyroscope
 - Provides angular velocity
 - Integrate for angular position
 - Problem: Drifts slowly over time (e.g., 1deg/hour), called the bias
- 3-axes MEMS accelerometer
 - Provides accelerations (including gravity)
- Can we use these sensors to estimate our position?


Inertial Measurement Unit

- IMU: Device that uses gyroscopes and accelerometers to estimate (relative) position, orientation, velocity and accelerations
- Integrate angular velocities to obtain absolute orientation
- Subtract gravity from acceleration
- Integrate acceleration to linear velocities
- Integrate linear velocities to position
- Note: All IMUs are subject to drift (position is integrated twice!), needs external reference

Example: AscTec Autopilot Board




GPS



GPS

- 24+ satellites, 12 hour orbit, 20.190 km height
- 6 orbital planes, 4+ satellites per orbit, 60deg distance



- Satellite transmits orbital location + time
- 50bits/s, msg has 1500 bits → 12.5 minutes

GPS

- Position from pseudorange
 - Requires measurements of 4 different satellites
 - Low accuracy (3-15m) but absolute
- Position from pseudorange + phase shift
 - Very precise (1mm) but highly ambiguous
 - Requires reference receiver (RTK/dGPS) to remove ambiguities

Range Sensors

- Sonar
- Laser range finder
- Time of flight camera
- Structured light (will be covered later)




Range Sensors

- Emit signal to determine distance along a ray
- Make use of propagation speed of ultrasound/light
- Traveled distance is given by $d = c \cdot t$
- Sound speed: 340m/s
- Light speed: 300.000km/s


Ultrasonic Range Sensors

- Range between 12cm and 5m
- Opening angle around 20 to 40 degrees
- Soft surfaces absorb sound
- Reflections → ghosts
- Lightweight and cheap




Laser Scanner

- Measures phase shift
- Pro: High precision, wide field of view, safety approved for collision detection
- Con: Relatively expensive + heavy




Laser Scanner

- 2D scanners




- 3D scanners




Camera

- Let's design a camera
 - Idea 1: put a piece of film in front of an object
 - Do we get a reasonable image?




Camera

- Add a barrier to block off most of the rays
 - This reduces blurring
 - The opening known as the **aperture**
 - How does this transform the image?




Camera Lens

- A lens focuses light onto the film
 - Rays passing through the optical center are not deviated




Camera Lens

- A lens focuses light onto the film
 - Rays passing through the center are not deviated
 - All rays parallel to the **Optical Axis** converge at the **Focal Point**



Camera Lens

- There is a specific distance at which objects are “in focus”
- Other points project to a “blur circle” in the image



Lens Distortions

- Radial distortion of the image
 - Caused by imperfect lenses
 - Deviations are most noticeable for rays that pass through the edge of the lens



Lens Distortions



- Radial distortion of the image
 - Caused by imperfect lenses
 - Deviations are most noticeable for rays that pass through the edge of the lens
- Typically compensated with a low-order polynomial

$$\hat{x}_c = x_c(1 + \kappa_1 r_c^2 + \kappa_2 r_c^4)$$

$$\hat{y}_c = y_c(1 + \kappa_1 r_c^2 + \kappa_2 r_c^4)$$



Digital Cameras

- Vignetting
- De-bayering
- Rolling shutter and motion blur
- Compression (JPG)
- Noise



Dead Reckoning and Odometry

- Estimating the position \mathbf{x}_t based on the issued controls (or IMU) readings \mathbf{u}_t
- Integrated over time $\mathbf{x}_t = f(\mathbf{x}_{t-1}, \mathbf{u}_t)$



Exercise Sheet 1

- Odometry sensor on Ardrone is an integrated package
- Sensors
 - Down-looking camera to estimate motion
 - Ultrasonic sensor to get height
 - 3-axes gyroscopes
 - 3-axes accelerometer
- IMU readings \mathbf{u}_t
 - Horizontal speed (v_x/v_y)
 - Height (z)
 - Roll, Pitch, Yaw
- Integrate these values to get robot pose $\mathbf{x}_t = f(\mathbf{x}_{t-1}, \mathbf{u}_t)$
 - Position ($x/y/z$)
 - Orientation (e.g., $r/p/y$)

Summary

- Linear Algebra
- 2D/3D Geometry
- Sensors

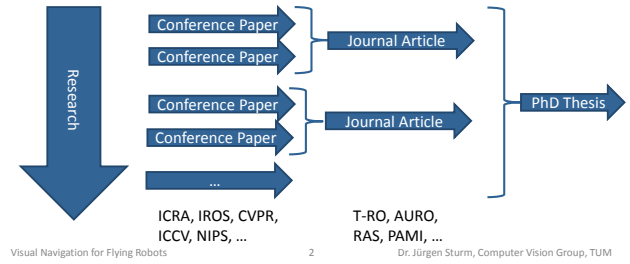
Visual Navigation for Flying Robots

Probabilistic Models and State Estimation

Dr. Jürgen Sturm

Organization

- Next week: Three scientific guest talks
- Recent research results from our group (2011/12)




Guest Talks

- An Evaluation of the RGB-D SLAM System (F. Endres, J. Hess, N. Engelhard, J. Sturm, D. Cremers, W. Burgard), In Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA), 2012.
- Real-Time Visual Odometry from Dense RGB-D Images (F. Steinbruecker, J. Sturm, D. Cremers), In Workshop on Live Dense Reconstruction with Moving Cameras at the Intl. Conf. on Computer Vision (ICCV), 2011.
- Camera-Based Navigation of a Low-Cost Quadcopter (J. Engel, J. Sturm, D. Cremers), Submitted to International Conference on Robotics and Systems (IROS), under review.

Visual Navigation for Flying Robots 3 Dr. Jürgen Sturm, Computer Vision Group, TUM

Perception


- Perception and models are strongly linked




Visual Navigation for Flying Robots 4 Dr. Jürgen Sturm, Computer Vision Group, TUM

Perception

- Perception and models are strongly linked
- Example: Human Perception




Visual Navigation for Flying Robots 5 Dr. Jürgen Sturm, Computer Vision Group, TUM



more on <http://michaelbach.de/ot/index.html>
Visual Navigation for Flying Robots 6 Dr. Jürgen Sturm, Computer Vision Group, TUM

Models in Human Perception

- Count the black dots



Visual Navigation for Flying Robots

7

Dr. Jürgen Sturm, Computer Vision Group, TUM

State Estimation

- Cannot observe world state directly
- Need to estimate the world state
- Robot maintains belief about world state
- Update belief according to observations and actions using models
- Sensor observations + sensor model
- Executed actions + action/motion model

Visual Navigation for Flying Robots

8

Dr. Jürgen Sturm, Computer Vision Group, TUM

State Estimation

What parts of the world state are (most) relevant for a flying robot?

Visual Navigation for Flying Robots

9

Dr. Jürgen Sturm, Computer Vision Group, TUM

State Estimation

What parts of the world state are (most) relevant for a flying robot?

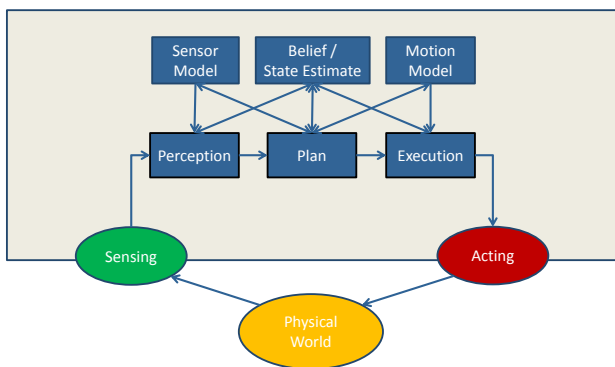
- Position
- Velocity
- Obstacles
- Map
- Positions and intentions of other robots/humans
- ...

Visual Navigation for Flying Robots

10

Dr. Jürgen Sturm, Computer Vision Group, TUM

Models and State Estimation



Visual Navigation for Flying Robots

11

Dr. Jürgen Sturm, Computer Vision Group, TUM

(Deterministic) Sensor Model

- Robot perceives the environment through its sensors

$$z = h(x)$$

↑
sensor
reading
↑
world
state

observation
function

- Goal: Infer the state of the world from sensor readings

$$x = h^{-1}(z)$$


Visual Navigation for Flying Robots

12

Dr. Jürgen Sturm, Computer Vision Group, TUM

(Deterministic) Motion Model

- Robot executes an action u (e.g., move forward at 1m/s)
- Update belief state according to motion model



Visual Navigation for Flying Robots

13

Dr. Jürgen Sturm, Computer Vision Group, TUM

Probabilistic Robotics

- Sensor observations are noisy, partial, potentially missing (why?)
- All models are partially wrong and incomplete (why?)
- Usually we have prior knowledge (why?)

Visual Navigation for Flying Robots

14

Dr. Jürgen Sturm, Computer Vision Group, TUM

Probabilistic Robotics

- Probabilistic sensor and motion models
- Integrate information from multiple sensors (multi-modal)

$$p(z | x) \quad p(x' | x, u)$$

$$p(x | z_{\text{vision}}, z_{\text{ultrasound}}, z_{\text{IMU}})$$

- Integrate information over time (filtering)

$$p(x | z_1, z_2, \dots, z_t)$$

$$p(x | u_1, z_1, \dots, u_t, z_t)$$

Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

Agenda for Today

- Motivation ✓
- Bayesian Probability Theory
- Bayes Filter
- Normal Distribution
- Kalman Filter

Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

The Axioms of Probability Theory

Notation: $P(A)$ refers to the probability that proposition A holds

1. $0 \leq P(A) \leq 1$
2. $P(\Omega) = 1 \quad P(\emptyset) = 0$
3. $P(A \cup B) = P(A) + P(B) - P(A \cap B)$


Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

A Closer Look at Axiom 3

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$



Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM

Discrete Random Variables

- X denotes a **random variable**
- X can take on a countable number of values in $\{x_1, x_2, \dots, x_n\}$
- $P(X = x_i)$ is the **probability** that the random variable X takes on value x_i
- $P(\cdot)$ is called the **probability mass function**
- **Example:** $P(\text{Room}) = \langle 0.7, 0.2, 0.08, 0.02 \rangle$
 $\text{Room} \in \{\text{office, corridor, lab, kitchen}\}$

Visual Navigation for Flying Robots

19



Dr. Jürgen Sturm, Computer Vision Group, TUM

Continuous Random Variables

- X takes on continuous values
- $p(X = x)$ or $p(x)$ is called the **probability density function (PDF)**

$$P(x \in [a, b]) = \int_a^b p(x) dx$$

- **Example**



Visual Navigation for Flying Robots

20

Dr. Jürgen Sturm, Computer Vision Group, TUM

Proper Distributions Sum To One

- Discrete case $\sum_x P(x) = 1$

- Continuous case $\int p(x) dx = 1$

Visual Navigation for Flying Robots

21

Dr. Jürgen Sturm, Computer Vision Group, TUM

Joint and Conditional Probabilities

- $P(X = x \text{ and } Y = y) = P(x, y)$

- If X and Y are **independent** then

$$P(x, y) = P(x)P(y)$$

- $P(x | y)$ is the probability of **x given y**

$$P(x | y)P(y) = P(x, y)$$

- If X and Y are independent then

$$P(x | y) = P(x)$$

Visual Navigation for Flying Robots

22

Dr. Jürgen Sturm, Computer Vision Group, TUM

Conditional Independence

- Definition of conditional independence

$$P(x, y | z) = P(x | z)P(y | z)$$

- Equivalent to $P(x | z) = P(x | y, z)$
 $P(y | z) = P(y | x, z)$

- **Note:** this does not necessarily mean that

$$P(x, y) = P(x)P(y)$$

Visual Navigation for Flying Robots

23

Dr. Jürgen Sturm, Computer Vision Group, TUM

Marginalization

- Discrete case $P(x) = \sum_y P(x, y)$

- Continuous case $p(x) = \int p(x, y) dy$

Visual Navigation for Flying Robots

24

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Marginalization

	x_1	x_2	x_3	x_4	$P_Y(Y) \downarrow$
Y_1	$\frac{1}{8}$	$\frac{1}{16}$	$\frac{1}{32}$	$\frac{1}{32}$	$\frac{1}{4}$
Y_2	$\frac{1}{16}$	$\frac{1}{8}$	$\frac{1}{32}$	$\frac{1}{32}$	$\frac{1}{4}$
Y_3	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{4}$
Y_4	$\frac{1}{4}$	0	0	0	$\frac{1}{4}$
$P_X(X) \rightarrow$	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{8}$	1

Visual Navigation for Flying Robots

25

Dr. Jürgen Sturm, Computer Vision Group, TUM

Law of Total Probability

▪ Discrete case
$$P(x) = \sum_y P(x, y) = \sum_y P(x | y)P(y)$$

▪ Continuous case
$$p(x) = \int p(x, y)dy = \int p(x | y)p(y)dy$$

Visual Navigation for Flying Robots

26

Dr. Jürgen Sturm, Computer Vision Group, TUM

Expected Value of a Random Variable

- Discrete case
$$E[X] = \sum_i x_i P(x_i)$$
- Continuous case
$$E[X] = \int x P(X = x) dx$$

- The expected value is the weighted average of all values a random variable can take on.
- Expectation is a linear operator

$$E[aX + b] = aE[X] + b$$

Visual Navigation for Flying Robots

27

Dr. Jürgen Sturm, Computer Vision Group, TUM

Covariance of a Random Variable

- Measures the squared expected deviation from the mean

$$\text{Cov}[X] = E[X - E[X]]^2 = E[X^2] - E[X]^2$$

Visual Navigation for Flying Robots

28

Dr. Jürgen Sturm, Computer Vision Group, TUM

The State Estimation Problem

We want to estimate the world state x

- From sensor measurements z
- and controls (or odometry readings) u

We need to model the relationship between these random variables, i.e.,

$$p(x | z) \quad p(x' | x, u)$$

Visual Navigation for Flying Robots

29

Dr. Jürgen Sturm, Computer Vision Group, TUM

Causal vs. Diagnostic Reasoning

- $P(x | z)$ is diagnostic
- $P(z | x)$ is causal
- Often causal knowledge is easier to obtain
- Bayes rule allows us to use causal knowledge:

$$P(x | z) = \frac{\overset{\text{observation likelihood}}{\downarrow} P(z | x) \overset{\text{prior on world state}}{\downarrow} P(x)}{\underset{\text{prior on sensor observations}}{\uparrow} P(z)}$$

Visual Navigation for Flying Robots

30

Dr. Jürgen Sturm, Computer Vision Group, TUM

Bayes Formula

$$P(x, z) = P(x | z)P(z) = P(z | x)P(x)$$

⇒

$$P(x | z) = \frac{P(z | x)P(x)}{P(z)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

Normalization


- Direct computation of $P(z)$ can be difficult
- Idea: Compute improper distribution, normalize afterwards
- Step 1: $L(x | z) = P(z | x)P(x)$
- Step 2: $P(z) = \sum_x P(z | x)P(x) = \sum_x L(x | z)$
(Law of total probability)
- Step 3: $P(x | z) = L(x | z) / P(z)$

Bayes Rule with Background Knowledge

$$P(x | y, z) = \frac{P(y | x, z)P(x | z)}{P(y | z)}$$

Example: Sensor Measurement

- Quadcopter seeks the landing zone
- Landing zone is marked with many bright lamps
- Quadcopter has a brightness sensor



Example: Sensor Measurement

- Binary sensor $Z \in \{\text{bright}, \neg\text{bright}\}$
- Binary world state $X \in \{\text{home}, \neg\text{home}\}$
- Sensor model $P(Z = \text{bright} | X = \text{home}) = 0.6$
 $P(Z = \text{bright} | X = \neg\text{home}) = 0.3$
- Prior on world state $P(X = \text{home}) = 0.5$
- Assume: Robot observes light, i.e., $Z = \text{bright}$
- What is the probability $P(X = \text{home} | Z = \text{bright})$ that the robot is above the landing zone?

Example: Sensor Measurement

- Sensor model $P(Z = \text{bright} | X = \text{home}) = 0.6$
 $P(Z = \text{bright} | X = \neg\text{home}) = 0.3$
- Prior on world state $P(X = \text{home}) = 0.5$
- Probability after observation (using Bayes)

$$\begin{aligned} P(X = \text{home} | Z = \text{bright}) &= \frac{P(\text{bright} | \text{home})P(\text{home})}{P(\text{bright} | \text{home})P(\text{home}) + P(\text{bright} | \neg\text{home})P(\neg\text{home})} \\ &= \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{0.3}{0.3 + 0.15} = 0.67 \end{aligned}$$

Example: Sensor Measurement

- Sensor model $P(Z = \text{bright} \mid X = \text{home}) = 0.6$
 $P(Z = \text{bright} \mid X = \neg\text{home}) = 0.3$
- Prior on world state $P(X = \text{home}) = 0.5$
- Probability after observation (using Bayes)

$$\begin{aligned}
 P(X = \text{home} \mid Z = \text{noise}) &= \frac{P(\text{bright} \mid \text{home})P(\text{home})}{P(\text{bright} \mid \text{home})P(\text{home}) + P(\text{bright} \mid \neg\text{home})P(\neg\text{home})} \\
 &= \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{0.3}{0.3 + 0.15} = 0.67
 \end{aligned}$$

Visual Navigation for Flying Robots

37

Dr. Jürgen Sturm, Computer Vision Group, TUM

Combining Evidence

- Suppose our robot obtains another observation z_2 (either from the same or a different sensor)
- How can we integrate this new information?
- More generally, how can we estimate $p(x \mid z_1, z_2, \dots)$?

Visual Navigation for Flying Robots

38

Dr. Jürgen Sturm, Computer Vision Group, TUM

Combining Evidence

- Suppose our robot obtains another observation z_2 (either from the same or a different sensor)
- How can we integrate this new information?
- More generally, how can we estimate $p(x \mid z_1, z_2, \dots)$?
- Bayes formula gives us:

$$P(x \mid z_1, \dots, z_n) = \frac{P(z_n \mid x, z_1, \dots, z_{n-1})P(x \mid z_1, \dots, z_{n-1})}{P(z_n \mid z_1, \dots, z_{n-1})}$$

Visual Navigation for Flying Robots

39

Dr. Jürgen Sturm, Computer Vision Group, TUM

Recursive Bayesian Updates

$$P(x \mid z_1, \dots, z_n) = \frac{P(z_n \mid x, z_1, \dots, z_{n-1})P(x \mid z_1, \dots, z_{n-1})}{P(z_n \mid z_1, \dots, z_{n-1})}$$

Visual Navigation for Flying Robots

40

Dr. Jürgen Sturm, Computer Vision Group, TUM

Recursive Bayesian Updates

$$P(x \mid z_1, \dots, z_n) = \frac{P(z_n \mid x, z_1, \dots, z_{n-1})P(x \mid z_1, \dots, z_{n-1})}{P(z_n \mid z_1, \dots, z_{n-1})}$$

- **Markov Assumption:**
 z_n is independent of z_1, \dots, z_{n-1} if we know x

Visual Navigation for Flying Robots

41

Dr. Jürgen Sturm, Computer Vision Group, TUM

Recursive Bayesian Updates

$$P(x \mid z_1, \dots, z_n) = \frac{P(z_n \mid x, z_1, \dots, z_{n-1})P(x \mid z_1, \dots, z_{n-1})}{P(z_n \mid z_1, \dots, z_{n-1})}$$

- **Markov Assumption:**
 z_n is independent of z_1, \dots, z_{n-1} if we know x

$$\begin{aligned}
 P(x \mid z_1, \dots, z_n) &= \frac{P(z_n \mid x)P(x \mid z_1, \dots, z_{n-1})}{P(z_n \mid z_1, \dots, z_{n-1})} \\
 &= \eta P(z_n \mid x)P(x \mid z_1, \dots, z_{n-1}) \\
 &= \eta_{1:n} \prod_{i=1, \dots, n} P(z_i \mid x)P(x)
 \end{aligned}$$

Visual Navigation for Flying Robots

42

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Second Measurement

- **Sensor model** $P(Z_2 = \text{marker} \mid X = \text{home}) = 0.8$
 $P(Z_2 = \text{marker} \mid X = \neg\text{home}) = 0.1$
- **Previous estimate** $P(X = \text{home} \mid Z_1 = \text{bright}) = 0.67$
- **Assume robot does not observe marker**
- **What is the probability of being home?**

$$\begin{aligned}
 &P(X = \text{home} \mid Z_1 = \text{bright}, Z_2 = \neg\text{marker}) \\
 &= \frac{P(\neg\text{marker} \mid \text{home})P(\text{home} \mid \text{bright})}{P(\neg\text{marker} \mid \text{home})P(\text{home} \mid \text{bright}) + P(\neg\text{marker} \mid \neg\text{home})P(\neg\text{home} \mid \text{bright})} \\
 &= \frac{0.2 \cdot 0.67}{0.2 \cdot 0.67 + 0.9 \cdot 0.33} = 0.31
 \end{aligned}$$

Example: Second Measurement

- **Sensor model** $P(Z_2 = \text{marker} \mid X = \text{home}) = 0.8$
 $P(Z_2 = \text{marker} \mid X = \neg\text{home}) = 0.1$
- **Previous estimate** $P(X = \text{home} \mid Z_1 = \text{bright}) = 0.67$
- **Assume robot does not observe marker**
- **What is the probability of being home?**

$$\begin{aligned}
 &P(X = \text{home} \mid Z_1 = \text{bright}, Z_2 = \neg\text{marker}) \\
 &= \frac{P(\neg\text{marker} \mid \text{home})P(\text{home} \mid \text{bright})}{P(\neg\text{marker} \mid \text{home})P(\text{home} \mid \text{bright}) + P(\neg\text{marker} \mid \neg\text{home})P(\neg\text{home} \mid \text{bright})} \\
 &= \frac{0.2 \cdot 0.67}{0.2 \cdot 0.67 + 0.9 \cdot 0.33} = 0.31
 \end{aligned}$$

The second observation lowers the probability that the robot is above the landing zone!

Actions (Motions)

- Often the world is dynamic since
 - actions carried out by the robot...
 - actions carried out by other agents...
 - or just time passing by...
 ...change the world
- How can we incorporate actions?

Typical Actions

- Quadcopter accelerates by changing the speed of its motors
- Position also changes when quadcopter does “nothing” (and drifts..)
- Actions are never carried out with absolute certainty
- In contrast to measurements, actions generally increase the uncertainty of the state estimate

Action Models


- To incorporate the outcome of an action u into the current state estimate (“belief”), we use the conditional pdf

$$p(x' \mid u, x)$$

- This term specifies the probability that executing the action u in state x will lead to state x'

Example: Take-Off

- **Action:** $u \in \{\text{takeoff}\}$
- **World state:** $x \in \{\text{ground}, \text{air}\}$



Integrating the Outcome of Actions

- Discrete case

$$P(x' | u) = \sum_x P(x' | u, x)P(x)$$

- Continuous case

$$p(x' | u) = \int p(x' | u, x)p(x)dx$$

Visual Navigation for Flying Robots

49

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Take-Off

- Prior belief on robot state: $P(x = \text{ground}) = 1.0$ (robot is located on the ground)
- Robot executes "take-off" action
- What is the robot's belief after one time step?

$$\begin{aligned} P(x' = \text{ground}) &= \sum_x P(x' = \text{ground} | u, x)P(x) \\ &= P(x' = \text{ground} | u, x = \text{ground})P(x = \text{ground}) \\ &\quad + P(x' = \text{ground} | u, x = \text{air})P(x = \text{air}) \\ &= 0.1 \cdot 1.0 + 0.01 \cdot 0.0 = 0.1 \end{aligned}$$

- Question: What is the probability at $t=2$?


Visual Navigation for Flying Robots

50

Dr. Jürgen Sturm, Computer Vision Group, TUM

Markov Chain

- A Markov chain is a stochastic process where, given the present state, the past and the future states are independent



Visual Navigation for Flying Robots

51

Dr. Jürgen Sturm, Computer Vision Group, TUM

Markov Assumption

- Observations depend only on current state

$$P(z_t | x_{0:t}, z_{1:t-1}, u_{1:t}) = P(z_t | x_t)$$
- Current state depends only on previous state and current action

$$P(x_t | x_{0:t-1}, z_{1:t}, u_{1:t}) = P(x_t | x_{t-1}, u_t)$$

- Underlying assumptions
 - Static world
 - Independent noise
 - Perfect model, no approximation errors

Visual Navigation for Flying Robots

52

Dr. Jürgen Sturm, Computer Vision Group, TUM

Bayes Filter

- Given:
 - Stream of observations z and actions u :

$$\mathbf{d}_t = (u_1, z_1, \dots, u_t, z_t)^\top$$
 - Sensor model $P(z | x)$
 - Action model $P(x' | x, u)$
 - Prior probability of the system state $P(x)$
- Wanted:
 - Estimate of the state x of the dynamic system
 - Posterior of the state is also called **belief**

$$\text{Bel}(x_t) = P(x_t | u_1, z_1, \dots, u_t, z_t)$$

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Bayes Filter

For each time step, do

- Apply motion model

$$\overline{\text{Bel}}(x_t) = \sum_{x_{t-1}} P(x_t | x_{t-1}, u_t) \text{Bel}(x_{t-1})$$

- Apply sensor model

$$\text{Bel}(x_t) = \eta P(z_t | x_t) \overline{\text{Bel}}(x_t)$$

Note: Bayes filters also work on continuous state spaces (replace sum by integral)


Visual Navigation for Flying Robots

54

Dr. Jürgen Sturm, Computer Vision Group, TUM


Example: Localization

- Discrete state $x \in \{1, 2, \dots, w\} \times \{1, 2, \dots, h\}$
- Belief distribution can be represented as a grid
- This is also called a **histogram filter**




Example: Localization

- Action $u \in \{\text{north, east, south, west}\}$
- Robot can move one cell in each time step
- Actions are not perfectly executed



Example: Localization


- Action $u \in \{\text{north, east, south, west}\}$
- Robot can move one cell in each time step
- Actions are not perfectly executed
- Example: move east



60% success rate, 10% to stay/move too far/
move one up/move one down

Example: Localization


- Observation $z \in \{\text{marker}, \neg\text{marker}\}$
- One (special) location has a marker
- Marker is sometimes also detected in neighboring cells



Example: Localization


- Let's start a simulation run... (shades are hand-drawn, not exact!)

- $t=0$
- Prior distribution (initial belief)
- Assume we know the initial location (if not, we could initialize with a uniform prior)



Example: Localization

- $t=1$, $u=$ east, $z=$ no-marker
- Bayes filter step 1: Apply motion model




Visual Navigation for Flying Robots

61

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Localization

- $t=1$, $u=$ east, $z=$ no-marker
- Bayes filter step 2: Apply observation model




Visual Navigation for Flying Robots

62

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Localization

- $t=2$, $u=$ east, $z=$ marker
- Bayes filter step 2: Apply motion model




Visual Navigation for Flying Robots

63

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Localization

- $t=2$, $u=$ east, $z=$ marker
- Bayes filter step 1: Apply observation model



Visual Navigation for Flying Robots

64

Dr. Jürgen Sturm, Computer Vision Group, TUM

Bayes Filter - Summary

- Markov assumption allows efficient recursive Bayesian updates of the belief distribution
- Useful tool for estimating the state of a dynamic system
- Bayes filter is the basis of many other filters
 - **Kalman filter**
 - Particle filter
 - Hidden Markov models
 - Dynamic Bayesian networks
 - Partially observable Markov decision processes (POMDPs)

Visual Navigation for Flying Robots

65

Dr. Jürgen Sturm, Computer Vision Group, TUM

Kalman Filter

- Bayes filter with continuous states
- State represented with a normal distribution
- Developed in the late 1950's
- Kalman filter is very efficient (only requires a few matrix operations per time step)
- Applications range from economics, weather forecasting, satellite navigation to robotics and many more
- Most relevant Bayes filter variant in practice
→ exercise sheet 2

Visual Navigation for Flying Robots

66


Dr. Jürgen Sturm, Computer Vision Group, TUM

Normal Distribution

- Univariate normal distribution

$$X \sim \mathcal{N}(\mu, \sigma)$$

$$p(X = x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2} \frac{(x - \mu)^2}{\sigma^2}\right)$$



Visual Navigation for Flying Robots

67

Dr. Jürgen Sturm, Computer Vision Group, TUM

Normal Distribution


- Multivariate normal distribution

$$X \sim \mathcal{N}(\mu, \Sigma)$$

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mu, \Sigma)$$

$$= \frac{1}{(2\pi)^{d/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^\top \Sigma^{-1}(\mathbf{x} - \mu)\right)$$

- Example: 2-dimensional normal distribution



Visual Navigation for Flying Robots

68

Dr. Jürgen Sturm, Computer Vision Group, TUM

Properties of Normal Distributions

- Linear transformation \rightarrow remains Gaussian

$$X \sim \mathcal{N}(\mu, \Sigma), Y \sim AX + B$$

$$\Rightarrow Y \sim \mathcal{N}(A\mu + B, A\Sigma A^\top)$$

- Intersection of two Gaussians \rightarrow remains Gaussian

$$X_1 \sim \mathcal{N}(\mu_1, \Sigma_1), X_2 \sim \mathcal{N}(\mu_2, \Sigma_2)$$

$$\Rightarrow p(X_1, X_2) = \mathcal{N}\left(\frac{\Sigma_2}{\Sigma_1 + \Sigma_2}\mu_1 + \frac{\Sigma_1}{\Sigma_1 + \Sigma_2}\mu_2, \frac{1}{\Sigma_1^{-1} + \Sigma_2^{-1}}\right)$$


Visual Navigation for Flying Robots

69

Dr. Jürgen Sturm, Computer Vision Group, TUM

Linear Process Model

- Consider a time-discrete stochastic process (Markov chain)



Visual Navigation for Flying Robots

70

Dr. Jürgen Sturm, Computer Vision Group, TUM

Linear Process Model

- Consider a time-discrete stochastic process
 - Represent the estimated state (belief) by a Gaussian
- $$x_t \sim \mathcal{N}(\mu_t, \Sigma_t)$$

Visual Navigation for Flying Robots

71

Dr. Jürgen Sturm, Computer Vision Group, TUM

Linear Process Model

- Consider a time-discrete stochastic process
- Represent the estimated state (belief) by a Gaussian
- Assume that the system evolves linearly over time, then

$$x_t = Ax_{t-1}$$

Visual Navigation for Flying Robots

72

Dr. Jürgen Sturm, Computer Vision Group, TUM

Linear Process Model

- Consider a time-discrete stochastic process
- Represent the estimated state (belief) by a Gaussian $x_t \sim \mathcal{N}(\mu_t, \Sigma_t)$
- Assume that the system evolves linearly over time and depends linearly on the controls

$$x_t = Ax_{t-1} + Bu_t$$

Visual Navigation for Flying Robots

73

Dr. Jürgen Sturm, Computer Vision Group, TUM

Linear Process Model

- Consider a time-discrete stochastic process
- Represent the estimated state (belief) by a Gaussian $x_t \sim \mathcal{N}(\mu_t, \Sigma_t)$
- Assume that the system evolves linearly over time, depends linearly on the controls, and has zero-mean, normally distributed process noise

$$x_t = Ax_{t-1} + Bu_t + \epsilon_t$$

with $\epsilon_t \sim \mathcal{N}(0, Q)$

Visual Navigation for Flying Robots

74

Dr. Jürgen Sturm, Computer Vision Group, TUM

Linear Observations

- Further, assume we make observations that depend linearly on the state

$$z_t = Cx_t$$

Visual Navigation for Flying Robots

75

Dr. Jürgen Sturm, Computer Vision Group, TUM

Linear Observations

- Further, assume we make observations that depend linearly on the state and that are perturbed by zero-mean, normally distributed observation noise

$$z_t = Cx_t + \delta_t$$

with $\delta_t \sim \mathcal{N}(0, R)$

Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM

Kalman Filter

Estimates the state x_t of a discrete-time controlled process that is governed by the linear stochastic difference equation

$$x_t = Ax_{t-1} + Bu_t + \epsilon_t$$

and (linear) measurements of the state

$$z_t = Cx_t + \delta_t$$

with $\delta_t \sim \mathcal{N}(0, R)$ and $\epsilon_t \sim \mathcal{N}(0, Q)$

Visual Navigation for Flying Robots

77

Dr. Jürgen Sturm, Computer Vision Group, TUM

Variables and Dimensions

- State $x \in \mathbb{R}^n$
- Controls $u \in \mathbb{R}^l$
- Observations $z \in \mathbb{R}^k$
- Process equation

$$x_t = \underbrace{A}_{n \times n} x_{t-1} + \underbrace{B}_{n \times l} u_t + \epsilon$$

- Measurement equation

$$z_t = \underbrace{C}_{n \times k} x_t + \delta_t$$

Visual Navigation for Flying Robots

78

Dr. Jürgen Sturm, Computer Vision Group, TUM

Kalman Filter

- Initial belief is Gaussian

$$\text{Bel}(x_0) = \mathcal{N}(x_0; \mu_0, \Sigma_0)$$

- Next state is also Gaussian (linear transformation)

$$x_t \sim \mathcal{N}(Ax_{t-1} + Bu_t, Q)$$

- Observations are also Gaussian

$$z_t \sim \mathcal{N}(Cx_t, R)$$

From Bayes Filter to Kalman Filter

For each time step, do

- Apply motion model

$$\overline{\text{Bel}}(x_t) = \int \underbrace{p(x_t | x_{t-1}, u_t)}_{\mathcal{N}(x_t; Ax_t + Bu_t, Q)} \underbrace{\text{Bel}(x_{t-1})}_{\mathcal{N}(x_{t-1}; \mu_{t-1}, \Sigma_{t-1})} dx_{t-1}$$

From Bayes Filter to Kalman Filter

For each time step, do

- Apply motion model

$$\begin{aligned} \overline{\text{Bel}}(x_t) &= \int \underbrace{p(x_t | x_{t-1}, u_t)}_{\mathcal{N}(x_t; Ax_{t-1} + Bu_t, Q)} \underbrace{\text{Bel}(x_{t-1})}_{\mathcal{N}(x_{t-1}; \mu_{t-1}, \Sigma_{t-1})} dx_{t-1} \\ &= \mathcal{N}(x_t; A\mu_{t-1} + Bu_t, A\Sigma A^\top + Q) \\ &= \mathcal{N}(x_t; \bar{\mu}_t, \bar{\Sigma}_t) \end{aligned}$$

From Bayes Filter to Kalman Filter

For each time step, do

- Apply sensor model

$$\begin{aligned} \text{Bel}(x_t) &= \eta \underbrace{p(z_t | x_t)}_{\mathcal{N}(z_t; Cx_t, R)} \underbrace{\overline{\text{Bel}}(x_t)}_{\mathcal{N}(x_t; \bar{\mu}_t, \bar{\Sigma}_t)} \\ &= \mathcal{N}(x_t; \bar{\mu}_t + K_t(z_t - C\bar{\mu}_t), (I - K_t C)\bar{\Sigma}_t) \\ &= \mathcal{N}(x_t; \mu_t, \Sigma_t) \end{aligned}$$

$$\text{with } K_t = \bar{\Sigma}_t C^\top (C\bar{\Sigma}_t C^\top + R)^{-1}$$

Kalman Filter

For each time step, do

- Apply motion model

$$\begin{aligned} \bar{\mu}_t &= A\mu_{t-1} + Bu_t \\ \bar{\Sigma}_t &= A\Sigma A^\top + Q \end{aligned}$$

- Apply sensor model

$$\begin{aligned} \mu_t &= \bar{\mu}_t + K_t(z_t - C\bar{\mu}_t) \\ \Sigma_t &= (I - K_t C)\bar{\Sigma}_t \end{aligned}$$

$$\text{with } K_t = \bar{\Sigma}_t C^\top (C\bar{\Sigma}_t C^\top + R)^{-1}$$

For the interested readers:
See Probabilistic Robotics for
full derivation (Chapter 3)

Kalman Filter

- Highly efficient: Polynomial in the measurement dimensionality k and state dimensionality n :

$$O(k^{2.376} + n^2)$$

- Optimal for linear Gaussian systems!**
- Most robotics systems are **nonlinear!**

Nonlinear Dynamical Systems

- Most realistic robotic problems involve nonlinear functions
- Motion function

$$x_t = g(u_t, x_{t-1})$$

- Observation function

$$z_t = h(x_t)$$

Visual Navigation for Flying Robots

85

Dr. Jürgen Sturm, Computer Vision Group, TUM

Taylor Expansion

- Solution: Linearize both functions
- Motion function

$$\begin{aligned} g(x_{t-1}, u_t) &\approx g(\mu_{t-1}, u_t) + \frac{\partial g(\mu_{t-1}, u_t)}{\partial x_{t-1}}(x_{t-1} - \mu_{t-1}) \\ &= g(\mu_{t-1}, u_t) + G_t(x_{t-1} - \mu_{t-1}) \end{aligned}$$

- Observation function

$$\begin{aligned} h(x_t) &\approx h(\bar{\mu}_t) + \frac{\partial h(\bar{\mu}_t)}{\partial x_t}(x_t - \mu_t) \\ &= h(\bar{\mu}_t) + H_t(x_t - \mu_t) \end{aligned}$$

Visual Navigation for Flying Robots

86

Dr. Jürgen Sturm, Computer Vision Group, TUM

Extended Kalman Filter

For each time step, do

- Apply motion model

$$\begin{aligned} \bar{\mu}_t &= g(\mu_{t-1}, u_t) \\ \bar{\Sigma}_t &= G_t \Sigma_t G_t^\top + Q \quad \text{with } G_t = \frac{\partial g(\mu_{t-1}, u_t)}{\partial x_{t-1}} \end{aligned}$$

- Apply sensor model

$$\begin{aligned} \mu_t &= \bar{\mu}_t + K_t(z_t - h(\bar{\mu}_t)) \\ \Sigma_t &= (I - K_t H_t) \bar{\Sigma}_t \end{aligned}$$

with $K_t = \bar{\Sigma}_t H_t^\top (H_t \bar{\Sigma}_t H_t^\top + R)^{-1}$ and $H_t = \frac{\partial h(\bar{\mu}_t)}{\partial x_t}$

For the interested readers:
See Probabilistic Robotics for
full derivation (Chapter 3)

Visual Navigation for Flying Robots

87

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- 2D case
- State $\mathbf{x} = (x \ y \ \psi)^\top$
- Odometry $\mathbf{u} = (\dot{x} \ \dot{y} \ \dot{\psi})^\top$
- Observations of visual marker $\mathbf{z} = (x \ y \ \psi)^\top$ (relative to robot pose)

Visual Navigation for Flying Robots

88

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Motion Function and its derivative

$$g(\mathbf{x}, \mathbf{u}) = \begin{pmatrix} x + (\cos(\psi)\dot{x} - \sin(\psi)\dot{y})\Delta t \\ y + (\sin(\psi)\dot{x} + \cos(\psi)\dot{y})\Delta t \\ \psi + \dot{\psi}\Delta t \end{pmatrix}$$

$$G = \frac{\partial g(\mathbf{x}, \mathbf{u})}{\partial \mathbf{x}} = \begin{pmatrix} 1 & 0 & (-\sin(\psi)\dot{x} - \cos(\psi)\dot{y})\Delta t \\ 0 & 1 & (\cos(\psi)\dot{x} + \sin(\psi)\dot{y})\Delta t \\ 0 & 0 & 1 \end{pmatrix}$$

Visual Navigation for Flying Robots

89

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Observation Function (→ Sheet 2)

$$h(\mathbf{x}) = \dots$$

$$H = \frac{\partial h(\mathbf{x})}{\partial \mathbf{x}} = \dots$$


Visual Navigation for Flying Robots

90

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Dead reckoning (no observations)
- Large process noise in $x+y$




Visual Navigation for Flying Robots

91 Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Dead reckoning (no observations)
- Large process noise in $x+y+yaw$




Visual Navigation for Flying Robots

92 Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Now with observations (limited visibility)
- Assume robot knows correct starting pose




Visual Navigation for Flying Robots

93

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- What if the initial pose ($x+y$) is wrong?




Visual Navigation for Flying Robots

94

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- What if the initial pose ($x+y+yaw$) is wrong?




Visual Navigation for Flying Robots

95

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- If we are aware of a bad initial guess, we set the initial sigma to a large value (large uncertainty)




Visual Navigation for Flying Robots

96

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example



Visual Navigation for Flying Robots

97

Dr. Jürgen Sturm, Computer Vision Group, TUM

Summary

- Observations and actions are inherently noisy
- Knowledge about state is inherently uncertain
- Probability theory
- Probabilistic sensor and motion models
- Bayes Filter, Histogram Filter, Kalman Filter, Examples

Visual Navigation for Flying Robots

98

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

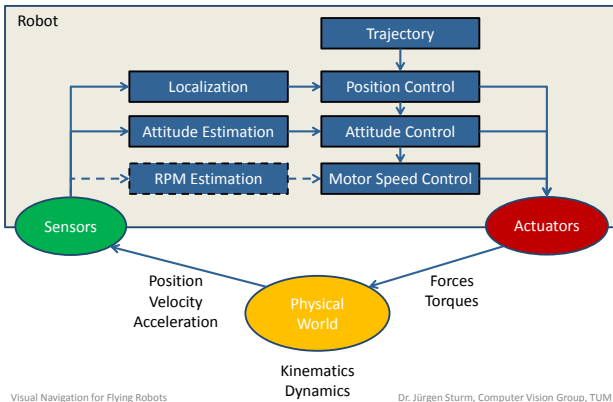
Robot Control

Dr. Jürgen Sturm

Organization - Exam


- Oral exams **in teams** (2-3 students)
- At least 15 minutes per student
→ individual grades
- Questions will address
 - Material from the lecture
 - Material from the exercise sheets
 - Your mini-project

Control Architecture





DC Motors

- Maybe you built one in school
- Stationary permanent magnet
- Electromagnet induces torque
- Split ring switches direction of current



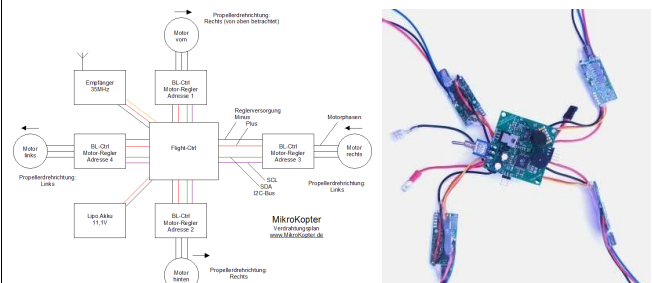
Brushless Motors

- Used in most quadrocopters
- Permanent magnets on the axis
- Electromagnets on the outside
- Requires motor controller to switch currents
→ Does not require brushes (less maintenance)



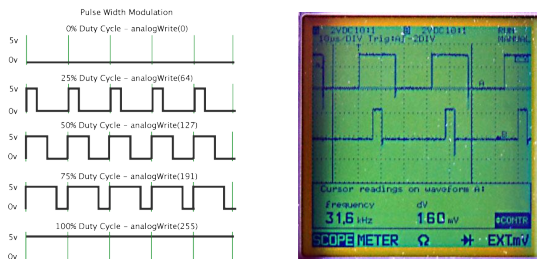
Attitude + Motor Controller Boards

- Example: Mikrokopter Plattform



Pulse Width Modulation (PWM)

- Protocol used to control motor speed
- Remote controls typically output PWM




Visual Navigation for Flying Robots

7

Dr. Jürgen Sturm, Computer Vision Group, TUM

I2C Protocol

- Serial data line (SDA) + serial clock line (SCL)
- All devices connected in parallel
- 7-10 bit address, 100-3400 kbit/s speed
- Used by Mikrocopter for motor control

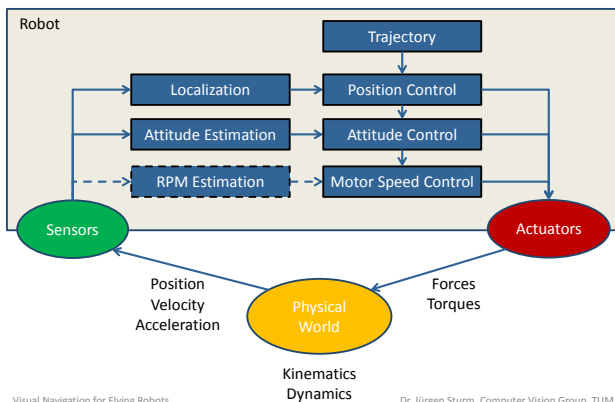


Visual Navigation for Flying Robots

8

Dr. Jürgen Sturm, Computer Vision Group, TUM

Control Architecture



Visual Navigation for Flying Robots

Dr. Jürgen Sturm, Computer Vision Group, TUM

Kinematics and Dynamics

- Kinematics
 - Integrate acceleration to get velocity
 - Integrate velocity to get position
- Dynamics
 - Actuators induce forces and torques
 - Forces induce linear acceleration
 - Torques induce angular acceleration
- What types of forces do you know?
- What types of torques do you know?

Visual Navigation for Flying Robots

10

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: 1D Kinematics

- State $\mathbf{x} = (x \ \dot{x} \ \ddot{x})^T \in \mathbb{R}^3$
- Action $u \in \mathbb{R}$
- Process model

$$\mathbf{x}_t = \begin{pmatrix} 1 & \Delta t & 0 \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{pmatrix} \mathbf{x}_{t-1} + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} u_t$$

- Kalman filter
- How many states do we need for 3D?

Visual Navigation for Flying Robots

11

Dr. Jürgen Sturm, Computer Vision Group, TUM

Dynamics - Essential Equations

- Force (Kraft)

$$m\ddot{\mathbf{x}} = \sum_i \mathbf{F}_i$$

- Torque (Drehmoment)

$$J\dot{\boldsymbol{\alpha}} = \sum_i \boldsymbol{\tau}_i$$

Visual Navigation for Flying Robots

12

Dr. Jürgen Sturm, Computer Vision Group, TUM

Forces

- Gravity $F_{\text{grav}} = mg$
- Friction
 - Stiction (static friction) $F_{\text{stiction}} = c_s \text{sign } \dot{x}$
 - Damping (viscous friction) $F_{\text{damping}} = D\dot{x}$
- Spring $F_{\text{spring}} = K(x - x_{\text{eq}})$
- Magnetic force
- ...



Visual Navigation for Flying Robots

13

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Spring-Damper System

- Combination of spring and damper
- Forces $F = F_{\text{damping}} + F_{\text{spring}}$
- Resulting dynamics $m\ddot{x} = D\dot{x} + K(x - x_{\text{eq}})$




Visual Navigation for Flying Robots

14

Dr. Jürgen Sturm, Computer Vision Group, TUM

Torques

- Definition $\tau = F \times r$
- Torques sum up $\tau_{\text{net}} = \sum \tau_i$
- Torque results in angular acceleration $\tau = J\alpha$ (with $\alpha = \frac{d\omega}{dt}$, J moment of inertia)
- Friction same as before...




Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

Dynamics of a Quadcopter

- Each propeller induces force and torque by accelerating air
- Gravity pulls quadcopter downwards




Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Vertical Acceleration

- Thrust $F_{\text{thrust}} = F_1 + F_2 + F_3 + F_4$




Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

Vertical and Horizontal Acceleration

- Thrust $F_{\text{thrust}} = F_1 + F_2 + F_3 + F_4$




Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM

Vertical and Horizontal Acceleration

- Thrust $F_{\text{thrust}} = F_1 + F_2 + F_3 + F_4$
- Acceleration $\ddot{\mathbf{x}}_{\text{global}} = (R_{RPY} F_{\text{thrust}} - F_{\text{grav}})/m$




Visual Navigation for Flying Robots

19

Dr. Jürgen Sturm, Computer Vision Group, TUM

Pitch (and Roll)

- Attitude changes when opposite motors generate unequal thrust
- Induced torque $\tau = (F_1 - F_3) \times r$
- Induced angular acceleration




Visual Navigation for Flying Robots

20

Dr. Jürgen Sturm, Computer Vision Group, TUM

Yaw

- Each propeller induces torque due to rotation and the interaction with the air
- Induced torque $\tau = \tau_1 - \tau_2 + \tau_3 - \tau_4$
- Induced angular acceleration

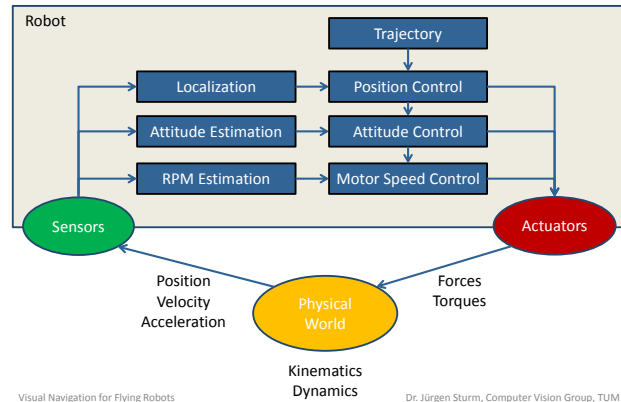


Visual Navigation for Flying Robots

21

Dr. Jürgen Sturm, Computer Vision Group, TUM

Cascaded Control



Visual Navigation for Flying Robots

Kinematics Dynamics

Dr. Jürgen Sturm, Computer Vision Group, TUM

Assumptions of Cascaded Control

- Dynamics of inner loops is so fast that it is not visible from outer loops
- Dynamics of outer loops is so slow that it appears as static to the inner loops

Visual Navigation for Flying Robots

23

Dr. Jürgen Sturm, Computer Vision Group, TUM

Cascaded Control Example

- Motor control happens on motor boards (controls every motor tick)
- Attitude control implemented on micro-controller with hard real-time (at 1000 Hz)
- Position control (at 10 – 250 Hz)
- Trajectory (waypoint) control (at 0.1 – 1 Hz)

Visual Navigation for Flying Robots

24

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feedback Control - Generic Idea


Desired value
35°

Visual Navigation for Flying Robots

25

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feedback Control - Generic Idea

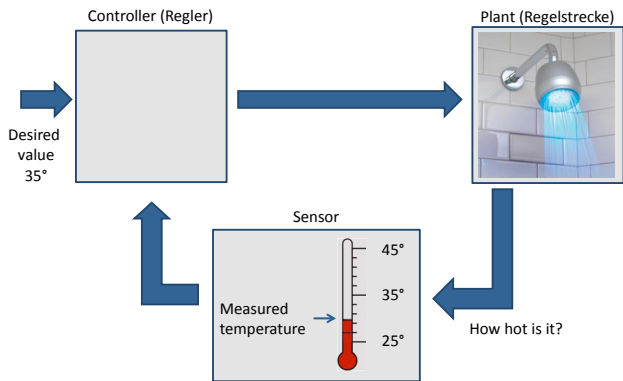


Visual Navigation for Flying Robots

26

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feedback Control - Generic Idea

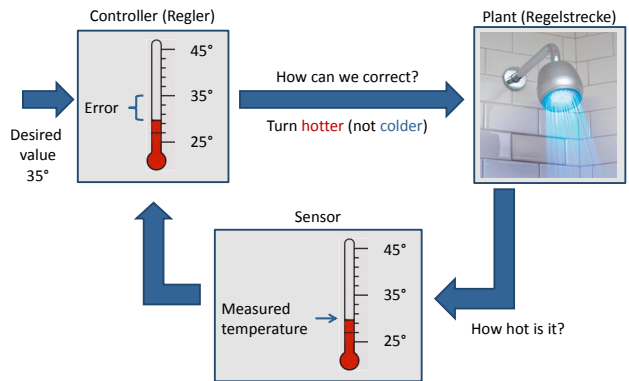


Visual Navigation for Flying Robots

27

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feedback Control - Generic Idea




Visual Navigation for Flying Robots

28

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feedback Control - Example




Visual Navigation for Flying Robots

29

Dr. Jürgen Sturm, Computer Vision Group, TUM

Measurement Noise

- What effect has noise in the measurements?



- Poor performance for $K=1$
- How can we fix this?


Visual Navigation for Flying Robots

30

Dr. Jürgen Sturm, Computer Vision Group, TUM

Proper Control with Measurement Noise

- Lower the gain... (K=0.15)




Visual Navigation for Flying Robots

31

Dr. Jürgen Sturm, Computer Vision Group, TUM

What do High Gains do?

- High gains are always problematic (K=2.15)




Visual Navigation for Flying Robots

32

Dr. Jürgen Sturm, Computer Vision Group, TUM

What happens if sign is messed up?

- Check K=-0.5




Visual Navigation for Flying Robots

33

Dr. Jürgen Sturm, Computer Vision Group, TUM

Saturation

- In practice, often the set of admissible controls u is bounded
- This is called (control) saturation




Visual Navigation for Flying Robots

34

Dr. Jürgen Sturm, Computer Vision Group, TUM

Block Diagram




Visual Navigation for Flying Robots

35

Dr. Jürgen Sturm, Computer Vision Group, TUM

Delays

- In practice most systems have delays
- Can lead to overshoots/oscillations/de-stabilization



Visual Navigation for Flying Robots


36

Dr. Jürgen Sturm, Computer Vision Group, TUM

- One solution: lower gains (why is this bad?)

Delays

- What is the total dead time of this system?




Visual Navigation for Flying Robots

37

Dr. Jürgen Sturm, Computer Vision Group, TUM

Delays

- What is the total dead time of this system?



Visual Navigation for Flying Robots


38

Dr. Jürgen Sturm, Computer Vision Group, TUM

- Can we distinguish delays in the measurement from delays in actuation?

Delays

- What is the total dead time of this system?



Visual Navigation for Flying Robots


39

Dr. Jürgen Sturm, Computer Vision Group, TUM

- Can we distinguish delays in the measurement from delays in actuation? No!

Smith Predictor

- Allows for higher gains
- Requires (accurate) model of plant




Visual Navigation for Flying Robots

40

Dr. Jürgen Sturm, Computer Vision Group, TUM

Smith Predictor

- Plant model is available
- 5 seconds delay
- Results in perfect compensation
- Why is this unrealistic in practice?




Visual Navigation for Flying Robots

41

Dr. Jürgen Sturm, Computer Vision Group, TUM

Smith Predictor

- Time delay (and plant model) is often not known accurately (or changes over time)
- What happens if time delay is **overestimated**?




Visual Navigation for Flying Robots

42

Dr. Jürgen Sturm, Computer Vision Group, TUM

Smith Predictor

- Time delay (and plant model) is often not known accurately (or changes over time)
- What happens if time delay is **underestimated**?

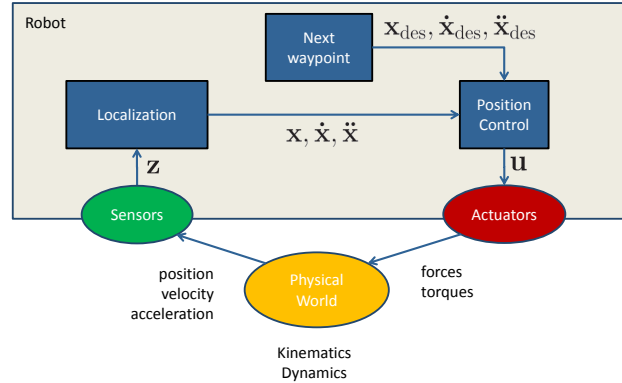


Visual Navigation for Flying Robots

43

Dr. Jürgen Sturm, Computer Vision Group, TUM

Position Control



Visual Navigation for Flying Robots

44

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rigid Body Kinematics

- Consider a rigid body
- Free floating in 1D space, no gravity
- How does this system evolve over time?


Visual Navigation for Flying Robots

45

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rigid Body Kinematics

- Consider a rigid body
- Free floating in 1D space, no gravity
- How does this system evolve over time?
- Example: $x_0 = 0, \dot{x}_0 = 0$




Visual Navigation for Flying Robots

46

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rigid Body Kinematics

- Consider a rigid body
- Free floating in 1D space, no gravity
- How does this system evolve over time?
- Example: $x_0 = 0, \dot{x}_0 = 0.1$



Visual Navigation for Flying Robots

47

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rigid Body Kinematics

- Consider a rigid body
- Free floating in 1D space, no gravity
- In each time instant, we can apply a force F
- Results in acceleration $\ddot{x} = F/m$
- Desired position $x_{des} = 1$

Visual Navigation for Flying Robots

48

Dr. Jürgen Sturm, Computer Vision Group, TUM

P Control

- What happens for this control law?

$$u_t = K(x_{\text{des}} - x_{t-1})$$
- This is called proportional control

Visual Navigation for Flying Robots


49

Dr. Jürgen Sturm, Computer Vision Group, TUM

P Control

- What happens for this control law?

$$u_t = K(x_{\text{des}} - x_{t-1})$$
- This is called proportional control



Visual Navigation for Flying Robots


50

Dr. Jürgen Sturm, Computer Vision Group, TUM

PD Control

- What happens for this control law?

$$u_t = K_P(x_{\text{des}} - x_{t-1}) + K_D(\dot{x}_{\text{des}} - \dot{x}_{t-1})$$
- Proportional-Derivative control



Visual Navigation for Flying Robots


51

Dr. Jürgen Sturm, Computer Vision Group, TUM

PD Control

- What happens for this control law?

$$u_t = K_P(x_{\text{des}} - x_{t-1}) + K_D(\dot{x}_{\text{des}} - \dot{x}_{t-1})$$
- What if we set **higher** gains?



Visual Navigation for Flying Robots


52

Dr. Jürgen Sturm, Computer Vision Group, TUM

PD Control

- What happens for this control law?

$$u_t = K_P(x_{\text{des}} - x_{t-1}) + K_D(\dot{x}_{\text{des}} - \dot{x}_{t-1})$$
- What if we set **lower** gains?




Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

PD Control

- What happens when we add gravity?



Visual Navigation for Flying Robots


54

Dr. Jürgen Sturm, Computer Vision Group, TUM

Gravity compensation

- Add as an additional term in the control law

$$u_t = K_P(x_{des} - x_{t-1}) + K_D(\dot{x}_{des} - \dot{x}_{t-1}) + F_{grav}$$
- Any known (inverse) dynamics can be included




Visual Navigation for Flying Robots

55

Dr. Jürgen Sturm, Computer Vision Group, TUM

PD Control

- What happens when we have systematic errors? (noise with non-zero mean)
- Example: unbalanced quadcopter, wind, ...
- Does the robot ever reach its desired location?



Visual Navigation for Flying Robots


56

Dr. Jürgen Sturm, Computer Vision Group, TUM

PID Control

- Idea: Estimate the system error (bias) by integrating the error

$$u_t = K_P(x_{des} - x_t) + K_D(\dot{x}_{des} - \dot{x}_t) + K_I \int_{-\infty}^t x_{des} - x_t dt$$
- Proportional+Derivative+Integral Control



Visual Navigation for Flying Robots

57

Dr. Jürgen Sturm, Computer Vision Group, TUM

PID Control

- Idea: Estimate the system error (bias) by integrating the error

$$u_t = K_P(x_{des} - x_t) + K_D(\dot{x}_{des} - \dot{x}_t) + K_I \int_{-\infty}^t x_{des} - x_t dt$$
- Proportional+Derivative+Integral Control
- For steady state systems, this can be reasonable
- Otherwise, it may create havoc or even disaster (wind-up effect)


Visual Navigation for Flying Robots

58

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Wind-up effect

- Quadcopter gets stuck in a tree → does not reach steady state
- What is the effect on the I-term?




Visual Navigation for Flying Robots

59

Dr. Jürgen Sturm, Computer Vision Group, TUM

De-coupled Control

- So far, we considered only single-input, single-output systems (SISO)
- Real systems have multiple inputs + outputs
- MIMO (multiple-input, multiple-output)
- In practice, control is often de-coupled




Visual Navigation for Flying Robots

60

Dr. Jürgen Sturm, Computer Vision Group, TUM

How to Choose the Coefficients?

- Gains too large: overshooting, oscillations
- Gains too small: long time to converge
- Heuristic methods exist
- In practice, often tuned manually



Visual Navigation for Flying Robots


61

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Ardrone

Cascaded control

- Inner loop runs on embedded PC and stabilizes flight
- Outer loop runs externally and implements position control



Visual Navigation for Flying Robots

62

Dr. Jürgen Sturm, Computer Vision Group, TUM

Ardrone: Inner Control Loop

- Plant input: motor torques

$$\mathbf{u}_{\text{inner}} = (\tau_1 \quad \tau_2 \quad \tau_3 \quad \tau_4)^\top$$

- Plant output: roll, pitch, yaw rate, z velocity

$$\mathbf{x}_{\text{inner}} = (\omega_x \quad \omega_y \quad \dot{\omega}_z \quad \dot{z})^\top$$

attitude
(measured using gyro +
accelerometer)
z velocity
(measured using ultrasonic
distance sensor + attitude)

Visual Navigation for Flying Robots

63

Dr. Jürgen Sturm, Computer Vision Group, TUM

Ardrone: Outer Control Loop

- Outer loop sees inner loop as a plant (black box)
- Plant input: roll, pitch, yaw rate, z velocity

$$\mathbf{u}_{\text{outer}} = (\omega_x \quad \omega_y \quad \dot{\omega}_z \quad \dot{z})^\top$$

- Plant output:

$$\mathbf{x}_{\text{outer}} = (x \quad y \quad z \quad \psi)^\top$$

Visual Navigation for Flying Robots

64

Dr. Jürgen Sturm, Computer Vision Group, TUM

Mechanical Equivalent

- PD Control is equivalent to adding spring-dampers between the desired values and the current position



Visual Navigation for Flying Robots

65

Dr. Jürgen Sturm, Computer Vision Group, TUM

PID Control – Summary

PID is the most used control technique in practice

- P control → simple proportional control, often enough
- PI control → can compensate for bias (e.g., wind)
- PD control → can be used to reduce overshoot (e.g., when acceleration is controlled)
- PID control → all of the above

Visual Navigation for Flying Robots

66

Dr. Jürgen Sturm, Computer Vision Group, TUM

Optimal Control

What other control techniques do exist?

- Linear-quadratic regulator (LQR)
- Reinforcement learning
- Inverse reinforcement learning
- ... and many more

Optimal Control

- Find the controller that provides the best performance
- Need to define a measure of performance
- What would be a good performance measure?
 - Minimize the error?
 - Minimize the controls?
 - Combination of both?

Linear Quadratic Regulator

Given:

- Discrete-time **linear** system

$$x_{k+1} = Ax_k + Bu_k$$

- **Quadratic** cost function

$$J = \sum_{k=0}^{\infty} (x_k^T Q x_k + u_k^T R u_k)$$

Goal: Find the controller with the lowest cost →
LQR control

Reinforcement Learning

- In principle, any measure can be used
- Define reward for each state-action pair
$$R(x_t, u_t)$$
- Find the policy (controller) that maximizes the expected future reward
- Compute the expected future reward based on
 - Known process model
 - Learned process model (from demonstrations)

Inverse Reinforcement Learning

- Parameterized reward function
- Learn these parameters from expert demonstrations and refine
- Example: [Abbeel and Ng, ICML 2010]




Interesting Papers at ICRA 2012

- Flying robots are a hot topic in the robotics community
- 4 (out of 27) sessions on flying robots, 4 sessions on localization and mapping
- Robots: quadcopters, nano quadcopters, fixed-wing airplanes
- Sensors: monocular cameras, Kinect, motion capture, laser-scanners

Autonomous Indoor 3D Exploration with a Micro-Aerial Vehicle

Shaojie Shen, Nathan Michael, and Vijay Kumar

- Map a previously unknown building
- Find good exploration frontiers in partial map



Visual Navigation for Flying Robots


73

Dr. Jürgen Sturm, Computer Vision Group, TUM

Decentralized Formation Control with Variable Shapes for Aerial Robots

Matthew Turpin, Nathan Michael, and Vijay Kumar

- Move in formation (e.g., to traverse a window)
- Avoid collisions
- Dynamic role switching



Visual Navigation for Flying Robots



74

Dr. Jürgen Sturm, Computer Vision Group, TUM

Versatile Distributed Pose Estimation and Sensor Self-Calibration for an Autonomous MAV

Stephan Weiss, Markus W. Achtelik, Margarita Chli, Roland Siegwart


- IMU, camera
- EKF for pose, velocity, sensor bias, scale, inter-sensor calibration



Visual Navigation for Flying Robots

75

Dr. Jürgen Sturm, Computer Vision Group, TUM



On-board Velocity Estimation and Closed-loop Control of a Quadrotor UAV based on Optical Flow

Volker Grabe, Heinrich H. Bühlhoff, and Paolo Robuffo Giordano


- Ego-motion from optical flow using homography constraint
- Use for velocity control



Visual Navigation for Flying Robots

76


Dr. Jürgen Sturm, Computer Vision Group, TUM



Autonomous Landing of a VTOL UAV on a Moving Platform Using Image-based Visual Servoing

Daewon Lee, Tyler Ryan and H. Jin. Kim


- Tracking and landing on a moving platform
- Switch between tracking and landing behavior



Visual Navigation for Flying Robots

77


Dr. Jürgen Sturm, Computer Vision Group, TUM



Resonant Wireless Power Transfer to Ground Sensors from a UAV

Brent Griffin and Carrick Detweiler


- Quadcopter transfers power to light a LED



Visual Navigation for Flying Robots

78


Dr. Jürgen Sturm, Computer Vision Group, TUM



Using Depth in Visual Simultaneous Localisation and Mapping

Sebastian A. Scherer, Daniel Dube and Andreas Zell

- Combine PTAM with Kinect
- Monocular SLAM: scale drift
- Kinect: has small maximum range



Visual Navigation for Flying Robots

79

Dr. Jürgen Sturm, Computer Vision Group, TUM

ICRA Papers

- Will put them in our paper repository
- Remember password (or ask by mail)
- See course website

Visual Navigation for Flying Robots

80

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

Visual Motion Estimation

Dr. Jürgen Sturm

Organization: Exam

- Registration deadline: June 30
- Course ends: July 19
- Examination dates: t.b.a. (mid August)
 - Oral team exam
 - Sign up for a time slot starting from Mid July
 - List will be placed on blackboard in front of our secretary

Motivation




Visual Motion Estimation

- Quick geometry recap
- Image filters
- 2D image alignment
- Corner detectors
- Kanade-Lucas-Tomasi tracker
- 3D motion estimation


Angular and linear velocities

- Linear velocity $\mathbf{v} = (v_x, v_y, v_z)^T \in \mathbb{R}^3$
- Angular velocity $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^T \in \mathbb{R}^3$
- Linear and angular velocity together form a twist $\boldsymbol{\xi} = (\mathbf{v}^T, \boldsymbol{\omega}^T)^T$




Angular and linear velocities

- Linear velocity $\mathbf{v} = (v_x, v_y, v_z)^T \in \mathbb{R}^3$
- Angular velocity $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^T \in \mathbb{R}^3$
- Now consider a 3D point $\mathbf{p} \in \mathbb{R}^3$ of a rigid body moving with twist $\boldsymbol{\xi} = (\mathbf{v}^T, \boldsymbol{\omega}^T)^T$



Angular and linear velocities

- Linear velocity $\mathbf{v} = (v_x, v_y, v_z)^\top \in \mathbb{R}^3$
- Angular velocity $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^\top \in \mathbb{R}^3$
- Now consider a 3D point $\mathbf{p} \in \mathbb{R}^3$ of a rigid body moving with twist $\boldsymbol{\xi} = (\mathbf{v}^\top, \boldsymbol{\omega}^\top)^\top$
- What is the velocity $\dot{\mathbf{p}}$ at point \mathbf{p} ?



Visual Navigation for Flying Robots


7

Dr. Jürgen Sturm, Computer Vision Group, TUM

Angular and linear velocities

- Linear velocity $\mathbf{v} = (v_x, v_y, v_z)^\top \in \mathbb{R}^3$
- Angular velocity $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^\top \in \mathbb{R}^3$
- Now consider a 3D point $\mathbf{p} \in \mathbb{R}^3$ of a rigid body moving with twist $\boldsymbol{\xi} = (\mathbf{v}^\top, \boldsymbol{\omega}^\top)^\top$
- What is the velocity $\dot{\mathbf{p}}$ at point \mathbf{p} ?

$$\begin{aligned} \mathbf{p}(t) &= R(t)\mathbf{p}(0) + \mathbf{t}(t) \\ &= \exp([\boldsymbol{\omega}]_{\times}t)\mathbf{p}(0) + \mathbf{v}t \end{aligned}$$



Visual Navigation for Flying Robots

8


Dr. Jürgen Sturm, Computer Vision Group, TUM

Angular and linear velocities

- Linear velocity $\mathbf{v} = (v_x, v_y, v_z)^\top \in \mathbb{R}^3$
- Angular velocity $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^\top \in \mathbb{R}^3$
- Now consider a 3D point $\mathbf{p} \in \mathbb{R}^3$ of a rigid body moving with twist $\boldsymbol{\xi} = (\mathbf{v}^\top, \boldsymbol{\omega}^\top)^\top$
- What is the velocity $\dot{\mathbf{p}}$ at point \mathbf{p} ?

$$\begin{aligned} \mathbf{p}(t) &= R(t)\mathbf{p}(0) + \mathbf{t}(t) \\ &= \exp([\boldsymbol{\omega}]_{\times}t)\mathbf{p}(0) + \mathbf{v}t \end{aligned}$$

$$\begin{aligned} \Rightarrow \dot{\mathbf{p}}(t) &= \exp([\boldsymbol{\omega}]_{\times}t)[\boldsymbol{\omega}]_{\times}\mathbf{p}(0) + \mathbf{v} \\ \dot{\mathbf{p}}(0) &= [\boldsymbol{\omega}]_{\times}\mathbf{p}(0) + \mathbf{v} \end{aligned}$$




Visual Navigation for Flying Robots

9

Dr. Jürgen Sturm, Computer Vision Group, TUM

Recap: Perspective Projection




Visual Navigation for Flying Robots

10

Dr. Jürgen Sturm, Computer Vision Group, TUM

Recap: Perspective Projection



Visual Navigation for Flying Robots

11

Dr. Jürgen Sturm, Computer Vision Group, TUM

3D to 2D Perspective Projection

- 3D point \mathbf{p} (in the camera frame)
- 2D point \mathbf{x} (on the image plane)
- Pin-hole camera model

$$\tilde{\mathbf{x}} = \lambda \bar{\mathbf{x}} = \mathbf{p}$$

- Remember, $\tilde{\mathbf{x}}$ is homogeneous, need to normalize

$$\tilde{\mathbf{x}} = \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \end{pmatrix} \Rightarrow \mathbf{x} = \begin{pmatrix} \tilde{x}/\tilde{z} \\ \tilde{y}/\tilde{z} \end{pmatrix}$$


Visual Navigation for Flying Robots

12

Dr. Jürgen Sturm, Computer Vision Group, TUM

Camera Intrinsics

- So far, 2D point is given in meters on image plane
- But: we want 2D point be measured in pixels (as the sensor does)



13

Camera Intrinsics

- Need to apply some scaling/offset

$$\tilde{\mathbf{x}} = \underbrace{\begin{pmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix}}_{\text{intrinsics } K} \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}}_{\text{projection}} \tilde{\mathbf{p}}$$

- Focal length f_x, f_y
- Camera center c_x, c_y
- Skew s

14

Image Plane

- Pixel coordinates $\mathbf{x} \in \Omega$
- Image plane $\Omega \subset \mathbb{R}^2$
- Example:
 - Discrete case $\mathbf{x} \in [0, W) \times [0, H) \subset \mathbb{N}_0^2$ (default in this course)
 - Continuous case $\mathbf{x} \in [0, 1] \times [0, 1] \subset \mathbb{R}^2$

Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Functions

- We can think of an image as a function $f : \Omega \mapsto \mathbb{R}$
- $f(\mathbf{x})$ gives the intensity at position \mathbf{x}
- Color images are vector-valued functions

$$f(\mathbf{x}) = \begin{pmatrix} r(\mathbf{x}) \\ g(\mathbf{x}) \\ b(\mathbf{x}) \end{pmatrix}$$

Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Functions

- Realistically, the image function is only defined on a rectangle and has finite range

$$f : [0, W - 1] \times [0, H - 1] \mapsto [0, 1]$$

- Image can be represented as a matrix
- Alternative notations

$$F_{ij}, f(i, j), f(x, y), f(\mathbf{x}), \dots$$


$\uparrow \uparrow$
 often (row,column)
 $\uparrow \uparrow$
 often (column,row)

	j	
	\rightarrow	
i	\downarrow	
		111 115 113 111 112 111 112 111
		135 130 137 130 145 146 149 142
		163 168 188 196 206 202 206 207
		180 184 206 219 202 200 195 193
		189 193 214 216 104 99 83 77
		191 201 217 220 103 99 81 68
		195 205 216 222 113 88 69 83
		199 203 223 228 108 68 71 77

Visual Navigation for Flying Robots

17

Example



Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM

Digital Images


- Light intensity is sampled by CCD/CMOS sensor on a regular grid
- Electric charge of each cell is quantized and gamma compressed (for historical reasons)

$$V = B^{\frac{1}{\gamma}} \text{ with } \gamma = 2.2$$

- CRTs / monitors do the inverse $B = V^{\gamma}$
 - Almost all images are gamma compressed
- Double brightness results only in a 37% higher intensity value (!)

Aliasing

- High frequencies in the scene and a small fill factor on the chip can lead to (visually) displeasing effects
- Examples



Rolling Shutter

- Most CMOS sensors have a rolling shutter
- Rows are read out sequentially
- Sensitive to camera and object motion
- Can we correct for this?



Image Filtering

- We want to remove unwanted sources of variation, and keep the information relevant for whatever task we need to solve




- Example tasks: de-noising, (de-)blurring, computing derivatives, edge detection, ...

Linear Filtering

- Each output is a linear combination of all the input values

$$g(i, j) = \sum_{k, l} h(i, j, k, l) f(k, l)$$

- In matrix form



Spatially Invariant Filtering

- We are often interested in spatially invariant operations

$$g(i, j) = f * h = \sum_{k, l} h(i - k, j - l) f(k, l)$$

- Example

111	115	117	111	112	111	112	111
135	138	137	139	145	146	149	147
163	168	188	196	206	202	206	207
180	184	206	219	202	200	195	193
189	193	214	216	104	79	83	77
191	201	217	220	103	59	67	68
195	205	216	222	113	68	69	83
199	203	223	228	108	68	71	77


$$* \begin{bmatrix} -1 & 2 & -1 \\ -1 & 2 & -1 \\ -1 & 2 & -1 \end{bmatrix} = ?$$

Spatially Invariant Filtering

- We are often interested in spatially invariant operations

$$g(i, j) = f * h = \sum_{k,l} h(i - k, j - l) f(k, l)$$

- Example



Visual Navigation for Flying Robots

25

Dr. Jürgen Sturm, Computer Vision Group, TUM

Important Filters

- Impulses
- Shifts
- Blur
 - Gaussian
 - Bilateral filter
 - Motion blur
- Edges
 - Finite difference filter
 - Derivative filter
 - Oriented filters
 - Gabor filter
- ...


Visual Navigation for Flying Robots

26

Dr. Jürgen Sturm, Computer Vision Group, TUM

Impulse

$$g(i, j) = f * h = \sum_{k,l} h(i - k, j - l) f(k, l)$$




Visual Navigation for Flying Robots

27

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image shift (translation)

$$g(i, j) = f * h = \sum_{k,l} h(i - k, j - l) f(k, l)$$




Visual Navigation for Flying Robots

28

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image rotation

$$g(i, j) = f * h = \sum_{k,l} h(i - k, j - l) f(k, l)$$



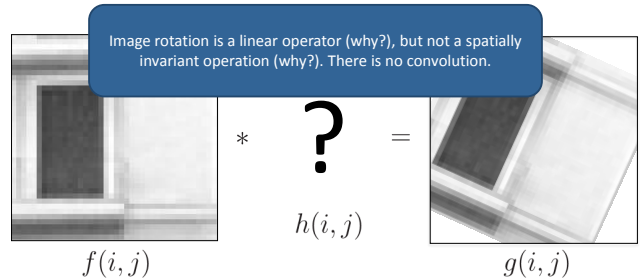
Visual Navigation for Flying Robots

29

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image rotation

$$g(i, j) = f * h = \sum_{k,l} h(i - k, j - l) f(k, l)$$




Visual Navigation for Flying Robots

30

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rectangular Filter

$$g(i, j) = f * h = \sum_{k, l} h(i - k, j - l) f(k, l)$$




Visual Navigation for Flying Robots

31

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rectangular Filter

$$g(i, j) = f * h = \sum_{k, l} h(i - k, j - l) f(k, l)$$




Visual Navigation for Flying Robots

32

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rectangular Filter

$$g(i, j) = f * h = \sum_{k, l} h(i - k, j - l) f(k, l)$$



Visual Navigation for Flying Robots

33


Dr. Jürgen Sturm, Computer Vision Group, TUM

Gaussian Blur

- Gaussian distribution

$$g_{\sigma}(i, j) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{i^2 + j^2}{2\sigma^2}\right)$$

- Example of resulting kernel




Visual Navigation for Flying Robots

34

Dr. Jürgen Sturm, Computer Vision Group, TUM



Gaussian Blur



Visual Navigation for Flying Robots

Image Gradient

- The image gradient $\nabla f = \left(\frac{\partial f}{\partial x} \quad \frac{\partial f}{\partial y}\right)^T$ points in the direction of increasing intensity (steepest ascend)



Visual Navigation for Flying Robots

36

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Gradient

- The image gradient $\nabla f = \left(\frac{\partial f}{\partial x} \quad \frac{\partial f}{\partial y} \right)^\top$ points in the direction of increasing intensity (steepest ascend)



$$\nabla f = \left(\frac{\partial f}{\partial x}, 0 \right)^\top$$



$$\nabla f = \left(0, \frac{\partial f}{\partial y} \right)^\top$$



$$\nabla f = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right)^\top$$

Visual Navigation for Flying Robots

37

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Gradient

- Gradient direction (related to edge orientation)

$$\theta = \text{atan2} \left(\frac{\partial f}{\partial y}, \frac{\partial f}{\partial x} \right)$$

- Gradient magnitude (edge strength)

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x} \right)^2 + \left(\frac{\partial f}{\partial y} \right)^2}$$

Visual Navigation for Flying Robots

38

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Gradient

How can we differentiate a digital image $f(x, y)$?

- Option 1: Reconstruct a continuous image, then take gradient
- Option 2: Take discrete derivative (finite difference filter)
- Option 3: Convolve with derived Gaussian (derivative filter)

Visual Navigation for Flying Robots

39


Dr. Jürgen Sturm, Computer Vision Group, TUM

Finite difference

- First-order central difference

$$\frac{\partial f}{\partial x}(x, y) \approx \frac{f(x+1, y) - f(x-1, y)}{2}$$

- Corresponding convolution kernel: $\begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$



Visual Navigation for Flying Robots

40


Dr. Jürgen Sturm, Computer Vision Group, TUM

Finite difference

- First-order central difference (half pixel)

$$\frac{\partial f}{\partial x}(x, y) \approx f(x+0.5, y) - f(x-0.5, y)$$

- Corresponding convolution kernel: $\begin{bmatrix} -1 & 1 \end{bmatrix}$



Visual Navigation for Flying Robots

41

Dr. Jürgen Sturm, Computer Vision Group, TUM

Second-order Derivative

- Differentiate again to get second-order central difference

$$\frac{\partial^2 f}{\partial x^2} \approx f(x+1) - 2f(x) + f(x-1)$$

- Corresponding convolution kernel: $\begin{bmatrix} 1 & -2 & 1 \end{bmatrix}$


Visual Navigation for Flying Robots

42

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example


$$g(i, j) = f * h = \sum_{k,l} h(i - k, j - l) f(k, l)$$



Visual Navigation for Flying Robots

Example

$$g(i, j) = f * h = \sum_{k,l} h(i - k, j - l) f(k, l)$$



Visual Navigation for Flying Robots

(Dense) Motion Estimation

- 2D motion





- 3D motion



Visual Navigation for Flying Robots

Problem Statement

- Given:** two camera images f_0, f_1
- Goal:** estimate the camera motion \mathbf{u}




- For the moment, let's assume that the camera only moves in the xy-plane, i.e., $\mathbf{u} = (u \ v)^\top$
- Extension to 3D follows

Visual Navigation for Flying Robots

General Idea

- Define an error metric $E(\mathbf{u})$ that defines how well the two images match given a motion vector
- Find the motion vector with the lowest error

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} E(\mathbf{u})$$



Visual Navigation for Flying Robots

Error Metrics for Image Comparison

- Sum of Squared Differences (SSD)

$$E_{SSD}(\mathbf{u}) = \sum_i (f_1(\mathbf{x}_i + \mathbf{u}) - f_0(\mathbf{x}_i))^2 = \sum_i e_i^2$$

with displacement $\mathbf{u} = (u \ v)^\top$
and residual errors $e_i = f_1(\mathbf{x}_i + \mathbf{u}) - f_0(\mathbf{x}_i)$

Visual Navigation for Flying Robots

Robust Error Metrics

- SSD metric is sensitive to outliers
- Solution: apply a (more) robust error metric

$$E_{\text{SRD}}(\mathbf{u}) = \sum_i \rho(f_1(\mathbf{x}_i + \mathbf{u}) - f_0(\mathbf{x}_i)) = \sum_i \rho(e_i)$$

Visual Navigation for Flying Robots

49

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robust Error Metrics

- Sum of Absolute Differences

$$\rho_{\text{SAD}}(e) = |e|$$

- Sum of truncated errors

$$\rho_{\text{trunc}}(e) = \begin{cases} e^2 & \text{if } |e| < b \\ b^2 & \text{otherwise} \end{cases}$$

- Geman-McClure function (Huber norm)


$$\rho_{\text{huber}}(e) = \frac{e^2}{1 + e^2/b^2}$$

Visual Navigation for Flying Robots

50

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robust Error Metrics



Visual Navigation for Flying Robots

51

Dr. Jürgen Sturm, Computer Vision Group, TUM

Windowed SSD

- Images (and image patches) have finite size
- Standard SSD has a bias towards smaller overlaps (less error terms)
- Solution: divide by the overlap area
- Root mean square error

$$E_{\text{RMS}}(\mathbf{u}) = \sqrt{E_{\text{SSD}}/A}$$

Visual Navigation for Flying Robots

52

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exposure Differences

- Images might be taken with different exposure (auto shutter, white balance, ...)
- Bias and gain model

$$f_1(\mathbf{x} + \mathbf{u}) = (1 + \alpha)f_0(\mathbf{x}) + \beta$$

- With SSD we get

$$\begin{aligned} E_{\text{BG}}(\mathbf{u}) &= \sum_i (f_1(\mathbf{x}_i + \mathbf{u}) - (1 + \alpha)f_0(\mathbf{x}_i) + \beta)^2 \\ &= \sum_i \alpha f_0(\mathbf{x}) + \beta - e_i^2 \end{aligned}$$

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Cross-Correlation

- Maximize the product (instead of minimizing the differences)

$$E_{\text{CC}}(\mathbf{u}) = - \sum_i f_0(\mathbf{x}_i) f_1(\mathbf{x}_i + \mathbf{u})$$

- Normalized cross-correlation (between -1..1)

$$\begin{aligned} E_{\text{NCC}}(\mathbf{u}) &= \\ &= - \sum_i \frac{(f_0(\mathbf{x}_i) - \text{mean} f_0)(f_1(\mathbf{x}_i + \mathbf{u}) - \text{mean} f_1)}{\sqrt{\text{var} f_0 \text{var} f_1}} \end{aligned}$$

Visual Navigation for Flying Robots

54

Dr. Jürgen Sturm, Computer Vision Group, TUM

General Idea

1. Define an error metric $E(\mathbf{u})$ that defines how well the two images match given a motion vector
2. Find the motion vector with the lowest error

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} E(\mathbf{u})$$



Visual Navigation for Flying Robots

55

Dr. Jürgen Sturm, Computer Vision Group, TUM

Finding the minimum

- Full search (e.g., ± 16 pixels)
- Gradient descent
- Hierarchical motion estimation

Visual Navigation for Flying Robots


56

Dr. Jürgen Sturm, Computer Vision Group, TUM

Hierarchical motion estimation

- Construct image pyramid

$$f_k^{(l+1)}(\mathbf{x}_i) \leftarrow f_k^{(l)}(2\mathbf{x}_i)$$



- Estimate motion on coarse level
- Use as initialization for next finer level

$$\hat{\mathbf{u}}^{(l-1)} \leftarrow 2\mathbf{u}^{(l)}$$

Visual Navigation for Flying Robots

57

Dr. Jürgen Sturm, Computer Vision Group, TUM

Gradient Descent

- Perform gradient descent on the SSD energy function (Lucas and Kanade, 1981)
- Taylor expansion of energy function

$$E_{\text{LK-SSD}}(\mathbf{u} + \Delta\mathbf{u}) = \sum_i (f_1(\mathbf{x}_i + \mathbf{u} + \Delta\mathbf{u}) - f_0(\mathbf{x}_i))^2$$

$$\approx \sum_i (f_1(\mathbf{x}_i + \mathbf{u}) + J_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} - f_0(\mathbf{x}_i))^2$$

$$= \sum_i (J_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} + e_i)^2$$

$$\text{with } J_1(\mathbf{x}_i + \mathbf{u}) = \nabla f_1(\mathbf{x}_i + \mathbf{u}) = \left(\frac{\partial f_1}{\partial x}, \frac{\partial f_1}{\partial y} \right) (\mathbf{x}_i + \mathbf{u})$$

Visual Navigation for Flying Robots

58

Dr. Jürgen Sturm, Computer Vision Group, TUM

Least Squares Problem

- Goal: Minimize

$$E(\mathbf{u} + \Delta\mathbf{u}) = \sum_i (J_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} + e_i)^2$$

- Solution: Compute derivative (and set to zero)

$$\frac{\partial E(\mathbf{u} + \Delta\mathbf{u})}{\partial \Delta\mathbf{u}} = 2A\Delta\mathbf{u} + 2\mathbf{b}$$

$$\text{with } A = \sum_i J_1^\top(\mathbf{x}_i + \mathbf{u})J_1(\mathbf{x}_i + \mathbf{u})$$

$$\text{and } \mathbf{b} = \sum_i e_i J_1^\top(\mathbf{x}_i + \mathbf{u})$$

Visual Navigation for Flying Robots

59

Dr. Jürgen Sturm, Computer Vision Group, TUM

Least Squares Problem

1. Compute A,b from image gradients using

$$A = \begin{pmatrix} \sum f_x^2 & \sum f_x f_y \\ \sum f_x f_y & \sum f_y^2 \end{pmatrix} \quad \mathbf{b} = \begin{pmatrix} \sum f_x f_t \\ \sum f_y f_t \end{pmatrix}$$

$$\text{with } f_x = \frac{\partial f_1(\mathbf{x})}{\partial x}, f_y = \frac{\partial f_1(\mathbf{x})}{\partial y}$$

$$\text{and } f_t = \frac{\partial f_t(\mathbf{x})}{\partial t} [\approx f_1(\mathbf{x}) - f_0(\mathbf{x})]$$

2. Solve $A\Delta\mathbf{u} = -\mathbf{b}$

$$\Rightarrow \Delta\mathbf{u} = -A^{-1}\mathbf{b}$$

All of these computation
are super fast!

Visual Navigation for Flying Robots

60

Dr. Jürgen Sturm, Computer Vision Group, TUM

Covariance of the Estimated Motion

- Assuming (small) Gaussian noise in the images

$$f_{\text{obs}}(\mathbf{x}_i) = f_{\text{true}}(\mathbf{x}_i) + \epsilon_i$$

with $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$

- ... results in uncertainty in the motion estimate with covariance (e.g., useful for Kalman filter)

$$\Sigma_u = \sigma^2 A^{-1}$$

Visual Navigation for Flying Robots



61

Dr. Jürgen Sturm, Computer Vision Group, TUM

Optical Computer Mouse (since 1999)

- E.g., ADNS3080 from Agilent Technologies, 2005

- 6400 fps
- 30x30 pixels
- 4 USD



Visual Navigation for Flying Robots

62

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Patches

- Sometimes we are interested of the motion of a small image patches
- Problem:** some patches are easier to track than others
- What patches are easy/difficult to track?
- How can we recognize “good” patches?


Visual Navigation for Flying Robots

63

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Patches

- Sometimes we are interested of the motion of a small image patches
- Problem:** some patches are easier to track than others



Visual Navigation for Flying Robots

64

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Let's look at the shape of the energy functional



Visual Navigation for Flying Robots

65

Dr. Jürgen Sturm, Computer Vision Group, TUM

Corner Detection

$$A = \begin{pmatrix} \sum f_x^2 & \sum f_x f_y \\ \sum f_x f_y & \sum f_y^2 \end{pmatrix}$$

- Idea:** Inspect eigenvalues λ_1, λ_2 of Hessian A
 - λ_1, λ_2 small \rightarrow no point of interest
 - λ_1 large, λ_2 small \rightarrow edge
 - λ_1, λ_2 large \rightarrow corner
- Harris detector (does not need eigenvalues)

$$\lambda_1 \lambda_2 > \kappa (\lambda_1 + \lambda_2)^2 \Leftrightarrow \det(A) > \kappa \text{trace}^2(A)$$
- Shi-Tomasi (or Kanade-Lucas) $\min(\lambda_1, \lambda_2) > \kappa$

Visual Navigation for Flying Robots

66

Dr. Jürgen Sturm, Computer Vision Group, TUM

Corner Detection

1. For all pixels, computer corner strength
2. Non-maximal suppression (E.g., sort by strength, strong corner suppresses weaker corners in circle of radius r)



strongest responses



non-maximal suppression

Other Detectors

- Förstner detector (localize corner with sub-pixel accuracy)
- FAST corners (learn decision tree, minimize number of tests → super fast)
- Difference of Gaussians / DoG (scale-invariant detector)
- ...

Kanade-Lucas-Tomasi (KLT) Tracker

- Algorithm
 1. Find (Shi-Tomasi) corners in first frame and initialize tracks
 2. Track from frame to frame
 3. Delete track if error exceeds threshold
 4. Initialize additional tracks when necessary
 5. Repeat step 2-4
- KLT tracker is highly efficient (real-time on CPU) but provides only sparse motion vectors
- Dense optical flow methods require GPU


Example



3D Motion Estimation


(How) Can we recover the camera motion from the estimated flow field?

- Research paper: Grabe et al., ICRA 2012
<http://www9.in.tum.de/~sturmju/dirs/icra2012/data/papers/2025.pdf>




Approach [Grabe et al., ICRA'12]

- Compute optical flow
- Estimate homography between images
- Extract angular and (scaled) linear velocity
- Additionally employ information from IMU




Assumptions

1. The quadcopter moves slowly relative to the sampling rate
→ limited search radius



2. The environment is planar with normal N
→ image transformation is a homography



Visual Navigation for Flying Robots

73

Dr. Jürgen Sturm, Computer Vision Group, TUM

Apparent Velocity of a Point

- Stationary 3D point feature, given in camera frame

$$\mathbf{p} \in \mathbb{R}^3$$

- Moving camera with twist

$$\xi = (\mathbf{v}^\top, \boldsymbol{\omega}^\top)^\top \in \mathbb{R}^6$$

- Apparent velocity** of the point in camera frame

$$\dot{\mathbf{p}} = [\boldsymbol{\omega}]_\times \mathbf{p} + \mathbf{v}$$

Visual Navigation for Flying Robots

74


Dr. Jürgen Sturm, Computer Vision Group, TUM

Continuous Homography Matrix

- Assumption: All feature points are located on a plane

$$N^\top \mathbf{p} = d$$

with plane normal $N \in \mathbb{R}^3$
and distance $d \in \mathbb{R}$



Visual Navigation for Flying Robots

75

Dr. Jürgen Sturm, Computer Vision Group, TUM

Continuous Homography Matrix

- Rewrite this to $\frac{1}{d}N^\top \mathbf{p} = 1$ and plug it into the equation for the apparent velocity, we obtain

$$\dot{\mathbf{p}} = [\boldsymbol{\omega}]_\times \mathbf{p} + \mathbf{v} \frac{1}{d} N^\top \mathbf{p} = \underbrace{\left([\boldsymbol{\omega}]_\times + \mathbf{v} \frac{1}{d} N^\top \right)}_{H \in \mathbb{R}^{3 \times 3}} \mathbf{p} = H \mathbf{p}$$

- H is called the **continuous homography matrix**
- Note: H contains both the linear/angular velocity $(\mathbf{v}, \boldsymbol{\omega})$ and the scene structure (N, d)

Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM


Continuous Homography Constraint

- The camera observes point $\mathbf{p} \in \mathbb{R}^3$ at pixel $\mathbf{x} \in \mathbb{R}^2$ (assuming $K = I$ for simplicity)

$$\tilde{\mathbf{x}} = \lambda \bar{\mathbf{x}} = \mathbf{p}$$

- The KLT tracker estimates the motion \mathbf{u} of the feature track in the image

- Constraint: $\mathbf{u} = \dot{\tilde{\mathbf{x}}}$



Visual Navigation for Flying Robots

77

Dr. Jürgen Sturm, Computer Vision Group, TUM

Continuous Homography Constraint

- We now have

1. $\dot{\mathbf{p}} = H \mathbf{p}$

2. $\dot{\mathbf{p}} = \lambda \dot{\bar{\mathbf{x}}} + \dot{\lambda} \bar{\mathbf{x}}$ (time derivative of $\mathbf{p} = \lambda \bar{\mathbf{x}}$ and the optical flow constraint $\mathbf{u} = \dot{\tilde{\mathbf{x}}}$)

- Let's combine these two formulas...

Visual Navigation for Flying Robots

78

Dr. Jürgen Sturm, Computer Vision Group, TUM

Continuous Homography Constraint

- Combining these formulas gives us

$$\dot{\lambda}\bar{x} + \lambda\bar{u} = H\mathbf{p}$$

$$\lambda\bar{u} = H\mathbf{p} - \dot{\lambda}\bar{x}$$

$$\bar{u} = H\bar{x} - \frac{\dot{\lambda}}{\lambda}\bar{x}$$

- Multiply both sides with $[\bar{x}]_{\times}$ gives us

$$[\bar{x}]_{\times}\bar{u} = [\bar{x}]_{\times}H\bar{x} - \underbrace{[\bar{x}]_{\times}\frac{\dot{\lambda}}{\lambda}\bar{x}}_{=0}$$

$$\Rightarrow [\bar{x}]_{\times}\bar{u} = [\bar{x}]_{\times}H\bar{x}$$

Approach

- Result:** For all observed motions in the image, the continuous homography constraint holds

$$[\bar{x}]_{\times}\bar{u} = [\bar{x}]_{\times}H\bar{x}$$

- How can we use this to estimate the camera motion?!

Approach

- Result:** For all observed motions in the image, the continuous homography constraint holds

$$[\bar{x}]_{\times}\bar{u} = [\bar{x}]_{\times}H\bar{x}$$

- How can we use this to estimate the camera motion?

- Estimate H from at least 4 feature tracks
- Recover $(\mathbf{v}, \boldsymbol{\omega})$ and (N, d) from H

Remember: $H = [\boldsymbol{\omega}]_{\times} + \mathbf{v}_d^1 N^T$

Step 1: Estimate H

- Continuous homography constraint

$$[\bar{x}]_{\times}H\bar{x} = [\bar{x}]_{\times}\bar{u}$$

- Stack matrix H as a vector $\mathbf{h} \in \mathbb{R}^9$ and rewrite

$$M^T \mathbf{h} = [\bar{x}]_{\times}\bar{u}$$

→ Linear system of equations

- For several feature tracks

$$\begin{pmatrix} M_1^T \\ M_2^T \\ \vdots \end{pmatrix} \mathbf{h} = \begin{pmatrix} [\bar{x}]_{\times}\bar{u}_1^T \\ [\bar{x}]_{\times}\bar{u}_2^T \\ \vdots \end{pmatrix}$$

Step 1: Estimate H

- Linear set of equations

$$\underbrace{\begin{pmatrix} M_1^T \\ M_2^T \\ \vdots \end{pmatrix}}_A \mathbf{h} = \underbrace{\begin{pmatrix} [\bar{x}]_{\times}\bar{u}_1^T \\ [\bar{x}]_{\times}\bar{u}_2^T \\ \vdots \end{pmatrix}}_b$$

- Solve for \mathbf{h} using least squares

$$A\mathbf{h} = \mathbf{b}$$

$$\Rightarrow \mathbf{h} = (A^T A)^{-1} A^T \mathbf{b}$$

Step 2: Recover camera motion

Grabe et al. investigated three alternatives:

- Recover $(\boldsymbol{\omega}, \frac{\mathbf{v}}{d}, N)$ from $H = [\boldsymbol{\omega}]_{\times} + \mathbf{v}_d^1 N^T$ using the 8-point algorithm (not yet explained)
- Use angular velocity $\boldsymbol{\omega}$ from IMU to de-rotate observed feature tracks beforehand, then:

$$H = \mathbf{v}_d^1 N^T$$

- Additionally use gravity vector from IMU as plane normal $N = N_{\text{IMU}}$, then

$$\frac{\mathbf{v}}{d} = H(N^T N)^{-1}$$

Evaluation

- Comparison of estimated velocities with ground truth from motion capture system

Algorithm	Norm error	Std. deviation
Pure vision	0.134 $\frac{m}{s}$	0.094 $\frac{m}{s}$
Ang. vel. known	0.117 $\frac{m}{s}$	0.093 $\frac{m}{s}$
Normal known	0.113 $\frac{m}{s}$	0.088 $\frac{m}{s}$

- Comparison of actual velocity with desired velocity (closed-loop control)

Algorithm	Norm error	Std. deviation
Pure vision	0.084 $\frac{m}{s}$	0.139 $\frac{m}{s}$
Ang. vel. known	0.039 $\frac{m}{s}$	0.042 $\frac{m}{s}$
Normal known	0.028 $\frac{m}{s}$	0.031 $\frac{m}{s}$

Visual Navigation for Flying Robots

85

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Velocity Control

- All computations are carried out on-board (18fps)



[Grabe et al., ICRA '12]

Visual Navigation for Flying Robots

86

Dr. Jürgen Sturm, Computer Vision Group, TUM

Landing on a Moving Platform

- Similar approach, but with offboard computation



[Herissé et al., T-RO '12]


Visual Navigation for Flying Robots

87

Dr. Jürgen Sturm, Computer Vision Group, TUM

Commercial Solutions

- Helicommand 3D from Robbe
2(?) cameras, IMU, air pressure sensor, 450 EUR
- Parrot Mainboard + Navigation board
1 camera, IMU, ultrasound sensor, 210 USD



Visual Navigation for Flying Robots

88

Dr. Jürgen Sturm, Computer Vision Group, TUM

Lessons Learned Today

- How to estimate the translational motion from camera images
- Which image patches are easier to track than others
- How to estimate 3D motion from multiple feature tracks (and IMU data)

Visual Navigation for Flying Robots

89

Dr. Jürgen Sturm, Computer Vision Group, TUM

A Few Ideas for Your Mini-Project

- Person following (colored shirt or wearing a marker)
- Flying camera for taking group pictures (possibly using the OpenCV face detector)
- Fly through a hula hoop (brightly colored, white background)
- Navigate through a door (brightly colored)
- Navigate from one room to another (using ground markers)
- Avoid obstacles using optical flow
- Landing on a moving platform
- Your own idea here – be creative!**
- ...

Visual Navigation for Flying Robots

90

Dr. Jürgen Sturm, Computer Vision Group, TUM

Joggbot

- Follows a person wearing a visual marker



[<http://exertiongameslab.org/projects/joggbot>]

Visual Navigation for Flying Robots

Simultaneous Localization and Mapping (SLAM)

Dr. Jürgen Sturm

Organization: Exam Dates

- Registration deadline: June 30
- Course ends: July 19
- Examination dates: August 9+14 (Thu+Tue)
 - Oral team exam
 - Sign up for a time slot starting from now
 - List placed on blackboard in front of our secretary

VISNAV Oral Team Exam

Date and Time	Student Name	Student Name	Student Name
Tue, Aug. 9, 10am			
Tue, Aug. 9, 11am			
Tue, Aug. 9, 2pm			
Tue, Aug. 9, 3pm			
Tue, Aug. 9, 4pm			
Thu, Aug. 14, 10am			
Thu, Aug. 14, 11am			
Thu, Aug. 14, 2pm			
Thu, Aug. 14, 3pm			
Thu, Aug. 14, 4pm			

The SLAM Problem

SLAM is the process by which a robot **builds a map** of the environment and, at the same time, uses the map to **compute its location**

- Localization: inferring location given a map
- Mapping: inferring a map given a location

The SLAM Problem

Given:

- The robot's controls $\mathbf{u}_{1:t} = \langle \mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_t \rangle$
- (Relative) observations $\mathbf{z}_{1:t} = \langle \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_t \rangle$

Wanted:

- Map of features $\mathbf{m} = \langle \mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_k \rangle$
- Trajectory of the robot $\mathbf{x}_{1:t} = \langle \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t \rangle$


SLAM Applications

SLAM is central to a range of indoor, outdoor, in-air and underwater applications for both unmanned and autonomous vehicles.

Examples


- At home: vacuum cleaner, lawn mower
- Air: inspection, transportation, surveillance
- Underwater: reef/environmental monitoring
- Underground: search and rescue
- Space: terrain mapping, navigation

SLAM with Ceiling Camera (Samsung Hauzen RE70V, 2008)



Visual Navigation for Flying Robots 7 Dr. Jürgen Sturm, Computer Vision Group, TUM

SLAM with Laser + Line camera (Neato XV 11, 2010)



Visual Navigation for Flying Robots 8 Dr. Jürgen Sturm, Computer Vision Group, TUM

Localization, Path planning, Coverage (Neato XV11, \$300)



Visual Navigation for Flying Robots 9 Dr. Jürgen Sturm, Computer Vision Group, TUM

SLAM vs. SfM

- In Robotics: Simultaneous Localization and Mapping (SLAM)
 - Laser scanner, ultrasound, monocular/stereo camera
 - Typically in combination with an odometry sensor
 - Typically pre-calibrated sensors
- In Computer Vision: Structure from Motion (SfM), sometimes: Structure and Motion
 - Monocular/stereo camera
 - Sometimes uncalibrated sensors (e.g., Flickr images)

Visual Navigation for Flying Robots 10 Dr. Jürgen Sturm, Computer Vision Group, TUM


Agenda for Today

- **This week:** focus on monocular vision
 - Feature detection, descriptors and matching
 - Epipolar geometry
 - Robust estimation (RANSAC)
 - Examples (PTAM, Photo Tourism)
- **Next week:** focus on optimization (bundle adjustment), stereo cameras, Kinect
- **In two weeks:** map representations, mapping and (dense) 3D reconstruction

Visual Navigation for Flying Robots 11 Dr. Jürgen Sturm, Computer Vision Group, TUM

How Do We Build a Panorama Map?


- We need to match (align) images
- Global methods sensitive to occlusion, lighting, parallax effects
- How would you do it by eye?



Visual Navigation for Flying Robots 12 Dr. Jürgen Sturm, Computer Vision Group, TUM

Matching with Features

- Detect features in both images




Visual Navigation for Flying Robots

13

Dr. Jürgen Sturm, Computer Vision Group, TUM

Matching with Features

- Detect features in both images
- Find corresponding pairs



Visual Navigation for Flying Robots

14

Dr. Jürgen Sturm, Computer Vision Group, TUM

Matching with Features

- Detect features in both images
- Find corresponding pairs
- Use these pairs to align images




Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

Matching with Features

- Problem 1:
We need to detect the **same** point **independently** in both images



no chance to match!

→ We need a reliable detector

Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Matching with Features

- Problem 2:
For each point correctly recognize the corresponding one



→ We need a reliable and distinctive descriptor

Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

Ideal Feature Detector

- Always finds the same point on an object, regardless of changes to the image
- Insensitive (invariant) to changes in:
 - Scale
 - Lightning
 - Perspective imaging
 - Partial occlusion

Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM


Harris Detector

- Rotation invariance?



Harris Detector

- Rotation invariance?



- Remember from last week

$$A = \begin{pmatrix} \sum f_x^2 & \sum f_x f_y \\ \sum f_x f_y & \sum f_y^2 \end{pmatrix} \quad R = \lambda_1 \lambda_2 - \kappa (\lambda_1 + \lambda_2)^2$$

Harris Detector

- Rotation invariance



- Remember from last week

$$A = \begin{pmatrix} \sum f_x^2 & \sum f_x f_y \\ \sum f_x f_y & \sum f_y^2 \end{pmatrix} \quad R = \lambda_1 \lambda_2 - \kappa (\lambda_1 + \lambda_2)^2$$

- Ellipse rotates but its shape (i.e. eigenvalues) remains the same

→ Corner response R is invariant to rotation

Harris Detector


- Invariance to intensity change?

Harris Detector

- Partial invariance to additive and multiplicative intensity changes

- Only derivatives are used → invariance to intensity shift $I \rightarrow I + b$
- Intensity scale $I \rightarrow aI$:

Because of fixed intensity threshold on local maxima, only partial invariance




Harris Detector

- Invariant to scaling?

Harris Detector

- Not invariant to image scale




Visual Navigation for Flying Robots

25

Dr. Jürgen Sturm, Computer Vision Group, TUM

Difference Of Gaussians (DoG)

- Alternative corner detector that is additionally invariant to scale change
- Approach:
 - Run linear filter (diff. of two Gaussians, $\sigma_1 = 2\sigma_2$)
 - Do this at different scales
 - Search for a maximum both in space and scale

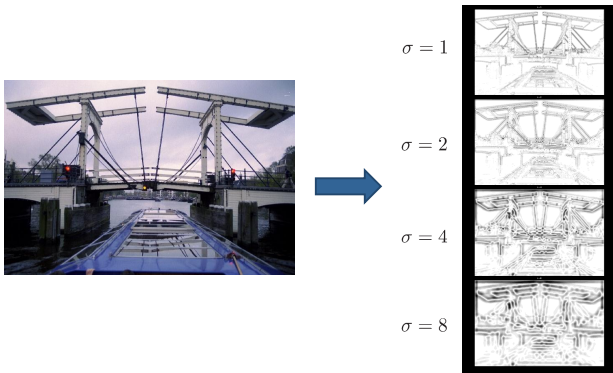


Visual Navigation for Flying Robots

26

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Difference of Gaussians




Visual Navigation for Flying Robots

27

Dr. Jürgen Sturm, Computer Vision Group, TUM

SIFT Detector

- Search for local maximum in space and scale
- Corner detections are invariant to scale change



Visual Navigation for Flying Robots

28

Dr. Jürgen Sturm, Computer Vision Group, TUM

SIFT Detector

- Detect maxima in scale-space
- Non-maximum suppression
- Eliminate edge points (check ratio of eigenvalues)
- For each maximum, fit quadratic function and compute center at sub-pixel accuracy

Visual Navigation for Flying Robots

29

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Input image 233x189 pixel
- 832 candidates DoG minima/maxima (visualization indicate scale, orient., location)
- 536 keypoints remain after thresholding on minimum contrast and principal curvature




Visual Navigation for Flying Robots

30

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feature Matching

- Now, we know how to find **repeatable** corners
- Next question: How can we match them?



Visual Navigation for Flying Robots

31


Dr. Jürgen Sturm, Computer Vision Group, TUM

Template Convolution

- Extract a small as a template



- Convolve image with this template



Visual Navigation for Flying Robots

32

Dr. Jürgen Sturm, Computer Vision Group, TUM

Template Convolution

Invariances

- Scaling: No
- Rotation: No (maybe rotate template?)
- Illumination: No (use bias/gain model?)
- Perspective projection: Not really


Visual Navigation for Flying Robots

33

Dr. Jürgen Sturm, Computer Vision Group, TUM

Scale Invariant Feature Transform (SIFT)

- Lowe, 2004: Transform patches into a canonical form that is invariant to translation, rotation, scale, and other imaging parameters



Visual Navigation for Flying Robots

34

Dr. Jürgen Sturm, Computer Vision Group, TUM

Scale Invariant Feature Transform (SIFT)

Approach

- Find SIFT corners (position + scale)
- Find dominant orientation and de-rotate patch
- Extract SIFT descriptor (histograms over gradient directions)


Visual Navigation for Flying Robots

35

Dr. Jürgen Sturm, Computer Vision Group, TUM

Select Dominant Orientation

- Create a histogram of local gradient directions computed at selected scale (36 bins)
- Assign canonical orientation at peak of smoothed histogram
- Each key now specifies stable 2D coordinates (x, y, scale, orientation)



Visual Navigation for Flying Robots

36

Dr. Jürgen Sturm, Computer Vision Group, TUM

SIFT Descriptor

- Compute image gradients over 16x16 window (green), weight with Gaussian kernel (blue)
- Create 4x4 arrays of orientation histograms, each consisting of 8 bins
- In total, SIFT descriptor has 128 dimensions




Image gradients

37

Keypoint descriptor

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feature Matching

Given features in I_1 , how to find best match in I_2 ?

- Define distance function that compares two features
- Test all the features in I_2 , find the one with the minimal distance

Visual Navigation for Flying Robots

38

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feature Distance

How to define the difference between features?

- Simple approach is Euclidean distance (or SSD)

$$d(\mathbf{d}_1, \mathbf{d}_2) = \|\mathbf{d}_1 - \mathbf{d}_2\|$$

Visual Navigation for Flying Robots

39

Dr. Jürgen Sturm, Computer Vision Group, TUM


Feature Distance

How to define the difference between features?

- Simple approach is Euclidean distance (or SSD)

$$d(\mathbf{d}_1, \mathbf{d}_2) = \|\mathbf{d}_1 - \mathbf{d}_2\|$$

- Problem: can give good scores to ambiguous (bad) matches



Visual Navigation for Flying Robots


40

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feature Distance

How to define the difference between features?

- Better approach $d(\mathbf{d}_1, \mathbf{d}_2) = \|\mathbf{d}_1 - \mathbf{d}_2\| / \|\mathbf{d}_1 - \mathbf{d}'_2\|$
with \mathbf{d}_2 best matching feature from I_2
 \mathbf{d}'_2 second best matching feature from I_2
- Gives small values for ambiguous matches



Visual Navigation for Flying Robots


41

Dr. Jürgen Sturm, Computer Vision Group, TUM

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$



Visual Navigation for Flying Robots


42

Dr. Jürgen Sturm, Computer Vision Group, TUM

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)



Visual Navigation for Flying Robots


43

Dr. Jürgen Sturm, Computer Vision Group, TUM

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)
 - Localize query in tree
 - Search nearby leaves until nearest neighbor is guaranteed found



Visual Navigation for Flying Robots

44


Dr. Jürgen Sturm, Computer Vision Group, TUM

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)

- Localize query in tree
- Search nearby leaves until nearest neighbor is guaranteed found



Visual Navigation for Flying Robots

45


Dr. Jürgen Sturm, Computer Vision Group, TUM

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)

- Localize query in tree
- Search nearby leaves until nearest neighbor is guaranteed found



Visual Navigation for Flying Robots

46


Dr. Jürgen Sturm, Computer Vision Group, TUM

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)

- Localize query in tree
- Search nearby leaves until nearest neighbor is guaranteed found
- Best-bin-first: use priority queue for unchecked leafs



Visual Navigation for Flying Robots


47

Dr. Jürgen Sturm, Computer Vision Group, TUM

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)
- Approximate search
 - Locality sensitive hashing
 - Approximate nearest neighbor




Visual Navigation for Flying Robots

48

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)
- Approximate search
 - Locality sensitive hashing
 - Approximate nearest neighbor



Visual Navigation for Flying Robots

49

Efficient Matching

For feature matching, we need to answer a large number of **nearest neighbor queries**

- Exhaustive search $O(n^2)$
- Indexing (k-d tree)
- Approximate search
- Vocabulary trees

Visual Navigation for Flying Robots

50

Other Descriptors

- SIFT (Scale Invariant Feature Transform) [Lowe, 2004]
- SURF (Speeded Up Robust Feature) [Bay et al., 2008]
- BRIEF (Binary robust independent elementary features) [Calonder et al., 2010]
- ORB (Oriented FAST and Rotated Brief) [Rublee et al, 2011]
- ...

Visual Navigation for Flying Robots


51

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: RGB-D SLAM

[Engelhard et al., 2011; Endres et al. 2012]

- Feature descriptor: SURF
- Feature matching: FLANN (approximate nearest neighbor)



I_1

I_2

Visual Navigation for Flying Robots

52

Dr. Jürgen Sturm, Computer Vision Group, TUM

Structure From Motion (SfM)

- Now we can compute point correspondences
- What can we use them for?

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Four Important SfM Problems

- **Camera calibration**
Known 3D points, observe corresponding 2D points, compute camera pose
- **Point triangulation**
Known camera poses, observe 2D point correspondences, compute 3D point
- **Motion estimation (epipolar geometry)**
Observe 2D point correspondences, compute camera pose (up to scale)
- **Bundle adjustment (next week!)**
Observe 2D point correspondences, compute camera pose and 3D points (up to scale)

Visual Navigation for Flying Robots

54

Dr. Jürgen Sturm, Computer Vision Group, TUM

Camera Calibration

- **Given:** n 2D/3D correspondences $\mathbf{x}_i \leftrightarrow \mathbf{p}_i$
- **Wanted:** $M = K(R \ t)$
such that $\tilde{\mathbf{x}}_i = M\mathbf{p}_i$
- The algorithm has two parts:
 1. Compute $M \in \mathbb{R}^{3 \times 4}$
 2. Decompose M into K, R, t via QR decomposition

Visual Navigation for Flying Robots

55

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 1: Estimate M

- $\tilde{\mathbf{x}}_i = M\mathbf{p}_i$
- Each correspondence generates two equations

$$x = \frac{m_{11}X + m_{12}Y + m_{13}Z + m_{14}W}{m_{31}X + m_{32}Y + m_{33}Z + m_{34}W} \quad y = \frac{m_{21}X + m_{22}Y + m_{23}Z + m_{24}W}{m_{31}X + m_{32}Y + m_{33}Z + m_{34}W}$$
- Multiplying out gives equations **linear** in the elements of M

$$(m_{31}X + m_{32}Y + m_{33}Z + m_{34}W)x = m_{11}X + m_{12}Y + m_{13}Z + m_{14}W$$

$$(m_{31}X + m_{32}Y + m_{33}Z + m_{34}W)y_j = m_{21}X + m_{22}Y + m_{23}Z + m_{24}W$$
- Re-arrange in matrix form...

Visual Navigation for Flying Robots

56

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 1: Estimate M

- Re-arranged in matrix form

$$\begin{pmatrix} X & Y & Z & 1 & 0 & 0 & 0 & 0 & -xX & -xY & -xZ & -x \\ 0 & 0 & 0 & 0 & X & Y & Z & 1 & -yX & -yY & -yZ & -y \end{pmatrix} \mathbf{m} = 0$$
 with $\mathbf{m} = (m_{11} \ m_{12} \ \dots \ m_{34}) \in \mathbb{R}^{12}$
- Concatenate equations for $n \geq 6$ correspondences

$$A\mathbf{m} = 0$$
- Wanted vector \mathbf{m} is in the null space of A
- Initial solution using SVD (vector with least singular value), refine using non-linear min.

Visual Navigation for Flying Robots

57

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 2: Recover K,R,t

- Remember $M = K(R \ t)$
- The first 3x3 submatrix is the product of an upper triangular and orthogonal (rot.) matrix

$$K = \begin{pmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix}$$

Procedure:

1. Factor M into KR using QR decomposition
2. Compute translation as $\mathbf{t} = K^{-1}(p_{14}, p_{24}, p_{34})^\top$


Visual Navigation for Flying Robots

58

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: ARToolkit Markers (1999)

1. Threshold image
2. Detect edges and fit lines
3. Intersect lines to obtain corners
4. Estimate projection matrix M
5. Extract camera pose R, t (assume K is known)



The final error between measured and projected points is typically less than 0.02 pixels


Visual Navigation for Flying Robots

59

Dr. Jürgen Sturm, Computer Vision Group, TUM

Triangulation

- **Given:** cameras $\{M_j = K_j(R_j \ t_j)\}$
point correspondence $\mathbf{x}_0, \mathbf{x}_1$
- **Wanted:** Corresponding 3D point \mathbf{p}



Visual Navigation for Flying Robots

60

Dr. Jürgen Sturm, Computer Vision Group, TUM

Triangulation

- Where do we expect to see $p = (X \ Y \ Z \ W)^T$?

$$\hat{x} = \frac{m_{11}X + m_{12}Y + m_{13}Z + m_{14}W}{m_{31}X + m_{32}Y + m_{33}Z + m_{34}W} \quad \hat{y} = \frac{m_{21}X + m_{22}Y + m_{23}Z + m_{24}W}{m_{31}X + m_{32}Y + m_{33}Z + m_{34}W}$$

- Minimize the residuals (e.g., using least squares)

$$p^* = \arg \min_p \sum_j d(x_j, \hat{x}_j)^2$$


Visual Navigation for Flying Robots

61

Dr. Jürgen Sturm, Computer Vision Group, TUM

Epipolar Geometry

- Consider two cameras that observe a 3D world point




Visual Navigation for Flying Robots

62

Dr. Jürgen Sturm, Computer Vision Group, TUM

Epipolar Geometry

- The line connecting both camera centers is called the **baseline**




Visual Navigation for Flying Robots

63

Dr. Jürgen Sturm, Computer Vision Group, TUM

Epipolar Geometry

- Given the image of a point in one view, what can we say about its position in another?



Visual Navigation for Flying Robots


64

Dr. Jürgen Sturm, Computer Vision Group, TUM

- A point in one image “generates” a line in another image (called the **epipolar line**)

Epipolar Geometry

- Left line in left camera frame $p_1 = d_1 \hat{x}_1$
 - Right line in right camera frame $p_2 = d_2 \hat{x}_2$
- where $\hat{x}_j = K^{-1} \bar{x}_j$ are the (local) ray directions



Visual Navigation for Flying Robots

65

Dr. Jürgen Sturm, Computer Vision Group, TUM

Epipolar Geometry

- Left line in **right** camera frame $p'_1 = R d_1 \hat{x}_1 + t$
 - Right line in right camera frame $p_2 = d_2 \hat{x}_2$
- where $\hat{x}_j = K^{-1} \bar{x}_j$ are the (local) ray directions
- Intersection of both lines

$$d_2 \hat{x}_2 = R d_1 \hat{x}_1 + t$$

$$d_2 [t]_{\times} \hat{x}_2 = d_1 [t]_{\times} R \hat{x}_1 + [t]_{\times} t = 0$$

$$0 = d_2 \hat{x}_2^T [t]_{\times} \hat{x}_2 = d_1 \hat{x}_2^T [t]_{\times} R \hat{x}_1$$

$$0 = \hat{x}_2^T [t]_{\times} R \hat{x}_1$$

$$0 = \hat{x}_2^T E \hat{x}_1$$

this is called the epipolar constraint

Visual Navigation for Flying Robots

66

Dr. Jürgen Sturm, Computer Vision Group, TUM

Epipolar Geometry

Note: The epipolar constraint holds for **every** pair of corresponding points $\mathbf{x}_1, \mathbf{x}_2$

$$\hat{\mathbf{x}}_2^\top E \hat{\mathbf{x}}_1 = 0$$

where E is called the essential matrix

$$E = [\mathbf{t}]_{\times} R \in \mathbb{R}^{3 \times 3}$$

8-Point Algorithm: General Idea

1. Estimate the essential matrix E from at least eight point correspondences
2. Recover the relative pose R, \mathbf{t} from E (up to scale)

Step 1: Estimate E

- Epipolar constraint $\hat{\mathbf{x}}_2^\top E \hat{\mathbf{x}}_1 = 0$
- Written out (with $\mathbf{x}_j = (x_j, y_j, 1)^\top$)

$$\begin{aligned} x_1 x_2 e_{11} + y_1 x_2 e_{12} + x_2 e_{13} + \\ x_1 y_2 e_{21} + y_1 y_2 e_{22} + y_2 e_{23} + \\ x_1 e_{31} + y_1 e_{32} + e_{33} = 0 \end{aligned}$$
- Stack the elements into two vectors

$$\left. \begin{aligned} \mathbf{z} &= (x_1 x_2 \ y_1 x_2 \ \dots \ 1)^\top \\ \mathbf{e} &= (e_{11} \ e_{12} \ \dots \ e_{33})^\top \end{aligned} \right\} \mathbf{z}^\top \mathbf{e} = 0$$

Step 1: Estimate E

- Each correspondence gives us one constraint

$$\left. \begin{aligned} \mathbf{z}_1^\top \mathbf{e} &= 0 \\ \mathbf{z}_2^\top \mathbf{e} &= 0 \\ &\vdots \\ \mathbf{z}_n^\top \mathbf{e} &= 0 \end{aligned} \right\} Z \mathbf{e} = 0$$

- Linear system with n equations
- \mathbf{e} is in the null-space of Z
- Solve using SVD (assuming $\|\mathbf{e}\| = 1$)

Normalized 8-Point Algorithm [Hartley 1997]

- Noise in the point observations is unequally distributed in the constraints, e.g.,


$$\begin{aligned} \text{double noise } & x_1 x_2 e_{11} + y_1 x_2 e_{12} + x_2 e_{13} + \\ & x_1 y_2 e_{21} + y_1 y_2 e_{22} + y_2 e_{23} + \\ \text{normal noise } & x_1 e_{31} + y_1 e_{32} + \underbrace{1}_{\text{noise free}} e_{33} = 0 \end{aligned}$$

- Estimation is sensitive to scaling
- Normalize all points to have zero mean and unit variance

Step 2: Recover R, t

- **Note:** The absolute distance between the two cameras can never be recovered from pure images measurements alone!!!

- Illustration



- We can only recover the translation $\hat{\mathbf{t}}$ up to scale

Step 2a: Recover t

- Remember: $E = [\mathbf{t}]_{\times} R$
- Therefore, \mathbf{t}^{\top} is in the null space of E

$$\mathbf{t}^{\top} E = \underbrace{\mathbf{t}^{\top} [\mathbf{t}]_{\times}}_{=0} R = 0$$

→ Recover $\hat{\mathbf{t}}$ (up to scale) using SVD

$$E = [\hat{\mathbf{t}}]_{\times} R = U \Sigma V^{\top}$$

$$= (\mathbf{u}_0 \quad \mathbf{u}_1 \quad \boxed{\hat{\mathbf{t}}}) \begin{pmatrix} 1 & & \\ & 1 & \\ & & 0 \end{pmatrix} (\mathbf{v}_0^{\top} \quad \mathbf{v}_1^{\top} \quad \mathbf{v}_2^{\top})$$

Visual Navigation for Flying Robots

73

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 2b: Recover R

Remember, the cross-product $[\hat{\mathbf{t}}]_{\times}$

- ... projects a vector onto a set of orthogonal basis vectors including $\hat{\mathbf{t}}$
- ... zeros out the $\hat{\mathbf{t}}$ component
- ... rotates the other two by 90°

$$[\hat{\mathbf{t}}]_{\times} = S Z R_{90^{\circ}} S^{\top}$$

$$= (s_0 \quad s_1 \quad \hat{\mathbf{t}}) \begin{pmatrix} 1 & & \\ & 1 & \\ & & 0 \end{pmatrix} \begin{pmatrix} 0 & -1 \\ 1 & 0 \\ & & 1 \end{pmatrix} \begin{pmatrix} s_0 \\ s_1 \\ \hat{\mathbf{t}} \end{pmatrix}$$

Visual Navigation for Flying Robots

74

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 2b: Recover R

- Plug this into the essential matrix equation

$$E = [\mathbf{t}]_{\times} R = S Z R_{90^{\circ}} S^{\top} R = U \Sigma V^{\top}$$

$$\underbrace{\hspace{10em}}_{=}$$

$$\underbrace{\hspace{10em}}_{=}$$

- By identifying $S = U$ and $Z = \Sigma$, we obtain

$$R_{90^{\circ}} U^{\top} R = V^{\top}$$

$$\boxed{R = U R_{90^{\circ}}^{\top} V^{\top}}$$

Visual Navigation for Flying Robots

75

Dr. Jürgen Sturm, Computer Vision Group, TUM

Summary: 8-Point Algorithm

Given: Image pair



Find: Camera motion R, \mathbf{t} (up to scale)

- Compute correspondences
- Compute essential matrix
- Extract camera motion

Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM

How To Deal With Outliers?



Problem: No matter how good the feature descriptor/matcher is, there is always a chance for bad point correspondences (=outliers)


Visual Navigation for Flying Robots

77

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robust Estimation

Example: Fit a line to 2D data containing outliers



There are two problems

- Fit** the line to the data $\arg \min_l \sum_i d_i^2$
- Classify** the data into inliers (valid points) and outliers (using some threshold)

Visual Navigation for Flying Robots

78

Dr. Jürgen Sturm, Computer Vision Group, TUM


RANdom Sample Consensus (RANSAC) [Fischler and Bolles, 1981]

Goal: Robustly fit a model to a data set S which contains outliers

Algorithm:

1. Randomly select a (minimal) subset of data points and instantiate the model from it
2. Using this model, classify the all data points as inliers or outliers
3. Repeat 1&2 for N iterations
4. Select the largest inlier set, and re-estimate the model from all points in this set

RANdom Sample Consensus (RANSAC)



- RANSAC is used very widely
- Many improvements/variants, e.g., MLESAC:

$$\arg \min_i \sum_i \rho(d_i) \text{ with } \rho(d) = \begin{cases} d^2 & \text{if } d \leq e(\text{inlier}) \\ e^2 & \text{if } d > e(\text{outlier}) \end{cases}$$

How Many Samples?

- For probability p of having no outliers, we need

to sample $N = \frac{\log(1-p)}{\log(1-(1-\epsilon)^s)}$ subsets

for subset size s and outlier ratio ϵ

- E.g., for $p=0.95$:


Sample size s	Proportion of outliers ϵ						
	5%	10%	20%	25%	30%	40%	50%
2	2	2	3	4	5	7	11
3	2	3	5	6	8	13	23
4	2	3	6	8	11	22	47
5	3	4	8	12	17	38	95
6	3	4	10	16	24	63	191
7	3	5	13	21	35	106	382
8	3	6	17	29	51	177	766

Two Examples

- PTAM
G. Klein and D. Murray, Parallel Tracking and Mapping for Small AR Workspaces, International Symposium on Mixed and Augmented Reality (ISMAR), 2007
<http://www.robots.ox.ac.uk/~gk/publications/KleinMurray2007ISMAR.pdf>
- Photo Tourism
N. Snavely, S. M. Seitz, R. Szeliski, Photo tourism: Exploring photo collections in 3D, ACM Transactions on Graphics (SIGGRAPH), 2006
http://phototour.cs.washington.edu/Photo_Tourism.pdf

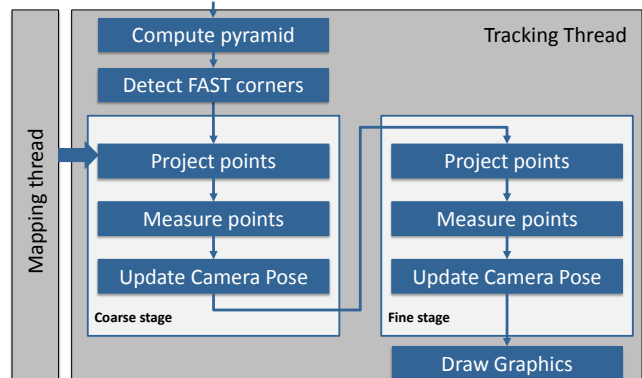
PTAM (2007)

- Architecture optimized for dual cores



- Tracking thread runs in real-time (30Hz)
- Mapping thread is not real-time

PTAM – Tracking Thread



PTAM – Feature Tracking

- Generate 8x8 matching template (warped from key frame to current pose estimate)
- Search a fixed radius around projected position
 - Using SSD
 - Only search at FAST corner points

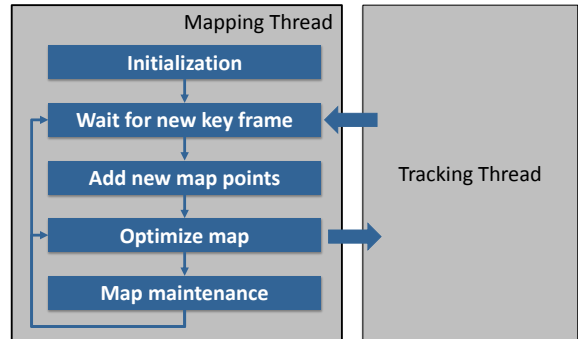


Visual Navigation for Flying Robots

85

Dr. Jürgen Sturm, Computer Vision Group, TUM

PTAM – Mapping Thread



Visual Navigation for Flying Robots

86

Dr. Jürgen Sturm, Computer Vision Group, TUM

PTAM – Example Timings

- Tracking thread

Total	19.2 ms
Key frame preparation	2.2 ms
Feature Projection	3.5 ms
Patch search	9.8 ms
Iterative pose update	3.7 ms

- Mapping thread

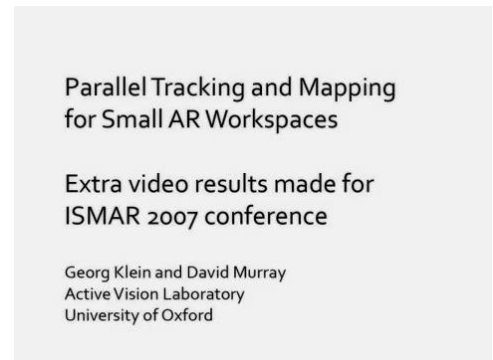
Key frames	2-49	50-99	100-149
Local Bundle Adjustment	170 ms	270 ms	440 ms
Global Bundle Adjustment	380 ms	1.7 s	6.9 s

Visual Navigation for Flying Robots

87

Dr. Jürgen Sturm, Computer Vision Group, TUM

PTAM Video



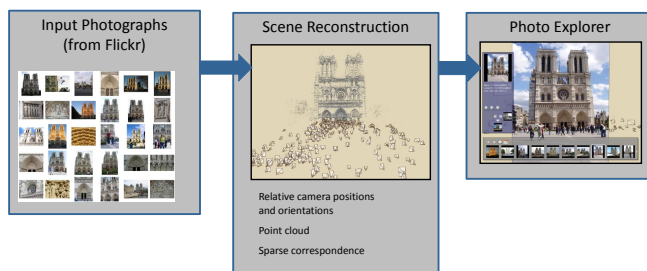
Visual Navigation for Flying Robots

88

Dr. Jürgen Sturm, Computer Vision Group, TUM

Photo Tourism (2006)

- Overview




Visual Navigation for Flying Robots

89

Dr. Jürgen Sturm, Computer Vision Group, TUM

Photo Tourism – Scene Reconstruction

- Processing pipeline



- Automatically estimate

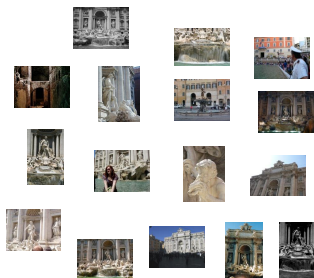
- Position, orientation and focal length of all cameras
- 3D positions of point features

Visual Navigation for Flying Robots

90

Dr. Jürgen Sturm, Computer Vision Group, TUM

Photo Tourism – Input Images



Visual Navigation for Flying Robots

91

Dr. Jürgen Sturm, Computer Vision Group, TUM

Photo Tourism – Feature Detection




Visual Navigation for Flying Robots

92

Dr. Jürgen Sturm, Computer Vision Group, TUM

Photo Tourism – Feature Matching



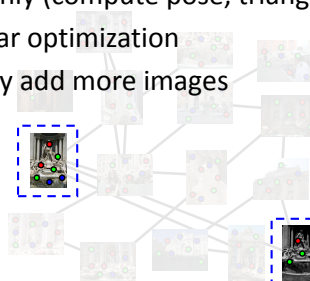
Visual Navigation for Flying Robots

93

Dr. Jürgen Sturm, Computer Vision Group, TUM

Incremental Structure From Motion

- To help get good initializations, start with two images only (compute pose, triangulate points)
- Non-linear optimization
- Iteratively add more images



Visual Navigation for Flying Robots

94

Dr. Jürgen Sturm, Computer Vision Group, TUM

Photo Tourism – Video



Visual Navigation for Flying Robots

95

Dr. Jürgen Sturm, Computer Vision Group, TUM

Lessons Learned Today

- ... how to detect and match feature points
- ... how to compute the camera pose and to triangulate points
- ... how to deal with outliers

Visual Navigation for Flying Robots

96

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

Bundle Adjustment and Stereo Correspondence

Dr. Jürgen Sturm

Project Proposal Presentations

- This Thursday
- Don't forget to put title, team name, team members on first slide
- Pitch has to fit in 5 minutes (+5 minutes discussion)
- 9 x (5+5) = 90 minutes
- Recommendation: use 3-5 slides

Visual Navigation for Flying Robots

2

Dr. Jürgen Sturm, Computer Vision Group, TUM

Agenda for Today

- Map optimization
 - Graph SLAM
 - Bundle adjustment
- Depth reconstruction
 - Laser triangulation
 - Structured light (Kinect)
 - Stereo cameras

Visual Navigation for Flying Robots

3

Dr. Jürgen Sturm, Computer Vision Group, TUM

Remember: 3D Transformations

- Representation as a homogeneous matrix

$$M = \begin{pmatrix} R & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{pmatrix} \in \text{SE}(3) \subset \mathbb{R}^{4 \times 4}$$

Pro: easy to concatenate and invert
Con: not minimal

- Representation as a twist coordinates

$$\xi = (v_x \ v_y \ v_z \ \omega_x \ \omega_y \ \omega_z)^\top \in \mathbb{R}^6$$

Pro: minimal
Con: need to convert to matrix for concatenation and inversion

Remember: 3D Transformations

- From twist coordinates to twist

$$\hat{\xi} = \begin{pmatrix} 0 & -\omega_z & \omega_y & v_x \\ \omega_z & 0 & -\omega_x & v_y \\ -\omega_y & \omega_x & 0 & v_z \\ 0 & 0 & 0 & 0 \end{pmatrix} \in \text{se}(3)$$

- Exponential map between $\text{se}(3)$ and $\text{SE}(3)$

$$M = \exp \hat{\xi} \quad \hat{\xi} = \log M$$

alternative notation: $M = \exp[\xi]^\wedge \quad \xi = [\log M]^\vee$

Visual Navigation for Flying Robots

5

Dr. Jürgen Sturm, Computer Vision Group, TUM

Remember: Rodrigues' formula

- **Given:** Twist coordinates

$$\xi = (\omega^\top, \mathbf{v}^\top)^\top = (\omega_x, \omega_y, \omega_z, v_x, v_y, v_z)^\top \\ = (t\bar{\omega}^\top, \mathbf{v}^\top)^\top \text{ with } \|\bar{\omega}\| = 1, t = \|\omega\|$$

- **Return:** Homogeneous transformation

$$R = I + [\bar{\omega}]_\times \sin(t) + [\bar{\omega}]_\times^2 (1 - \cos t)$$

$$\mathbf{t} = (I - R)[\bar{\omega}]_\times \mathbf{v} + \bar{\omega} \bar{\omega}^\top \mathbf{v} t$$

$$M = \begin{pmatrix} R & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{pmatrix}$$

Visual Navigation for Flying Robots

6

Dr. Jürgen Sturm, Computer Vision Group, TUM

Notation

- Camera poses in a minimal representation (e.g., twists)

$$\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_n$$

- ... as transformation matrices

$$M_1, M_2, \dots, M_n$$

- ... as rotation matrices and translation vectors

$$(R_1, \mathbf{t}_1), (R_2, \mathbf{t}_2), \dots, (R_n, \mathbf{t}_n)$$


Visual Navigation for Flying Robots

7

Dr. Jürgen Sturm, Computer Vision Group, TUM

Incremental Motion Estimation

- **Idea:** Estimate camera motion from frame to frame



Visual Navigation for Flying Robots

8

Dr. Jürgen Sturm, Computer Vision Group, TUM


Incremental Motion Estimation

- **Idea:** Estimate camera motion from frame to frame

- **Motion concatenation (for twists)**

$$\hat{\mathbf{c}}_j = \log(\exp \hat{\mathbf{c}}_i \exp \hat{\mathbf{z}}_{ij})$$

- **Motion composition operator (in general)**




Visual Navigation for Flying Robots

9

Dr. Jürgen Sturm, Computer Vision Group, TUM

Incremental Motion Estimation

- **Idea:** Estimate camera motion from frame to frame




Visual Navigation for Flying Robots

10

Dr. Jürgen Sturm, Computer Vision Group, TUM

Incremental Motion Estimation

- **Idea:** Estimate camera motion from frame to frame



Visual Navigation for Flying Robots

11

Dr. Jürgen Sturm, Computer Vision Group, TUM

Loop Closures

- **Idea:** Estimate camera motion from frame to frame

- **Problem:**

- Estimates are inherently noisy
- Error accumulates over time \rightarrow drift


Visual Navigation for Flying Robots

12

Dr. Jürgen Sturm, Computer Vision Group, TUM

Incremental Motion Estimation

- **Idea:** Estimate camera motion from frame to frame




Visual Navigation for Flying Robots

13

Dr. Jürgen Sturm, Computer Vision Group, TUM

Incremental Motion Estimation

- **Idea:** Estimate camera motion from frame to frame
- **Two ways to compute c_n :** $c_n = c_{n-1} \oplus z_{(n-1)n}$
 $c'_n = c_1 \oplus z_{1n}$




Visual Navigation for Flying Robots

14

Dr. Jürgen Sturm, Computer Vision Group, TUM

Loop Closures

- **Solution:** Use loop-closures to minimize the drift / minimize the error over all constraints



Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

Graph SLAM

[Thrun and Montemerlo, 2006; Olson et al., 2006]


- Use a graph to represent the model
- Every **node** in the graph corresponds to a pose of the robot during mapping
- Every **edge** between two nodes corresponds to a spatial constraint between them
- **Graph-based SLAM:** Build the graph and find the robot poses that **minimize the error** introduced by the constraints

Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Graph SLAM on Intel Dataset




Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

Graph SLAM Architecture



- Interleaving process of front-end and back-end
- A consistent map helps to determine new constraints by reducing the search space


Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM

Problem Definition

- **Given:** Set of observations $\mathbf{z}_{ij} \in \mathbb{R}^6$
- **Wanted:** Set of camera poses $\mathbf{c}_1, \dots, \mathbf{c}_n \in \mathbb{R}^6$
 → State vector $\mathbf{x} = (\mathbf{c}_1^\top, \dots, \mathbf{c}_n^\top)^\top \in \mathbb{R}^{6n}$



Visual Navigation for Flying Robots

19

Dr. Jürgen Sturm, Computer Vision Group, TUM

Map Error

- Real observation \mathbf{z}_{ij}
- Expected observation $\bar{\mathbf{z}}_{ij} = \mathbf{c}_j \ominus \mathbf{c}_i$
- Difference between observation and expectation

$$\mathbf{e}_{ij} = \mathbf{z}_{ij} \ominus \bar{\mathbf{z}}_{ij}$$

- Given the correct map, this difference is the result of sensor noise...

Visual Navigation for Flying Robots

20

Dr. Jürgen Sturm, Computer Vision Group, TUM

Error Function

- **Assumption:** Sensor noise is normally distributed

$$\mathbf{e}_{ij} \sim \mathcal{N}(\mathbf{0}, \Sigma_{ij})$$

- Error term for one observation (proportional to negative loglikelihood)

$$f_{ij}(\mathbf{x}) = \mathbf{e}_{ij}(\mathbf{x})^\top \Sigma_{ij}^{-1} \mathbf{e}_{ij}(\mathbf{x})$$

- Note: error is a scalar $f_{ij}(\mathbf{x}) \in \mathbb{R}$

Visual Navigation for Flying Robots

21

Dr. Jürgen Sturm, Computer Vision Group, TUM

Error Function

- Map error (over all observations)

$$f(\mathbf{x}) = \sum_{ij} f_{ij}(\mathbf{x}) = \sum_{ij} \mathbf{e}_{ij}(\mathbf{x})^\top \Sigma_{ij}^{-1} \mathbf{e}_{ij}(\mathbf{x})$$

- **Minimize this error** by optimizing the camera poses

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \sum_{ij} \mathbf{e}_{ij}(\mathbf{x})^\top \Sigma_{ij}^{-1} \mathbf{e}_{ij}(\mathbf{x})$$

- How can we solve this optimization problem?

Visual Navigation for Flying Robots

22

Dr. Jürgen Sturm, Computer Vision Group, TUM

Non-Linear Optimization Techniques

- Gradient descend
- Gauss-Newton
- Levenberg-Marquardt

Visual Navigation for Flying Robots

23

Dr. Jürgen Sturm, Computer Vision Group, TUM

Gauss-Newton Method

1. Linearize the error function
2. Compute its derivative
3. Set the derivative to zero
4. Solve the linear system
5. Iterate this procedure until convergence

Visual Navigation for Flying Robots

24

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 1: Linearize the Error Function

- Error function

$$f(\mathbf{x}) = \sum_{ij} f_{ij}(\mathbf{x}) = \sum_{ij} \mathbf{e}_{ij}(\mathbf{x})^\top \Sigma_{ij}^{-1} \mathbf{e}_{ij}(\mathbf{x})$$

- Evaluate the error function around the initial guess

$$f(\mathbf{x} + \Delta\mathbf{x}) = \sum_{ij} \mathbf{e}_{ij}(\mathbf{x} + \Delta\mathbf{x})^\top \Sigma_{ij}^{-1} \mathbf{e}_{ij}(\mathbf{x} + \Delta\mathbf{x})$$

Let's derive this term first...

Linearize the Error Function

- Approximate the error function around an initial guess \mathbf{x} using Taylor expansion

$$\mathbf{e}_{ij}(\mathbf{x} + \Delta\mathbf{x}) \simeq \mathbf{e}_{ij}(\mathbf{x}) + J_{ij}\Delta\mathbf{x}$$

with

$$J_{ij}(\mathbf{x}) = \left(\frac{\partial \mathbf{e}_{ij}(\mathbf{x})}{\partial c_1} \quad \frac{\partial \mathbf{e}_{ij}(\mathbf{x})}{\partial c_2} \quad \dots \quad \frac{\partial \mathbf{e}_{ij}(\mathbf{x})}{\partial c_n} \right)$$

Derivatives of the Error Terms

- Does one error function $\mathbf{e}_{ij}(\mathbf{x})$ depend on all state variables in \mathbf{x} ?

Derivatives of the Error Terms

- Does one error function $\mathbf{e}_{ij}(\mathbf{x})$ depend on all state variables in \mathbf{x} ?
 - No, $\mathbf{e}_{ij}(\mathbf{x})$ depends only on c_i and c_j

Derivatives of the Error Terms

- Does one error function $\mathbf{e}_{ij}(\mathbf{x})$ depend on all state variables in \mathbf{x} ?
 - No, $\mathbf{e}_{ij}(\mathbf{x})$ depends only on c_i and c_j
- Is there any consequence on the **structure** of the Jacobian?

Derivatives of the Error Terms

- Does one error function $\mathbf{e}_{ij}(\mathbf{x})$ depend on all state variables in \mathbf{x} ?
 - No, $\mathbf{e}_{ij}(\mathbf{x})$ depends only on c_i and c_j
- Is there any consequence on the **structure** of the Jacobian?
 - Yes, it will be non-zero only in the columns corresponding to c_i and c_j
 - Jacobian is **sparse**

$$J_{ij}(\mathbf{x}) = \left(\mathbf{0} \quad \dots \quad \frac{\partial \mathbf{e}_{ij}(\mathbf{x})}{\partial c_i} \quad \dots \quad \frac{\partial \mathbf{e}_{ij}(\mathbf{x})}{\partial c_j} \quad \dots \quad \mathbf{0} \right)$$

Linearizing the Error Function

$$\text{Linearize } f(\mathbf{x}) = \sum_{ij} \mathbf{e}_{ij}(\mathbf{x})^T \Sigma_{ij}^{-1} \mathbf{e}_{ij}(\mathbf{x})$$

$$\simeq \mathbf{c} + 2\mathbf{b}^T \Delta \mathbf{x} + \Delta \mathbf{x}^T H \Delta \mathbf{x}$$

$$\text{with } \mathbf{b}^T = \sum_{ij} \mathbf{e}_{ij}^T \Sigma_{ij}^{-1} J_{ij}$$

$$H = \sum_{ij} J_{ij}^T \Sigma_{ij}^{-1} J_{ij}$$

- What is the structure of \mathbf{b}^T and H ? (Remember: all J_{ij} 's are sparse)

Visual Navigation for Flying Robots

32

Dr. Jürgen Sturm, Computer Vision Group, TUM

Illustration of the Structure

$$\mathbf{b}_{ij}^T = \mathbf{e}_{ij}^T \Sigma_{ij}^{-1} J_{ij}$$

Illustration of the Structure

$$\mathbf{b}_{ij}^T = \mathbf{e}_{ij}^T \Sigma_{ij}^{-1} J_{ij}$$

$$H_{ij} = J_{ij}^T \Sigma_{ij}^{-1} J_{ij}$$

Illustration of the Structure

$$\mathbf{b}_{ij}^T = \mathbf{e}_{ij}^T \Sigma_{ij}^{-1} J_{ij}$$

$$H_{ij} = J_{ij}^T \Sigma_{ij}^{-1} J_{ij}$$

Illustration of the Structure

$$\mathbf{b} = \sum_{ij} \mathbf{b}_{ij}$$

$$H = \sum_{ij} H_{ij}$$

(Linear) Least Squares Minimization

1. Linearize error function

$$f(\mathbf{x} + \Delta \mathbf{x}) \simeq \mathbf{c} + 2\mathbf{b}^T \Delta \mathbf{x} + \Delta \mathbf{x}^T H \Delta \mathbf{x}$$

2. Compute the derivative

$$\frac{df(\mathbf{x} + \Delta \mathbf{x})}{d\Delta \mathbf{x}} = 2\mathbf{b} + 2H\Delta \mathbf{x}$$

3. Set derivative to zero

$$H\Delta \mathbf{x} = -\mathbf{b}$$

4. Solve this linear system of equations, e.g.,

$$\Delta \mathbf{x} = -H^{-1}\mathbf{b}$$

Visual Navigation for Flying Robots

37

Dr. Jürgen Sturm, Computer Vision Group, TUM

Gauss-Newton Method

Problem: $f(\mathbf{x})$ is non-linear!

Algorithm: Repeat until convergence

1. Compute the terms of the linear system

$$\mathbf{b}^\top = \sum_{ij} \mathbf{e}_{ij}^\top \Sigma_{ij}^{-1} J_{ij} \quad H = \sum_{ij} J_{ij}^\top \Sigma_{ij}^{-1} J_{ij}$$

2. Solve the linear system to get new increment

$$H \Delta \mathbf{x} = -\mathbf{b}$$

3. Update previous estimate

$$\mathbf{x} \leftarrow \mathbf{x} + \Delta \mathbf{x}$$

Visual Navigation for Flying Robots

38

Dr. Jürgen Sturm, Computer Vision Group, TUM

Sparsity of the Hessian

- The Hessian is
 - positive semi-definit
 - symmetric
 - sparse
- This allows the use of efficient solvers
 - Sparse Cholesky decomposition (~100M matrix elements)
 - Preconditioned conjugate gradients (~1.000M matrix elements)
 - ... many others


Visual Navigation for Flying Robots

39

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example in 1D

- Two camera poses $c_1, c_2 \in \mathbb{R}$
- State vector $\mathbf{x} = (c_1, c_2)^\top \in \mathbb{R}^2$
- One (distance) observation $z_{12} \in \mathbb{R}$
- Initial guess $c_1 = c_2 = 0$
- Observation $z_{12} = 1$
- Sensor noise $\Sigma_{12} = 0.5$



Visual Navigation for Flying Robots

40

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example in 1D

- Error $e_{12} = z_{12} - \bar{z}_{12}$

$$= z_{12} - (c_2 - c_1) = 1 - (0 - 0) = 1$$
- Jacobian $J_{12} = \begin{pmatrix} \frac{\partial e_{12}}{\partial c_1} & \frac{\partial e_{12}}{\partial c_2} \end{pmatrix} = \begin{pmatrix} 1 & -1 \end{pmatrix}$
- Build linear system of equations

$$\mathbf{b}^\top = \mathbf{e}_{12}^\top \Sigma^{-1} e_{12} = \begin{pmatrix} 2 & -2 \end{pmatrix}$$

$$H = J_{12}^\top \Sigma^{-1} J_{12} = \begin{pmatrix} 2 & -2 \\ -2 & 2 \end{pmatrix}$$
- Solve the system

$$\Delta \mathbf{x} = -H^{-1} \mathbf{b} \quad \text{but } \det H = 0 \text{ ???}$$

Visual Navigation for Flying Robots

41

Dr. Jürgen Sturm, Computer Vision Group, TUM

What Went Wrong?

- The constraint only specifies a **relative constraint** between two nodes
- Any poses for the nodes would be fine as long as their relative coordinates fit
- **One node needs to be fixed**
 - Option 1: Remove one row/column corresponding to the fixed pose
 - Option 2: Add to H , \mathbf{b} a linear constraint $1 \cdot \Delta c_1 = 0$
 - Option 3: Add the identity matrix to H (Levenberg-Marquardt)

Visual Navigation for Flying Robots

42

Dr. Jürgen Sturm, Computer Vision Group, TUM

Fixing One Node

- The constraint only specifies a **relative constraint** between two nodes
- Any poses for the nodes would be fine as long as their relative coordinates fit
- **One node needs to be fixed (here: Option 2)**

$$H = \begin{pmatrix} 2 & -2 \\ -2 & 2 \end{pmatrix} + \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad \text{additional constraint that sets } \Delta c_1 = 0$$

$$\Delta \mathbf{x} = -H^{-1} \mathbf{b}$$

$$\Delta \mathbf{x} = \begin{pmatrix} 0 & 1 \end{pmatrix}^\top$$

Visual Navigation for Flying Robots

43

Dr. Jürgen Sturm, Computer Vision Group, TUM

Levenberg-Marquardt Algorithm

- **Idea:** Add a damping factor

$$(H + \lambda I)\Delta\mathbf{x} = -\mathbf{b}$$

$$(J^T J + \lambda I)\Delta\mathbf{x} = -J^T \mathbf{e}$$

- What is the effect of this damping factor?
 - Small λ ?
 - Large λ ?

Levenberg-Marquardt Algorithm

- **Idea:** Add a damping factor

$$(H + \lambda I)\Delta\mathbf{x} = -\mathbf{b}$$

$$(J^T J + \lambda I)\Delta\mathbf{x} = -J^T \mathbf{e}$$

- What is the effect of this damping factor?
 - Small $\lambda \rightarrow$ same as least squares
 - Large $\lambda \rightarrow$ steepest descent (with small step size)
- **Algorithm**
 - If error decreases, accept $\Delta\mathbf{x}$ and reduce λ
 - If error increases, reject $\Delta\mathbf{x}$ and increase λ

Non-Linear Minimization

- One of the state-of-the-art solution to compute the maximum likelihood estimate
- Various open-source implementations available
 - g2o [Kuemmerle et al., 2011]
 - sba [Lourakis and Argyros, 2009]
 - iSAM [Kaess et al., 2008]
- Other extensions:
 - Robust error functions
 - Alternative parameterizations

Bundle Adjustment

- **Graph SLAM:** Optimize (only) the camera poses

$$\mathbf{x} = (\mathbf{c}_1^T, \dots, \mathbf{c}_n^T)^T \in \mathbb{R}^{6n}$$

- **Bundle Adjustment:** Optimize both 6DOF camera poses and 3D (feature) points

$$\mathbf{x} = \underbrace{(\mathbf{c}_1^T, \dots, \mathbf{c}_n^T)^T}_{\mathbf{c} \in \mathbb{R}^{6n}}, \underbrace{(\mathbf{p}_1^T, \dots, \mathbf{p}_m^T)^T}_{\mathbf{p} \in \mathbb{R}^{3m}} \in \mathbb{R}^{6n+3m}$$

- Typically $m \gg n$ (why?)

Error Function

- Camera pose $\mathbf{c}_i \in \mathbb{R}^6$
- Feature point $\mathbf{p}_j \in \mathbb{R}^3$
- Observed feature location $\mathbf{z}_{ij} \in \mathbb{R}^2$
- Expected feature location

$$g(\mathbf{c}_i, \mathbf{p}_j) = R_i^T (\mathbf{t}_i - \mathbf{p}_j)$$

$$h(\mathbf{c}_i, \mathbf{p}_j) = g_{x,y}(\mathbf{c}_i, \mathbf{p}_j) / g_z(\mathbf{c}_i, \mathbf{p}_j)$$

Error Function

- Difference between observation and expectation

$$\mathbf{e}_{ij} = \mathbf{z}_{ij} - h(\mathbf{c}_i, \mathbf{p}_j)$$


- Error function

$$f(\mathbf{c}, \mathbf{p}) = \sum_{ij} \mathbf{e}_{ij}^T \Sigma^{-1} \mathbf{e}_{ij}$$

- Covariance Σ is often chosen isotropic and on the order of one pixel

Illustration of the Structure

- Each camera sees several points
- Each point is seen by several cameras
- Cameras are independent of each other (given the points), same for the points



Visual Navigation for Flying Robots

50 Dr. Jürgen Sturm, Computer Vision Group, TUM

Primary Structure

- Characteristic structure

$$\begin{pmatrix} J_c^\top J_c & J_c^\top J_p \\ J_p^\top J_c & J_p^\top J_p \end{pmatrix} \begin{pmatrix} \Delta c \\ \Delta p \end{pmatrix} = \begin{pmatrix} -J_c^\top e_c \\ -J_p^\top e_p \end{pmatrix}$$

$$\begin{pmatrix} H_{cc} & H_{cp} \\ H_{pc} & H_{pp} \end{pmatrix} \begin{pmatrix} \Delta c \\ \Delta p \end{pmatrix} = \begin{pmatrix} -J_c^\top e_c \\ -J_p^\top e_p \end{pmatrix}$$

Visual Navigation for Flying Robots

51 Dr. Jürgen Sturm, Computer Vision Group, TUM

Primary Structure

- Insight:** H_{cc} and H_{pp} are block-diagonal (because each constraint depends only on one camera and one point)

$$\begin{pmatrix} \begin{matrix} \square & & \\ & \square & \\ & & \ddots \end{matrix} & & \\ & & \begin{matrix} \square & & \\ & \square & \\ & & \ddots \end{matrix} \end{pmatrix} \begin{pmatrix} \Delta c \\ \Delta p \end{pmatrix} = \begin{pmatrix} -J_c^\top e_c \\ -J_p^\top e_p \end{pmatrix}$$

- This can be efficiently solved using the Schur Complement

Visual Navigation for Flying Robots

52 Dr. Jürgen Sturm, Computer Vision Group, TUM

Schur Complement

- Given: Linear system

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} a \\ b \end{pmatrix}$$

- If D is invertible, then (using Gauss elimination)

$$\begin{aligned} (A - BD^{-1}C)x &= a - BD^{-1}b \\ y &= D^{-1}(b - Cx) \end{aligned}$$


- Reduced complexity**, i.e., invert one $p \times p$ and $p \times p$ matrix instead of one $(p + q) \times (p + q)$ matrix

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Hessian (Lourakis and Argyros, 2009)




Visual Navigation for Flying Robots

54

Dr. Jürgen Sturm, Computer Vision Group, TUM

From Sparse Maps to Dense Maps

- So far, we only looked at sparse 3D maps
 - We know where the (sparse) cameras are
 - We know where the (sparse) 3D feature points are
- How can we turn these models into volumetric 3D models?



Visual Navigation for Flying Robots

56

Dr. Jürgen Sturm, Computer Vision Group, TUM

From Sparse Maps to Dense Maps

- **Today:** Estimation of depth dense images (stereo cameras, laser triangulation, structured light/Kinect)
- **Next week:** Dense map representations and data fusion




Visual Navigation for Flying Robots

57


Dr. Jürgen Sturm, Computer Vision Group, TUM

Human Stereo Vision



Stereo Correspondence Constraints

- Given a point in the left image, where can the corresponding point be in the right image?




Visual Navigation for Flying Robots

59

Dr. Jürgen Sturm, Computer Vision Group, TUM

Reminder: Epipolar Geometry

- A point in one image “generates” a line in another image (called the **epipolar line**)
- Epipolar constraint $\hat{x}_2^T E \hat{x}_1 = 0$




Visual Navigation for Flying Robots

Dr. Jürgen Sturm, Computer Vision Group, TUM

Epipolar Plane

- All epipolar lines intersect at the epipoles
- An epipolar plane intersects the left and right image planes in epipolar lines



Visual Navigation for Flying Robots

61

Dr. Jürgen Sturm, Computer Vision Group, TUM

Epipolar Constraint




- This is useful because it reduces the correspondence problem to a 1D search along an epipolar line

Visual Navigation for Flying Robots

62

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Converging Cameras

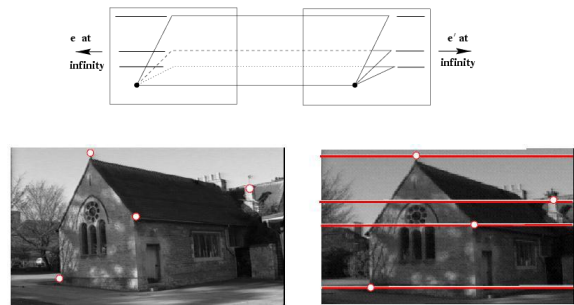


Visual Navigation for Flying Robots

63

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Parallel Cameras




Visual Navigation for Flying Robots

64

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rectification


- In practice, it is convenient if the image scanlines (rows) are the epipolar lines
- Reproject image planes onto a common plane parallel to the baseline (two 3x3 homographies)
- Afterwards pixel motion is horizontal



Visual Navigation for Flying Robots

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Rectification




Visual Navigation for Flying Robots

66

Dr. Jürgen Sturm, Computer Vision Group, TUM

Basic Stereo Algorithm

- For each pixel in the left image
 - Compare with every pixel on the same epipolar line in the right image
 - Pick pixel with minimum matching cost (noisy)
 - Better: match small blocks/patches (SSD, SAD, NCC)



left image

Visual Navigation for Flying Robots

67



right image

Dr. Jürgen Sturm, Computer Vision Group, TUM

Block Matching Algorithm

Input: Two images and camera calibrations

Output: Disparity (or depth) image

Algorithm:

1. Geometry correction (undistortion and rectification)
2. Matching cost computation along search window
3. Extrema extraction (at sub-pixel accuracy)
4. Post-filtering (clean up noise)

Visual Navigation for Flying Robots

68

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Input



- Output



Visual Navigation for Flying Robots

69


Dr. Jürgen Sturm, Computer Vision Group, TUM

What is the Influence of the Block Size?

- Common choices are 5x5 .. 11x11
- Smaller neighborhood: more details
- Larger neighborhood: less noise
- Suppress pixels with low confidence (e.g., check ratio best match vs. 2nd best match)



3x3



20x20


Visual Navigation for Flying Robots

70

Dr. Jürgen Sturm, Computer Vision Group, TUM

Problems with Stereo

- Block matching typically fails in regions with low texture
 - Global optimization/regularization (speciality of our research group)
 - Additional texture projection





Visual Navigation for Flying Robots

71

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: PR2 Robot with Projected Texture Stereo



Visual Navigation for Flying Robots

72

Dr. Jürgen Sturm, Computer Vision Group, TUM

Laser Triangulation

Idea:

- Well-defined light pattern (e.g., point or line) projected on scene
- Observed by a line/matrix camera or a position-sensitive device (PSD)
- Simple triangulation to compute distance


Visual Navigation for Flying Robots

73

Dr. Jürgen Sturm, Computer Vision Group, TUM

Laser Triangulation

- Function principle



- Depth triangulation $z = f \frac{L}{d}$
(note: same for stereo disparities)


Visual Navigation for Flying Robots

74

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Neato XV-11

- K. Konolige, "A low-cost laser distance sensor", ICRA 2008
- Specs: 360deg, 10Hz, 30 USD




Visual Navigation for Flying Robots

75

Dr. Jürgen Sturm, Computer Vision Group, TUM

How Does the Data Look Like?




Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM


Laser Triangulation

- Stripe laser + 2D camera
- Often used on conveyer belts (volume sensing)
- Large baseline gives better depth resolution but more occlusions → use two cameras



Visual Navigation for Flying Robots


77



Dr. Jürgen Sturm, Computer Vision Group, TUM

Structured Light

- Multiple stripes / 2D pattern
- Data association more difficult




Visual Navigation for Flying Robots

78

Dr. Jürgen Sturm, Computer Vision Group, TUM


Structured Light

- Multiple stripes / 2D pattern
- Data association more difficult
- Coding schemes
 - Temporal: Coded light



Visual Navigation for Flying Robots

79



Dr. Jürgen Sturm, Computer Vision Group, TUM

Structured Light

- Multiple stripes / 2D pattern
- Data association more difficult
- Coding schemes
 - Temporal: Coded light
 - Wavelength: Color
 - Spatial: Pattern (e.g., diffraction patterns)



Visual Navigation for Flying Robots


80



Dr. Jürgen Sturm, Computer Vision Group, TUM

Sensor Principle of Kinect

- Kinect projects a diffraction pattern (speckles) in near-infrared light
- CMOS IR camera observes the scene



Visual Navigation for Flying Robots

81

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Data

- Kinect provides color (RGB) and depth (D) video
- This allows for novel approaches for (robot) perception

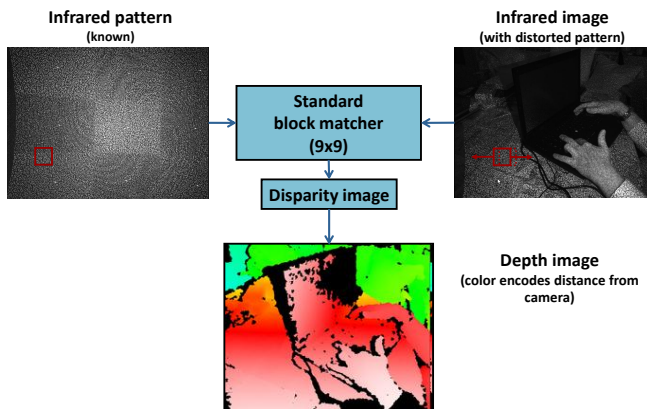


Visual Navigation for Flying Robots

82

Dr. Jürgen Sturm, Computer Vision Group, TUM

Sensor Principle of Kinect



- Pattern is memorized at a known depth
- For each pixel in the IR image
 - Extract 9x9 template from memorized pattern
 - Correlate with current IR image over 64 pixels and search for the maximum
 - Interpolate maximum to obtain sub-pixel accuracy (1/8 pixel)
 - Calculate depth by triangulation

Visual Navigation for Flying Robots

84

Dr. Jürgen Sturm, Computer Vision Group, TUM

Technical Specs

- Infrared camera has 640x480 @ 30 Hz
 - Depth correlation runs on FPGA
 - 11-bit depth image
 - 0.8m – 5m range
 - Depth sensing does not work in direct sunlight (why?)
- RGB camera has 640x480 @ 30 Hz
 - Bayer color filter
- Four 16-bit microphones with DSP for beam forming @ 16kHz
- Requires 12V (for motor), weighs 500 grams
- Human pose recognition runs on Xbox CPU and uses only 10-15% processing power @ 30 Hz
 (Paper: <http://research.microsoft.com/apps/pubs/default.aspx?id=145347>)

Visual Navigation for Flying Robots

85

Dr. Jürgen Sturm, Computer Vision Group, TUM

History

- 2005: Developed by PrimeSense (Israel)
- 2006: Offer to Nintendo and Microsoft, both companies declined
- 2007: Alex Kidman becomes new incubation director at Microsoft, decides to explore PrimeSense device. Johnny Lee assembles a team to investigate technology and develop game concepts
- 2008: The group around Prof. Andrew Blake and Jamie Shotton (Microsoft Research) develops pose recognition
- 2009: The group around Prof. Dieter Fox (Intel Labs / Univ. of Washington) works on RGB-D mapping and RGB-D object recognition
- Nov 4, 2010: Official market launch
- Nov 10, 2010: First open-source driver available
- 2011: First programming competitions (ROS 3D, PrimeSense), First workshops (RSS, Euron)
- 2012: First special Issues (JVCI, T-SMC)


Visual Navigation for Flying Robots

86

Dr. Jürgen Sturm, Computer Vision Group, TUM

Impact of the Kinect Sensor

- Sold >18M units, >8M in first 60 days (Guinness: “fastest selling consumer electronics device)
- Has become a “standard” sensor in robotics




Visual Navigation for Flying Robots

87

Dr. Jürgen Sturm, Computer Vision Group, TUM

Kinect: Applications



Visual Navigation for Flying Robots

88

Dr. Jürgen Sturm, Computer Vision Group, TUM

Open Research Questions

- How can RGB-D sensing facilitate in solving hard perception problems in robotics?
 - Interest points and feature descriptors?
 - Simultaneous localization and mapping?
 - Collision avoidance and visual navigation?
 - Object recognition and localization?
 - Human-robot interaction?
 - Semantic scene interpretation?

Visual Navigation for Flying Robots

89

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

Place Recognition, ICP, and Dense Reconstruction

Dr. Jürgen Sturm

Exercise Sheet 5

- Prepare mid-term presentation
- Proposed structure: 3 slides
 1. Remind people who you are and what you are doing (can be same slide as last time)
 2. Your work/achievements so far (video is a plus)
 3. Your plans for the next two weeks
- Hand in slides before July 3, 10am

Visual Navigation for Flying Robots

2

Dr. Jürgen Sturm, Computer Vision Group, TUM

Agenda for Today

- **Localization**
 - Visual place recognition
 - Scan matching and Iterative Closest Point
- **Mapping with known poses (3D reconstruction)**
 - Occupancy grids
 - Octrees
 - Signed distance field
 - Meshing


Visual Navigation for Flying Robots

3

Dr. Jürgen Sturm, Computer Vision Group, TUM

Remember: Loop Closures

- Use loop-closures to minimize the drift / minimize the error over all constraints




Visual Navigation for Flying Robots

4

Dr. Jürgen Sturm, Computer Vision Group, TUM

Loop Closures

How can we detect loop closures efficiently?



Visual Navigation for Flying Robots


5

Dr. Jürgen Sturm, Computer Vision Group, TUM

Loop Closures

How can we detect loop closures efficiently?

1. Compare with all previous images $O(n)$ (not efficient)



Visual Navigation for Flying Robots


6

Dr. Jürgen Sturm, Computer Vision Group, TUM

Loop Closures

How can we detect loop closures efficiently?

2. Use motion model and covariance to limit search radius (metric approach)



Visual Navigation for Flying Robots


7

Dr. Jürgen Sturm, Computer Vision Group, TUM

Loop Closures

How can we detect loop closures efficiently?

3. Appearance-based place recognition (using bag of words)




Visual Navigation for Flying Robots

8

Dr. Jürgen Sturm, Computer Vision Group, TUM

Appearance-based Place Recognition

Appearance can help to recover the pose estimate where metric approaches might fail



Visual Navigation for Flying Robots

9

Dr. Jürgen Sturm, Computer Vision Group, TUM

Analogy to Document Retrieval

Of all the sensory impressions proceeding to the brain, the visual experiences are the dominant ones. Our perception of the world around us is based essentially on visual impressions that reach the brain from the eyes. It was thought that the brain processes these impressions point by point, but Hubel and Wiesel have shown that the visual information is processed in the various layers of the visual cortex. The message about the image falls on the retina and undergoes a step-wise analysis in a system of nerve cells stored in columns. In this system each cell has its specific function and is responsible for a specific detail in the pattern of the retinal image.

sensory, brain, visual, perception, retinal, cerebral cortex, eye, cell, optical nerve, image Hubel, Wiesel

China is forecasting a trade surplus of \$90bn (£51bn) to \$100bn this year, a threefold increase on 2004's \$32bn. The Commerce Ministry said the surplus could be created by a predicted 30% increase in exports to \$750bn, compared with \$570bn in 2004. The figure, which has been widely questioned, is unfair, says the US government. China needed to increase the trade surplus by 2.1% in 2005. The dollar by 2.1% in 2005. The trade surplus within a narrow band, but the US government has to be allowed to trade freely. Beijing has made it clear that it will take the trade and tread carefully before allowing the yuan to rise further in value.

China, trade, surplus, commerce, exports, imports, US, yuan, bank, domestic, foreign, increase, trade, value

Visual Navigation for Flying Robots

10

Dr. Jürgen Sturm, Computer Vision Group, TUM

Object/Scene Recognition

- Analogy to documents: The content can be inferred from the frequency of visual words



object



bag of visual words

Visual Navigation for Flying Robots

11


Dr. Jürgen Sturm, Computer Vision Group, TUM

Bag of Visual Words

- Visual words = (independent) features



face



features

Visual Navigation for Flying Robots

12

Dr. Jürgen Sturm, Computer Vision Group, TUM

Bag of Visual Words

- Visual words = (independent) features
- Construct a dictionary of representative words

dictionary of visual words (codebook)




Visual Navigation for Flying Robots

13

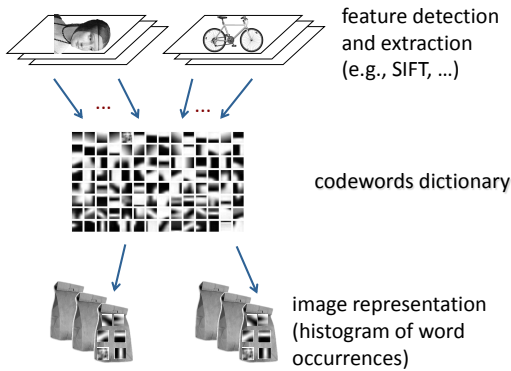
Dr. Jürgen Sturm, Computer Vision Group, TUM

Bag of Visual Words

- Visual words = (independent) features
- Construct a dictionary of representative words
- Represent the image based on a histogram of word occurrences (bag)



Overview




Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

Learning the Dictionary




Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Learning the Dictionary




Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

Learning the Dictionary



Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM

Learning the Visual Vocabulary




Visual Navigation for Flying Robots

19

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Image Representation

- Build the histogram by assigning each detected feature to the closest entry in the codebook



Visual Navigation for Flying Robots

20

Dr. Jürgen Sturm, Computer Vision Group, TUM


Object/Scene Recognition

- Compare histogram of new scene with those of known scenes, e.g., using

- simple histogram intersection

$$\text{score}(\mathbf{p}, \mathbf{q}) = \sum \min(p_i, q_i)$$

- naïve Bayes
- more advanced statistical methods




Visual Navigation for Flying Robots

21

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: FAB-MAP

[Cummins and Newman, 2008]




Visual Navigation for Flying Robots

22

Dr. Jürgen Sturm, Computer Vision Group, TUM

Timing Performance

- Inference: 25 ms for 100k locations
- SURF detection + quantization: 483 ms



Visual Navigation for Flying Robots

23

Dr. Jürgen Sturm, Computer Vision Group, TUM

Summary: Bag of Words

[Fei-Fei and Perona, 2005; Nister and Stewenius, 2006]

- Compact representation of content
- Highly efficient and scalable
- Requires training of a dictionary
- Insensitive to viewpoint changes/image deformations (inherited from feature descriptor)


Visual Navigation for Flying Robots

24

Dr. Jürgen Sturm, Computer Vision Group, TUM

Laser-based Motion Estimation

- So far, we looked at motion estimation (and place recognition) from **visual** sensors
- Today, we cover motion estimation from **range** sensors
 - Laser scanner (laser range finder, ultrasound)
 - Depth cameras (time-of-flight, Kinect ...)



Visual Navigation for Flying Robots

25


Dr. Jürgen Sturm, Computer Vision Group, TUM

Laser Scanner

- Measures angles and distances to closest obstacles

$$\mathbf{z} = (\theta_1, z_1, \dots, \theta_n, z_n) \in \mathbb{R}^{2n}$$
- Alternative representation: 2D point set (cloud)

$$\mathbf{z} = (x_1, y_1, \dots, x_n, y_n)^T \in \mathbb{R}^{2n}$$
- Probabilistic sensor model $p(z | x)$



Visual Navigation for Flying Robots

26

Dr. Jürgen Sturm, Computer Vision Group, TUM

Laser-based Motion Estimation

How can we best align two laser scans?

Visual Navigation for Flying Robots

27

Dr. Jürgen Sturm, Computer Vision Group, TUM

Laser-based Motion Estimation

How can we best align two laser scans?

- Exhaustive search
- Feature extraction (lines, corners, ...)
- Iterative minimization (ICP)

Visual Navigation for Flying Robots

28

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exhaustive Search

- Convolve first scan with sensor model
- Sweep second scan over likelihood map, compute correlation and select best pose




Visual Navigation for Flying Robots

29

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Exhaustive Search [Olson, '09]

- Multi-resolution correlative scan matching
- Real-time by using GPU
- Remember: SE(2) has 3 DOFs




Visual Navigation for Flying Robots

30

Dr. Jürgen Sturm, Computer Vision Group, TUM

Does Exhaustive Search Generalize To 3D As Well?



Visual Navigation for Flying Robots

31

Dr. Jürgen Sturm, Computer Vision Group, TUM

Iterative Closest Point (ICP)

- **Given:** Two corresponding point sets (clouds)

$$P = \{p_1, \dots, p_n\}$$

$$Q = \{q_1, \dots, q_n\}$$

- **Wanted:** Translation t and rotation R that minimize the sum of the squared error

$$E(R, t) = \frac{1}{n} \sum_{i=1}^n \|p_i - Rq_i - t\|^2$$

where p_i and q_i are corresponding points


Visual Navigation for Flying Robots

32

Dr. Jürgen Sturm, Computer Vision Group, TUM

Known Correspondences

Note: If the correct correspondences are known, both rotation and translation can be calculated in **closed form**.



Visual Navigation for Flying Robots

33

Dr. Jürgen Sturm, Computer Vision Group, TUM

Known Correspondences

- **Idea:** The center of mass of both point sets has to match

$$\bar{p} = \frac{1}{n} \sum_i p_i \quad \bar{q} = \frac{1}{n} \sum_i q_i$$

- Subtract the corresponding center of mass from every point
- Afterwards, the point sets are zero-centered, i.e., we only need to recover the rotation...

Visual Navigation for Flying Robots

34

Dr. Jürgen Sturm, Computer Vision Group, TUM

Known Correspondences

- Decompose the matrix

$$W = \sum_i (p_i - \bar{p})(q_i - \bar{q})^T = USV^T$$

using singular value decomposition (SVD)

- **Theorem**

If $\text{rank } W = 3$, the optimal solution of $E(R, t)$ is unique and given by

$$R = UV^T$$

$$t = \bar{p} - R\bar{q}$$

(for proof, see <http://hss.ulb.uni-bonn.de/2006/0912/0912.pdf>, p.34/35)


Visual Navigation for Flying Robots

35

Dr. Jürgen Sturm, Computer Vision Group, TUM

Unknown Correspondences

- If the correct correspondences are not known, it is generally impossible to determine the optimal transformation in one step



Visual Navigation for Flying Robots

36

Dr. Jürgen Sturm, Computer Vision Group, TUM

ICP Algorithm

[Besl & McKay, 92]

- **Algorithm:** Iterate until convergence
 - Find correspondences
 - Solve for R, t
- Converges if starting position is “close enough”




Visual Navigation for Flying Robots

37

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: ICP



Visual Navigation for Flying Robots

38

Dr. Jürgen Sturm, Computer Vision Group, TUM

ICP Variants

Many variants on all stages of ICP have been proposed:

- **Selecting** and **weighting** source points
- **Finding** corresponding points
- Rejecting certain (outlier) correspondences
- Choosing an **error metric**
- **Minimization**

Visual Navigation for Flying Robots

39

Dr. Jürgen Sturm, Computer Vision Group, TUM

Performance Criteria

- Various aspects of performance
 - Speed
 - Stability (local minima)
 - Tolerance w.r.t. noise and/or outliers
 - Basin of convergence (maximum initial misalignment)
- Choice depends on data and application

Visual Navigation for Flying Robots

40

Dr. Jürgen Sturm, Computer Vision Group, TUM

Selecting Source Points

- Use all points
- Uniform sub-sampling
- Random sampling
- Feature-based sampling
- Normal-space sampling
 - Ensure that samples have normals distributed as uniformly as possible


Visual Navigation for Flying Robots

41

Dr. Jürgen Sturm, Computer Vision Group, TUM

Spatially Uniform Sampling

- Density of points usually depends on the distance to the sensor → no uniform distribution
- Can lead to a bias in ICP



Visual Navigation for Flying Robots


42

Dr. Jürgen Sturm, Computer Vision Group, TUM

Feature-based Sampling

Detect interest points (same as with images)


- Decrease the number of correspondences
- Increase efficiency and accuracy
- Requires pre-processing



3D Scan (~200.000 Points)

Visual Navigation for Flying Robots


43




Extracted Features (~5.000 Points)

Dr. Jürgen Sturm, Computer Vision Group, TUM

Normal-Space Sampling



Uniform sampling



Normal-space sampling


Visual Navigation for Flying Robots

44

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Normal-Space Sampling


Normal-space sampling can help on mostly-smooth areas with sparse features



Random sampling

Visual Navigation for Flying Robots

45



Normal-space sampling

Dr. Jürgen Sturm, Computer Vision Group, TUM

Selection and Weighting

- Selection is a form of (binary) weighting
- Instead of re-sampling one can also use weighting
- Weighting strategy depends on the data
- Pre-processing / run-time trade-off

Visual Navigation for Flying Robots

46

Dr. Jürgen Sturm, Computer Vision Group, TUM

Finding Correspondences

Has greatest effect on convergence and speed

- Closest point
- Normal shooting
- Closest compatible point
- Projection
- Speed-up using kd-trees (or oct-trees)


Visual Navigation for Flying Robots

47

Dr. Jürgen Sturm, Computer Vision Group, TUM

Closest Point Matching

- Find closest point in the other point set
- Distance threshold



- Closest-point matching generally stable, but slow and requires pre-processing


Visual Navigation for Flying Robots

48

Dr. Jürgen Sturm, Computer Vision Group, TUM

Normal Shooting

- Project along normal, intersect other mesh



- Slightly better than closest point for smooth meshes, worse for noisy or complex meshes

Visual Navigation for Flying Robots

49

Dr. Jürgen Sturm, Computer Vision Group, TUM

Closest Compatible Point

- Can improve effectiveness of both the previous variants by only matching to **compatible** points
- Compatibility based on normals, colors, ...
- In the limit, degenerates to feature matching

Visual Navigation for Flying Robots

50

Dr. Jürgen Sturm, Computer Vision Group, TUM

Speeding Up Correspondence Search

Finding closest point is most expensive stage of the ICP algorithm

- Build index for one point set (kd-tree)
- Use simpler algorithm (e.g., projection-based matching)


Visual Navigation for Flying Robots

51

Dr. Jürgen Sturm, Computer Vision Group, TUM

Projection-based Matching

- Slightly worse performance per iteration
- Each iteration is one to two orders of magnitude faster than closest-point
- Requires point-to-plane error metric





Visual Navigation for Flying Robots

52

Dr. Jürgen Sturm, Computer Vision Group, TUM

Error Metrics

- Point-to-point
- Point-to-plane lets flat regions slide along each other



- Generalized ICP: Assign individual covariance to each data point [Segal, 2009]

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Minimization


- Only point-to-point metric has closed form solution(s)
- Other error metrics require non-linear minimization methods
 - Which non-linear minimization methods do you remember?
 - Which robust error metrics do you remember?

Visual Navigation for Flying Robots

54

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robust Error Metrics



Visual Navigation for Flying Robots

55

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Real-Time ICP on Range Images

[Rusinkiewicz and Levoy, 2001]

- Real-time scan alignment
- Range images from structure light system (projector and camera, temporal coding)



Visual Navigation for Flying Robots

56

Dr. Jürgen Sturm, Computer Vision Group, TUM

ICP: Summary

- ICP is a powerful algorithm for calculating the displacement between point clouds
- The overall speed depends most on the choice of matching algorithm
- ICP is (in general) only locally optimal → can get stuck in local minima

Visual Navigation for Flying Robots

57

Dr. Jürgen Sturm, Computer Vision Group, TUM

Agenda for Today

- Localization
 - Visual place recognition
 - Scan matching and Iterative Closest Point
- Mapping with known poses (3D reconstruction)
 - Occupancy grids
 - Octrees
 - Signed distance field
 - Meshing

Visual Navigation for Flying Robots

58

Dr. Jürgen Sturm, Computer Vision Group, TUM

Occupancy Grid

Idea:

- Represent the map \mathbf{m} using a grid
- Each cell is either free or occupied



$$\mathbf{m} = (m_1, \dots, m_n) \in \{\text{empty}, \text{occ}\}^n$$

- Robot maintains a belief $\text{Bel}(\mathbf{m})$ on map state

Goal: Estimate the belief from sensor observations

$$\text{Bel}(\mathbf{m}) = P(\mathbf{m} \mid \mathbf{z}_1, \dots, \mathbf{z}_t)$$

Visual Navigation for Flying Robots

59

Dr. Jürgen Sturm, Computer Vision Group, TUM

Occupancy Grid - Assumptions

- Map is static
- Cells have binary state (empty or occupied)
- All cells are independent of each other
- As a result, each cell m_i can be estimated independently from the sensor observations
- Will also drop index i (for the moment)

Visual Navigation for Flying Robots

60

Dr. Jürgen Sturm, Computer Vision Group, TUM

Mapping

- **Goal:** Estimate

$$\text{Bel}(m) = P(m \mid z_1, \dots, z_n)$$

- How can this be computed?

Visual Navigation for Flying Robots

61

Dr. Jürgen Sturm, Computer Vision Group, TUM

Binary Bayes Filter

- **Goal:** Estimate

$$\text{Bel}(m) = P(m \mid z_1, \dots, z_n)$$

- How can this be computed?

- E.g., using the Bayes Filter from Lecture 3

$$P(m \mid z_{1:t}) = \left(1 + \frac{1 - P(m \mid z_t)}{P(m \mid z_t)} \frac{1 - P(m \mid z_{1:t-1})}{P(m \mid z_{1:t-1})} \frac{P(m)}{1 - P(m)} \right)^{-1}$$

Visual Navigation for Flying Robots

62

Dr. Jürgen Sturm, Computer Vision Group, TUM

Binary Bayes Filter

- **Prior probability** that cell is occupied $P(m)$ (often 0.5)
- **Inverse sensor model** $P(m \mid z_t)$ is specific to the sensor used for mapping
- The **log-odds representation** can be used to increase speed and numerical stability

$$L(x) := \log \frac{p(x)}{p(\neg x)} = \log \frac{p(x)}{1 - p(x)}$$

Visual Navigation for Flying Robots

63

Dr. Jürgen Sturm, Computer Vision Group, TUM

Binary Bayes Filter using Log-Odds

- In each time step, compute

$$L(m \mid z_{1:t}) = \overset{\text{previous belief}}{L(m \mid z_{1:t-1})} + \overset{\text{inverse sensor model}}{L(m \mid z_t)} + \overset{\text{map prior}}{L(m)}$$

- When needed, compute current belief as

$$\text{Bel}_t(m) = 1 - \frac{1}{1 + \exp L(m \mid z_{1:t})}$$

Visual Navigation for Flying Robots

64

Dr. Jürgen Sturm, Computer Vision Group, TUM

Clamping Update Policy

- Often, the world is not “fully” static
- Consider an appearing/disappearing obstacle
- To change the state of a cell, the filter needs as many positive (negative) observations
- **Idea:** Clamp the beliefs to min/max values

$$L'(m \mid z_{1:t}) = \max(\min(L(m \mid z_{1:t}), l_{\max}), l_{\min})$$

Visual Navigation for Flying Robots

65

Dr. Jürgen Sturm, Computer Vision Group, TUM

Sensor Model

- For the Bayes filter, we need the inverse sensor model

$$p(m \mid z)$$

- Let's consider an ultrasound sensor
 - Located at (0,0)
 - Measures distance of 2.5m
 - How does the inverse sensor model look like?


Visual Navigation for Flying Robots

66

Dr. Jürgen Sturm, Computer Vision Group, TUM

Typical Sensor Model for Ultrasound

- Combination of a linear function (in x-direction) and a Gaussian (in y-direction)



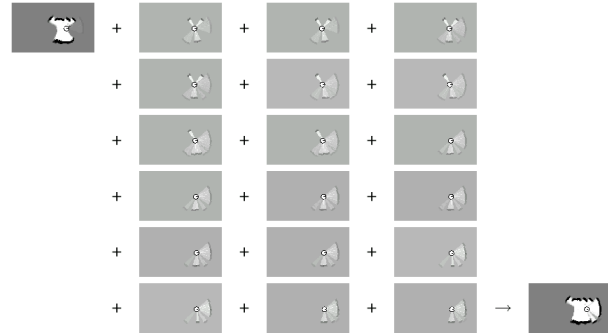
- Question: What about a laser scanner?

Visual Navigation for Flying Robots

67

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Updating the Occupancy Grid



Visual Navigation for Flying Robots

68

Dr. Jürgen Sturm, Computer Vision Group, TUM

Resulting Map



Note: The maximum likelihood map is obtained by clipping the occupancy grid map at a threshold of 0.5

Visual Navigation for Flying Robots

69

Dr. Jürgen Sturm, Computer Vision Group, TUM

Memory Consumption

- Consider we want to map a building with 40x40m at a resolution of 0.05cm
- How much memory do we need?

Visual Navigation for Flying Robots

70

Dr. Jürgen Sturm, Computer Vision Group, TUM

Memory Consumption

- Consider we want to map a building with 40x40m at a resolution of 0.05cm
- How much memory do we need?

$$\left(\frac{40}{0.05}\right)^2 = 640.000 \text{ cells} = 4.88\text{mb}$$

- And for 3D?

$$\left(\frac{40}{0.05}\right)^3 = 512.000.000 \text{ cells} = 3.8\text{gb}$$

- And what about a whole city?



Visual Navigation for Flying Robots

71

Dr. Jürgen Sturm, Computer Vision Group, TUM

Map Representation by Octrees

- Tree-based data structure
- Recursive subdivision of space into octants
- Volumes can be allocated as needed
- Multi-resolution



Visual Navigation for Flying Robots


72

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: OctoMap

[Wurm et al., 2011]

- Freiburg, building 79
44 x 18 x 3 m³, 0.05m resolution, 0.7mb on disk



Visual Navigation for Flying Robots


73

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: OctoMap

[Wurm et al., 2011]

- Freiburg computer science campus
292 x 167 x 28 m³, 0.2m resolution, 2mb on disk



Visual Navigation for Flying Robots


74

Dr. Jürgen Sturm, Computer Vision Group, TUM


Signed Distance Field (SDF)

[Curless and Levoy, 1996]

- Idea:** Instead of representing the cell occupancy, represent the distance of each cell to the surface
- Occupancy grid maps: explicit representation



- SDF: implicit representation



Visual Navigation for Flying Robots

75


Dr. Jürgen Sturm, Computer Vision Group, TUM

Signed Distance Field (SDF)

[Curless and Levoy, 1996]

Algorithm:

- Estimate the signed distance field
- Extract the surface using interpolation (surface is located at zero-crossing)




Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM

Weighting Function

- Weight each observation according to its confidence



- Weight can additionally be influenced by other modalities (reflectance values, ...)

Visual Navigation for Flying Robots

77

Dr. Jürgen Sturm, Computer Vision Group, TUM

Data Fusion

- Each voxel cell x in the SDF stores two values
 - Weighted sum of signed distances $D_t(x)$
 - Sum of all weights $W_t(x)$
- When new range image arrives, update every voxel cell according to

$$D_{t+1}(x) = D_t(x) + w_{t+1}(x)d_{t+1}(x)$$

$$W_{t+1}(x) = W_t(x) + w_{t+1}(x)$$


Visual Navigation for Flying Robots

78

Dr. Jürgen Sturm, Computer Vision Group, TUM

Two Nice Properties

- Noise cancels out over multiple measurements



- Zero-crossing can be extracted at sub-voxel accuracy (least squares estimate)

1D Example:
$$x^* = \frac{\sum D_t(x)}{\sum W_t(x)}$$


Visual Navigation for Flying Robots

79


Dr. Jürgen Sturm, Computer Vision Group, TUM

SDF Example

A cross section through a 3D signed distance function of a real scene



Surface with cross-section




SDF

Visual Navigation for Flying Robots

80

Dr. Jürgen Sturm, Computer Vision Group, TUM

SDF Fusion



Visual Navigation for Flying Robots

81

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visualizing Signed Distance Fields

Common approaches to iso surface extraction:

- Ray casting (GPU, fast)
 - For each camera pixel, shoot a ray and search for zero crossing
- Poligonization (CPU, slow)
 - E.g., using the marching cubes algorithm
 - Advantage: outputs triangle mesh



Visual Navigation for Flying Robots

82

Dr. Jürgen Sturm, Computer Vision Group, TUM

Ray Casting

- For each camera pixel, shoot a ray and search for the first zero crossing in the SDF
- Value in the SDF can be used to skip along when far from surface





Visual Navigation for Flying Robots

83

Dr. Jürgen Sturm, Computer Vision Group, TUM

Ray Casting

- Interpolation reduces artifacts
- Close to surface, gradient represents the surface normal



Visual Navigation for Flying Robots


84

Dr. Jürgen Sturm, Computer Vision Group, TUM


Marching Cubes

First in 2D, **marching squares**:

- Evaluate each cell separately
- Check which edges are inside/outside
- Generate triangles according to lookup table
- Locate vertices using least squares




Visual Navigation for Flying Robots



85

Dr. Jürgen Sturm, Computer Vision Group, TUM

Marching Cubes



Visual Navigation for Flying Robots


86

Dr. Jürgen Sturm, Computer Vision Group, TUM

KinectFusion

[Newcombe et al., 2011]

- Projective ICP with point-to-plane metric
- Truncated signed distance function (TSDF)
- Ray Casting



Visual Navigation for Flying Robots

87

Dr. Jürgen Sturm, Computer Vision Group, TUM

Lessons Learned Today

- How to quickly recognize previously seen places
- How to align point clouds
- How to estimate occupancy maps
- How to reconstruct triangle meshes at sub-voxel accuracy

Visual Navigation for Flying Robots

88

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

Motion Planning

Dr. Jürgen Sturm

in.tum.summer party & career forum

The Department of Informatics would like to invite its students and employees to its summer party and career forum.

July 4, 2012

3 pm – 6 pm **Career Forum:**


Presentations given by Google, Capgemini etc, stands, panel discussion: TUM alumni talk about their career paths in informatics

3 pm – 6 pm **Foosball Tournament**

Starting at 5 pm **Summer Party:**
BBQ, live band and lots of fun!

www.in.tum.de/2012summerparty

Motivation: Flying Through Forests




Visual Navigation for Flying Robots

3

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Planning Problem

- Given obstacles, a robot, and its motion capabilities, compute collision-free robot motions from the start to goal.



Visual Navigation for Flying Robots

4

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Planning Problem

What are good performance metrics?

Visual Navigation for Flying Robots

5

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Planning Problem

What are good performance metrics?

- Execution speed / path length
- Energy consumption
- Planning speed
- Safety (minimum distance to obstacles)
- Robustness against disturbances
- Probability of success
- ...


Visual Navigation for Flying Robots

6

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Planning Examples

Motion planning is sometimes also called the **piano mover's problem**

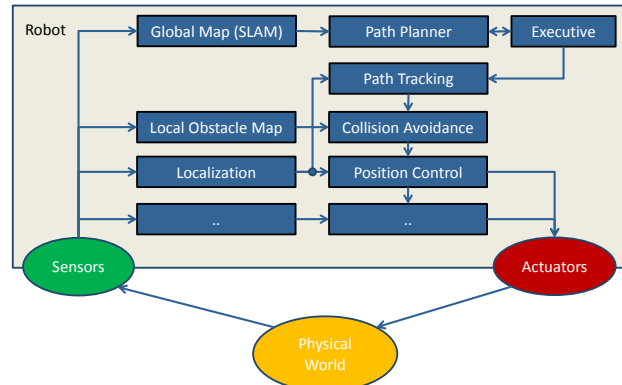


Visual Navigation for Flying Robots

7

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robot Architecture



Visual Navigation for Flying Robots

8

Dr. Jürgen Sturm, Computer Vision Group, TUM

Agenda for Today

- Configuration spaces
- Roadmap construction
- Search algorithms
- Path optimization and re-planning
- Path execution

Visual Navigation for Flying Robots

9

Dr. Jürgen Sturm, Computer Vision Group, TUM

Configuration Space

- Work space
 - Typically 3D pose (position + orientation) \rightarrow 6 DOF
- Configuration space
 - Reduced pose (position + yaw) \rightarrow 4 DOF
 - Full pose \rightarrow 6 DOF
 - Pose + velocity \rightarrow 12 DOF
 - Joint angles of manipulation robot
 - ...
- Planning takes place in **configuration space**


Visual Navigation for Flying Robots

10

Dr. Jürgen Sturm, Computer Vision Group, TUM

Configuration Space

- The configuration space (C-space) is the **space of all possible configurations**
- C-space topology is usually not Cartesian
- C-space is described as a topological manifold



Visual Navigation for Flying Robots

11

Dr. Jürgen Sturm, Computer Vision Group, TUM

Notation

- Configuration space $C \subset \mathbb{R}^d$
- Configuration $q \in C$
- Free space C_{free}
- Obstacle space C_{obs}

- Properties

$$C_{\text{free}} \cup C_{\text{obs}} = C$$

$$C_{\text{free}} \cap C_{\text{obs}} = \emptyset$$


Visual Navigation for Flying Robots

12

Dr. Jürgen Sturm, Computer Vision Group, TUM

Free Space Example

- What are admissible configurations for the robot? Equiv.: What is the free space?
- "Point" robot




Visual Navigation for Flying Robots

13

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- What are admissible configurations for the robot? Equiv.: What is the free space?
- "Point" robot




Visual Navigation for Flying Robots

14

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- What are admissible configurations for the robot? Equiv.: What is the free space?
- Circular robot




Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- What are admissible configurations for the robot? Equiv.: What is the free space?
- Circular robot




Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- What are admissible configurations for the robot? Equiv.: What is the free space?
- Large circular robot



Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

Computing the Free Space

- Free configuration space is obtained by sliding the robot along the edge of the obstacle regions "blowing them up" by the robot radius
- This operation is called the **Minowski sum**

$$A \oplus B = \{a + b \mid a \in A, b \in B\}$$

where $A, B \subset \mathbb{R}^d$


Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Minkowski Sum

- Triangular robot and rectangular obstacle




Visual Navigation for Flying Robots

19

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Polygonal robot, translation only



- C-space is obtained by sliding the robot along the edge of the obstacle regions

Visual Navigation for Flying Robots


20

Dr. Jürgen Sturm, Computer Vision Group, TUM

Basic Motion Planning Problem

Given

- Free space C_{free}
- Initial configuration q_I
- Goal configuration q_G



- Goal:** Find a continuous path

$$\tau : [0, 1] \rightarrow C_{\text{free}}$$

$$\text{with } \tau(0) = q_I, \tau(1) = q_G$$

Visual Navigation for Flying Robots

21

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Planning Sub-Problems

- C-Space discretization**
(generating a graph / roadmap)
- Search algorithm
(Dijkstra's algorithm, A*, ...)
- Re-planning
(D*, ...)
- Path tracking
(PID control, potential fields, funnels, ...)

Visual Navigation for Flying Robots

22

Dr. Jürgen Sturm, Computer Vision Group, TUM

C-Space Discretizations

Two competing paradigms

- Combinatorial planning**
(exact planning)
- Sampling-based planning**
(probabilistic/randomized planning)

Visual Navigation for Flying Robots

23

Dr. Jürgen Sturm, Computer Vision Group, TUM

Combinatorial Methods

- Mostly developed in the 1980s
- Extremely efficient for low-dimensional problems
- Sometimes difficult to implement
- Usually produce a road map in C_{free}
- Assume polygonal environments

Visual Navigation for Flying Robots


24

Dr. Jürgen Sturm, Computer Vision Group, TUM

Roadmaps

A **roadmap** is a graph in C_{free} where

- Each vertex is a configuration $q \in C_{\text{free}}$
- Each edge is a path $\tau : [0, 1] \rightarrow C_{\text{free}}$ for which $\tau(0)$ and $\tau(1)$ are vertices



Visual Navigation for Flying Robots

25

Dr. Jürgen Sturm, Computer Vision Group, TUM

(Desired) Properties of Roadmaps

▪ Accessibility

From anywhere in C_{free} , it is easy to compute a path that reaches at least one of the vertices

▪ Connectivity-preserving

If there exists a path between q_I and q_G in C_{free} then there must also exist a path in the road map



Visual Navigation for Flying Robots

26

Dr. Jürgen Sturm, Computer Vision Group, TUM

Roadmap Construction

We consider here three **combinatorial** methods:

- Trapezoidal decomposition
- Shortest path roadmap
- Regular grid
- ... but there are many more!

Afterwards, we consider two **sampling-based** methods:

- Probabilistic roadmaps (PRMs)
- Rapidly exploring random trees (RRTs)


Visual Navigation for Flying Robots

27

Dr. Jürgen Sturm, Computer Vision Group, TUM

Roadmap Construction

- Decompose horizontally in convex regions using plane sweep
- Sort vertices in x direction. Iterate over vertices while maintaining a vertically sorted list of edges




Visual Navigation for Flying Robots

28

Dr. Jürgen Sturm, Computer Vision Group, TUM

Roadmap Construction

- Place vertices
 - in the center of each trapezoid
 - on the edge between two neighboring trapezoids
- Resulting road map



Quick check on properties:
- Accessibility
- Connectivity-preserving?

Visual Navigation for Flying Robots


29

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Query

Compute path from q_I to q_G

- Identify **start** and **goal** trapezoid
- Connect **start** and **goal** location to center vertex
- Run search algorithm (e.g., Dijkstra)




Visual Navigation for Flying Robots

30

Dr. Jürgen Sturm, Computer Vision Group, TUM

Properties of Trapezoidal Decomposition

- + Easy to implement
- + Efficient computation
- + Scales to 3D
- Does not generate shortest path




Visual Navigation for Flying Robots

31

Dr. Jürgen Sturm, Computer Vision Group, TUM

Shortest-Path Roadmap

- Contains all vertices and edges that optimal paths follow when obstructed
- Imagine pulling a tight string between q_I and q_G




Visual Navigation for Flying Robots

32

Dr. Jürgen Sturm, Computer Vision Group, TUM

Roadmap Construction

- Vertices = all sharp corners (>180deg, red)
- Edges
 1. Two consecutive sharp corners on the same obstacle (light blue)
 2. Bitangent edges (when line connecting two vertices extends into free space, dark blue)



Visual Navigation for Flying Robots


33

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Query

Compute path from q_I to q_G

- Connect **start** and **goal** location to all visible roadmap vertices
- Run search algorithm (e.g., Dijkstra)




Visual Navigation for Flying Robots

34

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Query

- + Easy to construct in 2D
- + Generates shortest paths
- Optimal planning in 3D or more dim. is NP-hard




Visual Navigation for Flying Robots

35

Dr. Jürgen Sturm, Computer Vision Group, TUM

Approximate Decompositions

- Construct a regular grid
- High memory consumption (and number of tests)
- Any ideas?




Visual Navigation for Flying Robots

36

Dr. Jürgen Sturm, Computer Vision Group, TUM

Approximate Decompositions

- Construct a regular grid
- Use quadtree/octtree to save memory
- Sometimes difficult to determine status of cell




Visual Navigation for Flying Robots

37

Dr. Jürgen Sturmf, Computer Vision Group, TUM

Approximate Decompositions

- + Easy to construct
- High number of tests
- + Most used in practice



Visual Navigation for Flying Robots

38

Dr. Jürgen Sturmf, Computer Vision Group, TUM

Summary: Combinatorial Planning

- **Pro:** Find a solution when one exists (complete)
- **Con:** Become quickly intractable for higher dimensions
- **Alternative:** Sampling-based planning
Weaker guarantees but more efficient

Visual Navigation for Flying Robots

39

Dr. Jürgen Sturmf, Computer Vision Group, TUM

Sampling-based Methods

- Abandon the concept of explicitly characterizing C_{free} and C_{obs} and leave the algorithm **in the dark** when exploring C_{free}
- The only light is provided by a **collision-detection algorithm** that probes C to see whether some configuration lies in C_{free}
- We will have a look at
 - Probabilistic road maps (PRMs)
 - Rapidly exploring random trees (RRTs)

Visual Navigation for Flying Robots

40

Dr. Jürgen Sturmf, Computer Vision Group, TUM

Probabilistic Roadmaps (PRMs)

[Kavraki et al., 1992]

- **Vertex:** Take random sample from C , check whether sample is in C_{free}
- **Edge:** Check whether line-of-sight between two nearby vertices is collision-free
- Options for “nearby”: k-nearest neighbors or all neighbors within specified radius
- Add vertices and edges until roadmap is dense enough



Visual Navigation for Flying Robots

41

Dr. Jürgen Sturmf, Computer Vision Group, TUM

PRM Example

1. Sample vertex
2. Find neighbors
3. Add edges



Step 3: Check edges for collisions, e.g., using discretized line search



Visual Navigation for Flying Robots

42

Dr. Jürgen Sturmf, Computer Vision Group, TUM

Probabilistic Roadmaps

- + Probabilistic. complete
- + Scale well to higher dimensional C-spaces
- + Very popular, many extensions
- Do not work well for some problems (e.g., narrow passages)
- Not optimal, not complete



Visual Navigation for Flying Robots


43

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rapidly Exploring Random Trees

[Lavalle and Kuffner, 1999]

- **Idea:** Grow tree from start to goal location



Visual Navigation for Flying Robots

44

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rapidly Exploring Random Trees

Algorithm

1. Initialize tree with first node q_I
2. Pick a random target location (every 100th iteration, choose q_G)
3. Find closest vertex in roadmap
4. Extend this vertex towards target location
5. Repeat steps until goal is reached

- Why not pick q_G every time?

Visual Navigation for Flying Robots

45

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rapidly Exploring Random Trees

Algorithm

1. Initialize tree with first node q_I
2. Pick a random target location (every 100th iteration, choose q_G)
3. Find closest vertex in roadmap
4. Extend this vertex towards target location
5. Repeat steps until goal is reached

- Why not pick q_G every time?

- This will fail and run into C_{obs} instead of exploring

Visual Navigation for Flying Robots


46

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rapidly Exploring Random Trees

[Lavalle and Kuffner, 1999]

- **RRT:** Grow trees from start and goal location towards each other, stop when they connect




Visual Navigation for Flying Robots

47


Dr. Jürgen Sturm, Computer Vision Group, TUM

RRT Examples

- 2-DOF example



- 3-DOF example (2D translation + rotation)




Visual Navigation for Flying Robots

48

Dr. Jürgen Sturm, Computer Vision Group, TUM

Non-Holonomic Robots

- Some robots cannot move freely on the configuration space manifold
- Example: A car can not move sideways
 - 2-DOF controls (speed and steering)
 - 3-DOF configuration space (2D translation + rotation)




Visual Navigation for Flying Robots

49

Dr. Jürgen Sturm, Computer Vision Group, TUM

Non-Holonomic Robots

- RRTs can naturally consider such constraints during tree construction
- Example: Car-like robot



Visual Navigation for Flying Robots

50

Dr. Jürgen Sturm, Computer Vision Group, TUM

Rapidly Exploring Random Trees

- + Probabilistic. complete
- + Balance between greedy search and exploration
- + Very popular, many extensions
- Metric sensitivity
- Unknown rate of convergence
- Not optimal, not complete



Visual Navigation for Flying Robots

51

Dr. Jürgen Sturm, Computer Vision Group, TUM

Summary: Sampling-based Planning

- More efficient** in most **practical problems** but offer weaker guarantees
- Probabilistically complete** (given enough time it finds a solution if one exists, otherwise, it may run forever)
- Performance degrades in problems with **narrow passages**

Visual Navigation for Flying Robots

52

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Planning Sub-Problems

- C-Space discretization (generating a graph / roadmap)
- Search algorithms** (Dijkstra's algorithm, A*, ...)
- Re-planning** (D*, ...)
- Path tracking (PID control, potential fields, funnels, ...)

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Search Algorithms

- Given:** Graph G consisting of vertices and edges (with associated costs)
- Wanted:** find the best (shortest) path between two vertices
- What search algorithms do you know?

Visual Navigation for Flying Robots


54

Dr. Jürgen Sturm, Computer Vision Group, TUM

Uninformed Search


Breadth-first

- Complete
- Optimal if action costs equal
- Time and space $O(b^d)$



Depth-first

- Not complete in infinite spaces
- Not optimal
- Time $O(b^d)$
- Space $O(bd)$
(can forget explored subtrees)




Visual Navigation for Flying Robots

55

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Dijkstra's Algorithm

- Extension of breadth-first with arbitrary (non-negative) costs



www.combinatorica.com

Visual Navigation for Flying Robots

56

Dr. Jürgen Sturm, Computer Vision Group, TUM

Informed Search

Idea

- Select nodes for further expansion based on an evaluation function $f(n)$
- First explore the node with lowest value
- What is a good evaluation function?

Visual Navigation for Flying Robots

57

Dr. Jürgen Sturm, Computer Vision Group, TUM

Informed Search

Idea

- Select nodes for further expansion based on an evaluation function $f(n)$
- First explore the node with lowest value
- What is a good evaluation function?
- Often a combination of
 - Path cost so far $g(n)$
 - Heuristic function $h(n)$
(e.g., estimated distance to goal, but can also encode additional domain knowledge)

Visual Navigation for Flying Robots

58

Dr. Jürgen Sturm, Computer Vision Group, TUM

Informed Search

Greedy best-first search

- Simply expand the node closest to the goal
 $f(n) = h(n)$
- Not optimal, not complete

A* search

- Combines path cost with estimated goal distance
 $f(n) = g(n) + h(n)$
- Optimal and complete (if $h(n)$ never overestimates actual cost)

Visual Navigation for Flying Robots

59

Dr. Jürgen Sturm, Computer Vision Group, TUM

What is a Good Heuristic Function?

- Choice is problem/application-specific
- Two popular choices
 - Manhattan distance (neglecting obstacles)
 - Euclidean distance (neglecting obstacles)
 - Value iteration / Dijkstra (from the goal backwards)

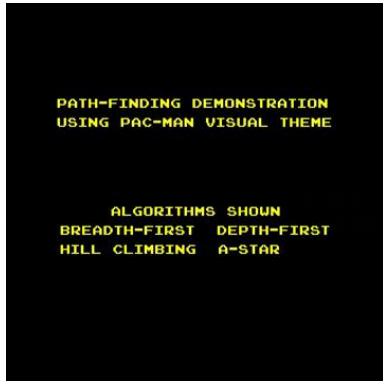


Visual Navigation for Flying Robots

60

Dr. Jürgen Sturm, Computer Vision Group, TUM

Comparison Search Algorithms



Visual Navigation for Flying Robots

61

Dr. Jürgen Sturm, Computer Vision Group, TUM

Problems on A* on Grids

1. The shortest path is often very **close to obstacles** (cutting corners)
 - Uncertain path execution increases the risk of collisions
 - Uncertainty can come from delocalized robot, imperfect map, or poorly modeled dynamic constraints
2. Trajectories are **aligned to grid** structure
 - Path looks unnatural
 - Paths are longer than the true shortest path in continuous space

Visual Navigation for Flying Robots

62

Dr. Jürgen Sturm, Computer Vision Group, TUM

Problems on A* on Grids

3. When the path turns out to be blocked during traversal, it needs to be **re-planned from scratch**
 - In unknown or dynamic environments, this can occur very often
 - Replanning in large state spaces is costly
 - Can we re-use (repair) the initial plan?

Let's look at solutions to these problems...

Visual Navigation for Flying Robots

63

Dr. Jürgen Sturm, Computer Vision Group, TUM

Map Smoothing

- **Problem:** Path gets close to obstacles
- **Solution:** Convolve the map with a kernel (e.g., Gaussian)



- Leads to non-zero probability around obstacles
- Evaluation function



$$f(n) = g(n) \cdot p_{occ}(n) + h(n)$$

Visual Navigation for Flying Robots

64

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Map Smoothing



Visual Navigation for Flying Robots

65

Dr. Jürgen Sturm, Computer Vision Group, TUM

Path Smoothing

- **Problem:** Paths are aligned to grid structure (because they have to lie in the roadmap)
- Paths look unnatural and are sub-optimal
- **Solution:** Smooth the path after generation
 - Traverse path and find pairs of nodes with direct line of sight; replace by line segment
 - Refine initial path using non-linear minimization (e.g., optimize for continuity/energy/execution time)
 - ...

Visual Navigation for Flying Robots

66


Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Path Smoothing

- Replace pairs of nodes by line segments



- Non-linear optimization



Visual Navigation for Flying Robots

67

Dr. Jürgen Sturm, Computer Vision Group, TUM

D* Search

- Problem:** In unknown, partially known or dynamic environments, the planned path may be blocked and we need to **replan**
- Can this be done efficiently, avoiding to replan the **entire path**?

Visual Navigation for Flying Robots

68

Dr. Jürgen Sturm, Computer Vision Group, TUM

D* Search

- Idea:** Incrementally repair path keeping its modifications local around robot pose
- Many variants:
 - D* (Dynamic A*) [Stentz, ICRA '94] [Stentz, IJCAI '95]
 - D* Lite [Koenig and Likhachev, AAAI '02]
 - Field D* [Ferguson and Stenz, JFR '06]

Visual Navigation for Flying Robots

69

Dr. Jürgen Sturm, Computer Vision Group, TUM

D* Search

Main concepts

- Invert search direction** (from goal to start)
 - Goal does not move, but robot does
 - Map changes (new obstacles) have only local influence close to current robot pose
- Mark** the changed node and all dependent nodes **as unclean** (=to be re-evaluated)
- Find shortest path** to start (using A*) while **re-using previous solution**


Visual Navigation for Flying Robots

70

Dr. Jürgen Sturm, Computer Vision Group, TUM

D* Example

- Situation at start



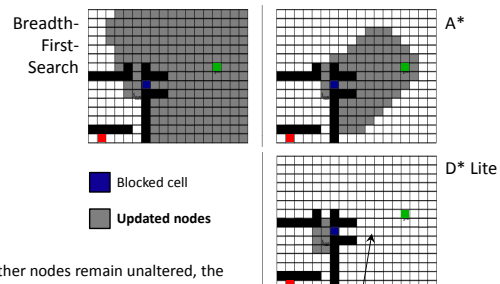
Visual Navigation for Flying Robots

71

Dr. Jürgen Sturm, Computer Vision Group, TUM

D* Example

- After discovery of blocked cell




Visual Navigation for Flying Robots

72

Dr. Jürgen Sturm, Computer Vision Group, TUM

D* Search

- D* is as optimal and complete as A*
- D* and its variants are widely used in practice
- Field D* was running on Mars rovers Spirit and Opportunity



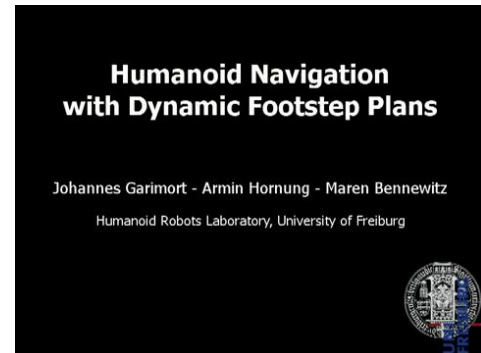
Visual Navigation for Flying Robots

73

Dr. Jürgen Sturm, Computer Vision Group, TUM

D* Lite for Footstep Planning

[Garimort et al., ICRA '11]



Visual Navigation for Flying Robots

74

Dr. Jürgen Sturm, Computer Vision Group, TUM

Real-Time Motion Planning

- What is the maximum time needed to re-plan in case of an obstacle detection?
- What if the robot has to react quickly to unforeseen, fast moving objects?
- Do we really need to re-plan for every obstacle on the way?

Visual Navigation for Flying Robots

75

Dr. Jürgen Sturm, Computer Vision Group, TUM

Real-Time Motion Planning

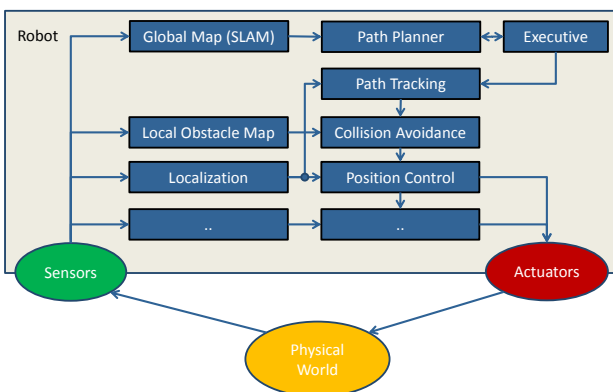
- What is the maximum time needed to re-plan in case of an obstacle detection?
In principle, re-planning with D can take arbitrarily long*
- What if the robot has to react quickly to unforeseen, fast moving objects?
Need a collision avoidance algorithm that runs in constant time!
- Do we really need to re-plan for every obstacle on the way?
Could trigger re-planning only if path gets obstructed (or robot predicts that re-planning reduces path length by p%)

Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robot Architecture



Visual Navigation for Flying Robots

77

Dr. Jürgen Sturm, Computer Vision Group, TUM

Layered Motion Planning

- An approximate **global planner** computes paths ignoring the kinematic and dynamic vehicle constraints (not real-time)
- An accurate **local planner** accounts for the constraints and generates feasible local trajectories in real-time (collision avoidance)

Visual Navigation for Flying Robots

78


Dr. Jürgen Sturm, Computer Vision Group, TUM

Local Planner

- **Given:** Path to goal (sequence of via points), range scan of the local vicinity, dynamic constraints
- **Wanted:** Collision-free, safe, and fast motion towards the goal (or next via point)
- Typical approaches:
 - Potential fields
 - Dynamic window approach

Navigation with Potential Fields


- Treat robot as a particle under the influence of a potential field
- **Pro:**
 - easy to implement
- **Con:**
 - suffers from local minima
 - no consideration of dynamic constraints



Dynamic Window Approach

[Simmons, 96], [Fox et al., 97], [Brock & Khatib, 99]


- Consider a 2D planar robot



Dynamic Window Approach

[Simmons, 96], [Fox et al., 97], [Brock & Khatib, 99]


- Consider a 2D planar robot + 2D environment



Dynamic Window Approach

[Simmons, 96], [Fox et al., 97], [Brock & Khatib, 99]


- Consider additionally dynamic constraints



Dynamic Window Approach

[Simmons, 96], [Fox et al., 97], [Brock & Khatib, 99]


- Navigation function (potential field)



Dynamic Window Approach

[Simmons, 96], [Fox et al., 97], [Brock & Khatib, 99]

- Navigation function (potential field)



Visual Navigation for Flying Robots

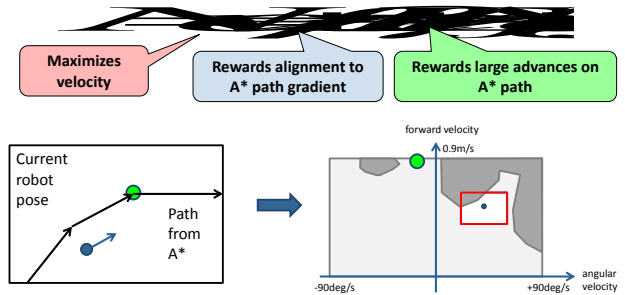
85

Dr. Jürgen Sturm, Computer Vision Group, TUM

Dynamic Window Approach

[Simmons, 96], [Fox et al., 97], [Brock & Khatib, 99]

- Navigation function (potential field)



Visual Navigation for Flying Robots


86

Dr. Jürgen Sturm, Computer Vision Group, TUM

Dynamic Window Approach

[Simmons, 96], [Fox et al., 97], [Brock & Khatib, 99]

- Discretize dynamic window and evaluate navigation function (note: window has fixed size = real-time!)
- Find the maximum and execute motion



Visual Navigation for Flying Robots

87

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Dynamic Window Approach

[Brock and Khatib, ICRA '99]



Visual Navigation for Flying Robots

88

Dr. Jürgen Sturm, Computer Vision Group, TUM

Problems of DWAs

- DWAs suffer from local minima (need tuning), e.g., robot does not slow down early enough to enter doorway:



- Can you think of a solution?
- Note:** General case requires global planning

Visual Navigation for Flying Robots

89

Dr. Jürgen Sturm, Computer Vision Group, TUM

Lessons Learned Today

- Motion planning problem and configuration spaces
- Roadmap construction
- Search algorithms and path optimization
- Local planning for path execution

Visual Navigation for Flying Robots

90

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

Planning under Uncertainty, Exploration and Coordination

Dr. Jürgen Sturm

Agenda for Today

- Planning under Uncertainty
- Exploration with a single robot
- Coordinated exploration with a team of robots
- Coverage

Visual Navigation for Flying Robots

2

Dr. Jürgen Sturm, Computer Vision Group, TUM

Agenda For Next Week

- **First half:** Good practices for experimentation, evaluation and benchmarking
- **Second half:** Time for your questions on course material

→ Prepare your questions (if you have)


Visual Navigation for Flying Robots

3

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motivation: Planning under Uncertainty

- Consider a robot with range-limited sensors and a feature-poor environment
- Which route should the robot take?



Visual Navigation for Flying Robots

4

Dr. Jürgen Sturm, Computer Vision Group, TUM

Reminder: Performance Metrics

- Execution speed / path length
- Energy consumption
- Planning speed
- Safety (minimum distance to obstacles)
- **Robustness against disturbances**
- **Probability of success**
- ...

Visual Navigation for Flying Robots

5

Dr. Jürgen Sturm, Computer Vision Group, TUM

Reminder: Belief Distributions

- In general, actions of the robot are not carried out perfectly
- Position estimation ability depends on map
- Let's look at the belief distributions...




Visual Navigation for Flying Robots

6

Dr. Jürgen Sturm, Computer Vision Group, TUM


Reminder: Belief Distributions

- Actions increase the uncertainty (in general)
- Observations decrease the uncertainty (always)
- Observations are not always available




Solution 1: Shape The Environment To Decrease Uncertainty

- Assume a robot without sensors
- What is a good navigation plan?




Solution 1: Shape The Environment To Decrease Uncertainty

- Plan 1: Take the shortest path
- What is the probability of success of plan 1?



Solution 1: Shape The Environment To Decrease Uncertainty

- What is the probability of success of plan 2?



Solution 1: Shape The Environment To Decrease Uncertainty

- **Pro:** Simple solution, need fewer/no sensors
- **Con:** Requires task specific design/engineering of both the robot and the environment
- Applications:
 - Docking station
 - Perception-less manipulation (on conveyer belts)
 - ...

Solution 2: Add (More/Better) Sensors



Solution 3: POMDPs

- Partially observable Markov decision process (POMDP)
- Considers uncertainty of the motion model and sensor model
- Finite/infinite time horizon
- Resulting policy is optimal
- One solution technique: Value iteration
- **Problem:** In general (and in practice) computationally intractable (PSPACE-hard)

Visual Navigation for Flying Robots

13

Dr. Jürgen Sturm, Computer Vision Group, TUM

Continuum of Possible Approaches to Motion Planning

Conventional path planner

POMDP

tractable
not robust

intractable
robust


maybe we can find something in between...

Visual Navigation for Flying Robots

14

Dr. Jürgen Sturm, Computer Vision Group, TUM

Remember: Motion Planning




Slides adopted from Nick Roy
Visual Navigation for Flying Robots

15

Goal: shortest path, subject to kinematic and environmental constraints

Remember: Motion Planning in High-Dimensional Configuration Spaces



Remember: Probabilistic Roadmaps

1. Add vertices (sampled in free space)
2. Add edges between neighboring vertices (when line of sight is not obstructed)
3. Find shortest path (Dijkstra, ...)

Slides adopted from Nick Roy
Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Remember: Motion Planning in High-Dimensional Configuration Spaces

- **Problem:** The roadmap does not consider the sensor capabilities of the robot
- Can the robot actually keep position at each vertex?
 - Can it localize at the vertex?
 - Given localization abilities, what is the probability of hitting into an obstacle?
- Can the robot robustly navigate between two vertices?
 - Line of sight is not enough
 - Robot might get lost or hit into an obstacle


Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Planning in Information Space

[Roy et al.]




2. Add edges between points where $p(x \in C_{\text{obst}}) < \epsilon$ along path
3. Perform graph search

Visual Navigation for Flying Robots

Motion Planning in Information Space

- **Problem:** Posterior distribution depends also on the path taken to the vertex
- Example




Visual Navigation for Flying Robots

19

Dr. Jürgen Sturm, Computer Vision Group, TUM

Belief Roadmap

[He et al., 2008]



1. Sample vertices from C_{free} , build graph and estimate belief dist. transfer functions
2. Propagate covariances by performing graph search

Slides adopted from Nick Roy
Visual Navigation for Flying Robots

Planning in Information Spaces

[He et al., 2008]

- **Given:** Roadmap
- **Goal:** Find path from start to goal nodes that results in minimum uncertainty at goal
- **Problem:** How can we estimate the belief distribution at the goal (efficiently)?

Visual Navigation for Flying Robots

21


Dr. Jürgen Sturm, Computer Vision Group, TUM

Planning in Information Spaces

[He et al., 2008]

How can we propagate the belief distribution along an edge?

1. Sample waypoints, use forward simulation to compute full posterior
2. Linearize model and use Kalman filter




Visual Navigation for Flying Robots

22

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Belief Roadmap

[He et al., 2008]



Visual Navigation for Flying Robots


23

Dr. Jürgen Sturm, Computer Vision Group, TUM

Belief Propagation

[He et al., 2008]

- The posterior distribution depends on the prior distribution



Visual Navigation for Flying Robots

24

Dr. Jürgen Sturm, Computer Vision Group, TUM

Planning in Information Spaces

[He et al., 2008]

- The posterior distribution at a vertex depends on the prior distribution (and thus on path to the vertex)
- Need to perform forward simulation (and belief prediction) along each edge for every start state
- Computing minimum cost path of 30 edges: ≈ 100 seconds

Summary: Planning Under Uncertainty

- Actions and observations are inherently noisy
- Planners neglecting this are not robust
- Consider the uncertainty during planning to increase robustness



Mission Planning

- **Goal:** Generate and execute a plan to accomplish a certain (navigation) task
- Example tasks
 - Exploration
 - Coverage
 - Surveillance
 - Tracking
 - ...

Task Planning


- **Goal:** Generate and execute a high level plan to accomplish a certain task
- Often symbolic reasoning (or hard-coded)
 - Propositional or first-order logic
 - Automated reasoning systems
 - Common programming languages: Prolog, LISP
- Multi-agent systems, communication
- Artificial Intelligence

Exploration and SLAM

- SLAM is typically passive, because it consumes incoming sensor data
- Exploration actively guides the robot to cover the environment with its sensors
- Exploration in combination with SLAM: Acting under pose and map uncertainty
- Uncertainty should/needs to be taken into account when selecting an action

Exploration

- By reasoning about control, the mapping process can be made much more effective
- Question: **Where to move next?**



- This is also called the **next-best-view problem**

Visual Navigation for Flying Robots

31

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exploration

- Choose the action that maximizes utility

$$a^* = \arg \max_{a \in A} U(m, a)$$

- Question: How can we define utility?

Visual Navigation for Flying Robots

32

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example

- Where should the robot go next?



Visual Navigation for Flying Robots


33

Dr. Jürgen Sturm, Computer Vision Group, TUM

Maximizing the Information Gain

- Pick the action a that maximizes the **information gain** given a map m

$$a^* = \arg \max_{a \in A} IG(m, a)$$



Visual Navigation for Flying Robots

34

Dr. Jürgen Sturm, Computer Vision Group, TUM

Information Theory

- Entropy** is a general measure for the uncertainty of a probability distribution
- Entropy = Expected amount of information needed to encode an outcome $X = x$

$$\begin{aligned} H(X) &= E(I(X)) \\ &= E(-\log p(X)) \\ &= -\sum_{i=1}^n p(x_i) \log p(x_i) \end{aligned}$$

Visual Navigation for Flying Robots

35

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Binary Random Variable

- Binary random variable $X \in \{0, 1\}$
- Probability distribution $P(X = 1) = p$
- How many bits do we need to transmit one sample of $p(X)$?
 - For $p=0$?
 - For $p=0.5$?
 - For $p=1$?


Visual Navigation for Flying Robots

36

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Binary Random Variable

- Binary random variable $X \in \{0, 1\}$
- Probability distribution $P(X = 1) = p$
- How many bits do we need to transmit one sample of $p(X)$?
- Answer:

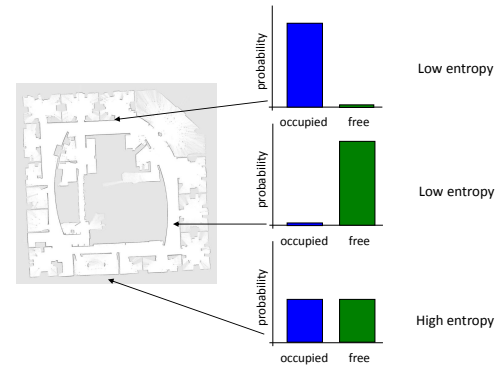


Visual Navigation for Flying Robots

37

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Map Entropy



The overall entropy is the sum of the individual entropy values

Visual Navigation for Flying Robots

38

Dr. Jürgen Sturm, Computer Vision Group, TUM

Information Theory

- Information gain** = Uncertainty reduction

$$IG(X, Y) = H(X) - H(X | Y)$$

- Conditional entropy**

$$H(X | Y) = \sum_{i,j} p(x_i, y_j) \log \frac{p(y_j)}{p(x_i, y_j)}$$

Visual Navigation for Flying Robots

39

Dr. Jürgen Sturm, Computer Vision Group, TUM

Maximizing the Information Gain

- To compute the information gain one needs to know the observations obtained when carrying out an action

$$a^* = \arg \max_{a \in A} IG(m, a)$$

- This quantity is not known! Reason about potential measurements


$$a^* = \arg \max_{a \in A} \int IG(m, z) p(z | a) dz$$

Visual Navigation for Flying Robots

40

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example



Visual Navigation for Flying Robots

41

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exploration Costs

- So far, we did not consider the cost of executing an action (e.g., time, energy, ...)

- Utility = uncertainty reduction – cost**

- Select the action with the highest expected utility

$$a^* = \arg \max_{a \in A} IG(m, a) - \alpha \cdot E(cost(m, a))$$


Visual Navigation for Flying Robots

42

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exploration

- For each location $\langle x, y \rangle$
 - Estimate the number of cells robot can sense (e.g., simulate laser beams using current map)
 - Estimate the cost of getting there




Visual Navigation for Flying Robots

43

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exploration

- Greedy strategy:** Select the candidate location with the highest utility, then repeat...



Visual Navigation for Flying Robots

44

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exploration Actions

- So far, we only considered reduction in map uncertainty
- In general, there are many sources of uncertainty that can be reduced by exploration
 - Map uncertainty (visit unexplored areas)
 - Trajectory uncertainty (loop closing)
 - Localization uncertainty (active re-localization by re-visiting known locations)

Visual Navigation for Flying Robots


45

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Active Loop Closing

[Stachniss et al., 2005]

- Reduce map uncertainty



- Reduce map + path uncertainty




Visual Navigation for Flying Robots

46

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Active Loop Closing

[Stachniss et al., 2005]



Visual Navigation for Flying Robots


47

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Active Loop Closing

[Stachniss et al., 2005]

- Entropy evolution




Visual Navigation for Flying Robots

48

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Reduce uncertainty in map, path, and pose [Stachniss et al., 2005]



Selected target location




Visual Navigation for Flying Robots

49

Dr. Jürgen Sturm, Computer Vision Group, TUM

Corridor Exploration [Stachniss et al., 2005]



- The decision-theoretic approach leads to **intuitive behaviors**: “re-localize before getting lost”
- Some animals show a similar behavior (dogs marooned in the tundra of north Russia)

Visual Navigation for Flying Robots

50

Dr. Jürgen Sturm, Computer Vision Group, TUM

Multi-Robot Exploration

Given: Team of robots with communication

Goal: Explore the environment as fast as possible



[Wurm et al., IROS 2011]

Visual Navigation for Flying Robots

51

Dr. Jürgen Sturm, Computer Vision Group, TUM

Complexity

- Single-robot exploration in known, graph-like environments is in general **NP-hard**
- Proof: Reduce traveling salesman problem to exploration
- Complexity of multi-robot exploration is **exponential** in the number of robots


Visual Navigation for Flying Robots

52

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motivation: Why Coordinate?

Robot 1



Robot 2



- Without coordination, two robots might choose the same exploration frontier

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Levels of Coordination

1. **No exchange of information**
2. **Implicit coordination:** Sharing a joint map
 - Communication of the individual maps and poses
 - Central mapping system
3. **Explicit coordination:** Determine better target locations to distribute the robots
 - Central planner for target point assignment
 - Minimize expected path cost / information gain / ...


Visual Navigation for Flying Robots

54


Dr. Jürgen Sturm, Computer Vision Group, TUM

Typical Trajectories

Implicit coordination:



Explicit coordination:




Visual Navigation for Flying Robots

55

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exploration Time

[Stachniss et al., 2006]



Visual Navigation for Flying Robots

56

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coordination Algorithm

In each time step:

- Determine set of exploration targets
 $S = \{s_1, \dots, s_n\}$
- Compute for each robot i and each target j the expected cost/utility C_{ij}
- Assign robots to targets using the **Hungarian algorithm**

Visual Navigation for Flying Robots

57

Dr. Jürgen Sturm, Computer Vision Group, TUM

Hungarian Algorithm

[Kuhn, 1955]

- Combinatorial optimization algorithm
- Solves the assignment problem in polynomial time $O(n^3)$
- General idea: Algorithm modifies the cost matrix until there is zero cost assignment


Visual Navigation for Flying Robots

58

Dr. Jürgen Sturm, Computer Vision Group, TUM

Hungarian Algorithm: Example

		targets			
		W	X	Y	Z
robots	a	3	2	3	2
	b	2	5	6	3
	c	7	1	3	5
	d	6	2	3	5




1. Compute the cost matrix (non-negative)

59 / 16

Hungarian Algorithm: Example

		targets			
		W	X	Y	Z
robots	a	3	2	3	2
	b	2	5	6	3
	c	7	1	3	5
	d	6	2	3	5



2. Find minimum element in each row

60 / 16

Hungarian Algorithm: Example

		targets				
		W	X	Y	Z	
robots	a	3	2	3	2	2
	b	2	5	6	3	2
	c	7	1	3	5	1
	d	6	2	3	5	2

		W	X
			(c)
(b)	(a)		(d)
Z	Y		

3. Subtract minimum from each row element

61 / 16

Hungarian Algorithm: Example

		targets			
		W	X	Y	Z
robots	a	1	0	1	0
	b	0	3	4	1
	c	6	0	2	4
	d	4	0	1	3

		W	X
			(c)
(b)	(a)		(d)
Z	Y		

0 0 1 0

4. Find minimum element in each column

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

62

Hungarian Algorithm: Example

		targets			
		W	X	Y	Z
robots	a	1	0	0	0
	b	0	3	3	1
	c	6	0	1	4
	d	4	0	0	3

		W	X
			(c)
(b)	(a)		(d)
Z	Y		

0 0 1 0

5. Subtract minimum from each column element

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

63

Hungarian Algorithm: Example

		W	X	Y	Z
robots	a	1	0	0	0
	b	0	3	3	1
	c	6	0	1	4
	d	4	0	0	3

		W	X
			(c)
(b)	(a)		(d)
Z	Y		

6a. Assign (if possible)

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

64

Hungarian Algorithm: Example

		W	X	Y	Z
robots	a	1	0	0	0
	b	0	3	3	1
	c	6	0	1	4
	d	4	0	0	3

6b. If no assignment is possible:

- Connect all 0's by lines
- Find the minimum in all remaining elements and subtract
- Repeat step 2 – 6

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

65

Hungarian Algorithm: Example

		targets			
		X	Y	X'	Y'
robots	a	2	3	2	3
	b	5	6	5	6
	c	1	3	1	3
	d	2	3	2	3

If there are not enough targets:
Copy targets to allow multiple assignments

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

66

Example: Segmentation-based Exploration [Wurm et al., IROS 2008]

- Two-layer hierarchical role assignments using Hungarian algorithm (1: rooms, 2: targets in room)
- Reduces exploration time and risk of interferences



Visual Navigation for Flying Robots

67

Dr. Jürgen Sturm, Computer Vision Group, TUM

Summary: Exploration

- Exploration aims at generating robot motions so that an **optimal map** is obtained
- Coordination** reduces exploration time
- Hungarian algorithm** efficiently solves the assignment problem (centralized, 1-step lookahead)
- Challenges (active research):
 - Limited bandwidth and **unreliable communication**
 - Decentralized planning** and task assignment

Visual Navigation for Flying Robots

68

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage Path Planning

- Given:** Known environment with obstacles
- Wanted:** The shortest trajectory that ensures complete (sensor) coverage



Visual Navigation for Flying Robots

69




[images from Xu et al., ICRA 2011]

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage Path Planning



Visual Navigation for Flying Robots




70

Dr. Jürgen Sturm, Computer Vision Group, TUM



70



Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage Path Planning: Applications

- For flying robots
 - Search and rescue
 - Area surveillance
 - Environmental inspection
 - Inspection of buildings (bridges)
- For service robots
 - Lawn mowing
 - Vacuum cleaning
- For manipulation robots
 - Painting
 - Automated farming

Visual Navigation for Flying Robots

71

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage Path Planning

- What is a good coverage strategy?
- What would be a good cost function?

Visual Navigation for Flying Robots

72

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage Path Planning

- What is a good coverage strategy?
- What would be a good cost function?
 - Amount of redundant traversals
 - Number of stops and rotations
 - Execution time
 - Energy consumption
 - Robustness
 - Probability of success
 - ...

Visual Navigation for Flying Robots

73

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage Path Planning

- Related to the traveling salesman problem (TSP):
“Given a weighted graph, compute a path that visits every vertex once”
- In general **NP-complete**
- Many approximations exist
- Many approximate (and exact) solvers exist


Visual Navigation for Flying Robots

74

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage of Simple Shapes

- Approximately optimal solution often easy to compute for simple shapes (e.g., trapezoids)




Visual Navigation for Flying Robots

75

Dr. Jürgen Sturm, Computer Vision Group, TUM

Idea

[Mannadiar and Rekleitis, ICRA 2011]




Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM

Idea

[Mannadiar and Rekleitis, ICRA 2011]




Visual Navigation for Flying Robots

77

Dr. Jürgen Sturm, Computer Vision Group, TUM

Idea

[Mannadiar and Rekleitis, ICRA 2011]



Visual Navigation for Flying Robots

78

Dr. Jürgen Sturm, Computer Vision Group, TUM

Coverage Based On Cell Decomposition [Mannadiar and Rekleitis, ICRA 2011]

Approach:

1. Decompose map into “simple” cells
2. Compute connectivity between cells and build graph
3. Solve coverage problem on reduced graph


Visual Navigation for Flying Robots

79

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 1: Boustrophedon Cellular Decomposition [Mannadiar and Rekleitis, ICRA 2011]

- Similar to trapezoidal decomposition
- Can be computed efficiently




Visual Navigation for Flying Robots

80

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 2: Build Reeb Graph [Mannadiar and Rekleitis, ICRA 2011]

- Vertices = Critical points (that triggered the split)
- Edges = Connectivity between critical points




Visual Navigation for Flying Robots

81

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 3: Compute Euler Tour [Mannadiar and Rekleitis, ICRA 2011]

- Extend graph so that vertices have even order
- Compute Euler tour (linear time)




Visual Navigation for Flying Robots

82

Dr. Jürgen Sturm, Computer Vision Group, TUM

Resulting Coverage Plan [Mannadiar and Rekleitis, ICRA 2011]

- Follow the Euler tour
- Use simple coverage strategy for cells
- Note: Cells are visited once or twice



Visual Navigation for Flying Robots

83

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robotic Cleaning of 3D Surfaces [Hess et al., IROS 2012]

- **Goal:** Cover entire surface while minimizing trajectory length in configuration space



- **Approach:**
 - Discretize 3D environment into patches
 - Build a neighborhood graph
 - Formulate the problem as generalized TSP (GTSP)


Visual Navigation for Flying Robots

84

Dr. Jürgen Sturm, Computer Vision Group, TUM

Robotic Cleaning of 3D Surfaces

[Hess et al., IROS 2012]



Visual Navigation for Flying Robots

85

Dr. Jürgen Sturm, Computer Vision Group, TUM

Lessons Learned Today

- How to generate plans that are robust to uncertainty in sensing and locomotion
- How to explore an unknown environment
 - With a single robot
 - With a team of robots
- How to generate plans that fully cover known environments

Visual Navigation for Flying Robots

86

Dr. Jürgen Sturm, Computer Vision Group, TUM

Video: SFLY Final Project Demo (2012)



sFly

Swarm of Micro Flying Robots

<http://www.sfly.org/>



Visual Navigation for Flying Robots

87

Dr. Jürgen Sturm, Computer Vision Group, TUM

Visual Navigation for Flying Robots

Experimentation, Evaluation and Benchmarking

Dr. Jürgen Sturm

Agenda for Today

- Course Evaluation
- Scientific research: The big picture
- Best practices in experimentation
- Datasets, evaluation criteria and benchmarks

- Time for questions

Course Evaluation

- Much positive feedback – thank you!!!
- We are also very happy with you as a group. Everybody seemed to be highly motivated!
- Suggestions for improvements (from course evaluation forms)
 - Workload was considered a bit too high
→ ECTS have been adjusted to 6 credits
 - ROS introduction lab course would be helpful
→ Will do this next time
- Any further suggestions/comments?

Scientific Research – General Idea

1. Observe phenomena
2. Formulate explanations and theories
3. Test them

Scientific Research – Methodology

1. Generate an idea
2. Develop an approach that solves the problem
3. Demonstrate the validity of your solution
4. Disseminate your results
5. At all stages: iteratively refine

Scientific Research in Student Projects

- How can you get involved in scientific research during your study?

Scientific Research in Student Projects

- How can you get involved in scientific research during your study?
 - Bachelor lab course (10 ECTS)
 - Bachelor thesis (15 ECTS)
 - Graduate lab course (10 ECTS)
 - Interdisciplinary project (16 ECTS)
 - Master thesis (30 ECTS)
 - Student research assistant (10 EUR/hour, typically 10 hours/week)

Visual Navigation for Flying Robots

7

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 1: Generate the Idea

- Be creative
- Follow your interests / preferences
- Examples:
 - Research question
 - Challenging problem
 - Relevant application
 - Promising method (e.g., try to transfer method from another field)

Visual Navigation for Flying Robots

8

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 1b: Find related work

- There is **always** related work
- Find related research papers
 - Use Google scholar, paper repositories, ...
 - Navigate the citation network
 - Read survey articles
- Browse through (recent) text books
- Ask your professor, colleagues, ...
- It's very unlikely that somebody else has already perfectly solved exactly your problem, so don't worry! Technology evolves very fast...

Visual Navigation for Flying Robots

9

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 2: Develop a Solution

- Practitioner
 - Start programming
 - Realize that it is not going to work, start over, ...
 - When it works, formalize it (try to find out why it works and what was missing before)
 - Empirically verify that it works
- Theorist
 - Formalize the problem
 - Find suitable method
 - (Theoretically) prove that it is right
 - (If needed) implement a proof-of-concept

Visual Navigation for Flying Robots

10

Dr. Jürgen Sturm, Computer Vision Group, TUM

Step 3: Validation

- What are your claims?
- How can you prove them?
 - Theoretical proof (mathematical problem)
 - Experimental validation
 - Qualitative (e.g., video)
 - Quantitative (e.g., many trials, statistical significance)
- Compare and discuss your results with respect to previous work/approaches

Visual Navigation for Flying Robots

11

Dr. Jürgen Sturm, Computer Vision Group, TUM

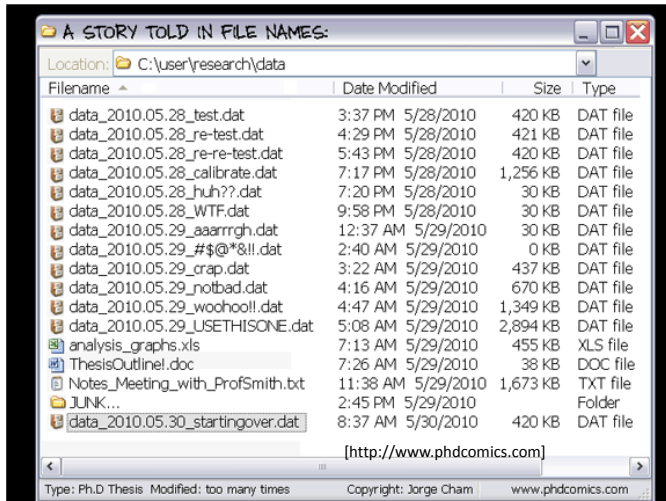
Step 4: Dissemination

- Good solution/expertise alone is not enough
- You need to convince other people in the field
- Usual procedure:
 1. Write research paper (usually 6-8 pages) **3-6 month**
 2. Submit PDF to an international conference or journal
 3. Paper will be peer-reviewed **3-6 month**
 4. Improve paper (if necessary)
 5. Give talk or poster presentation at conference **15 min.**
 6. Optionally: Repeat step 1-5 until PhD 😊 **3-5 years**

Visual Navigation for Flying Robots

12

Dr. Jürgen Sturm, Computer Vision Group, TUM



Step 5: Refinement

- Discuss your work with
 - Your colleagues
 - Your professor
 - Other colleagues at conferences
- Improve your approach and evaluation
 - Adopt notation to the standard
 - Get additional references/insights
 - Conduct more/additional experiments
- Simplify and generalize your approach
- Collaborate with other people (in other fields)

Visual Navigation for Flying Robots

14

Dr. Jürgen Sturm, Computer Vision Group, TUM

Scientific Research

- This was the big picture
- Today's focus is on best practices in experimentation
- **What do you think are the (desired) properties of a good scientific experiment?**

Visual Navigation for Flying Robots

15

Dr. Jürgen Sturm, Computer Vision Group, TUM

What are the desired properties of a good scientific experiment?

- Reproducibility / repeatability
 - Document the experimental setup
 - Choose (and motivate) an your evaluation criterion
- Experiments should allow you to validate/falsify competing hypotheses

Current trends:

- Make data available for review and criticism
- Same for software (open source)

Visual Navigation for Flying Robots

16

Dr. Jürgen Sturm, Computer Vision Group, TUM

Challenges

- Reproducibility is sometimes not easy to guarantee
- Any ideas why?

Visual Navigation for Flying Robots

17

Dr. Jürgen Sturm, Computer Vision Group, TUM

Challenges

- Randomized components/noise (beat with the law of large numbers/statistical tests)
- Experiment requires special hardware
 - Self-built, unique robot
 - Expensive lab equipment
 - ...
- Experiments cost time
- "(Video) Demonstrations will suffice"
- Technology changes fast

Visual Navigation for Flying Robots

18

Dr. Jürgen Sturm, Computer Vision Group, TUM

Benchmarks

- Effective and affordable way of conducting experiments
- Sample of a task domain
- Well-defined performance measurements
- Widely used in computer vision and robotics
- **Which benchmark problems do you know?**

Visual Navigation for Flying Robots

19

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Benchmark Problems

Computer Vision

- Middlebury datasets (optical flow, stereo, ...)
- Caltech-101, PASCAL (object recognition)
- Stanford bunny (3d reconstruction)

Robotics

- RoboCup competitions (robotic soccer)
- DARPA challenges (autonomous car)
- SLAM datasets

Visual Navigation for Flying Robots

20

Dr. Jürgen Sturm, Computer Vision Group, TUM

Image Denoising: Lenna Image

- 512x512 pixel standard image for image compression and denoising
- Lena Söderberg, Playboy magazine Nov. 1972
- Scanned by Alex Sawchuck at USC in a hurry for a conference paper



<http://www.cs.cmu.edu/~chuck/lennapp/>


Visual Navigation for Flying Robots

21

Dr. Jürgen Sturm, Computer Vision Group, TUM

Object Recognition: Caltech-101

- Pictures of objects belonging to 101 categories
- About 40-800 images per category
- Recognition, classification, categorization



Visual Navigation for Flying Robots

22

Dr. Jürgen Sturm, Computer Vision Group, TUM

RoboCup Initiative

- Evaluation of full system performance
- Includes perception, planning, control, ...
- Easy to understand, high publicity
- “By mid-21st century, a team of fully autonomous humanoid robot soccer players shall win the soccer game, complying with the official rule of the FIFA, against the winner of the most recent World Cup.”

Visual Navigation for Flying Robots

23

Dr. Jürgen Sturm, Computer Vision Group, TUM

RoboCup Initiative



Visual Navigation for Flying Robots

24

Dr. Jürgen Sturm, Computer Vision Group, TUM

SLAM Evaluation

- Intel dataset: laser + odometry [Haehnel, 2004]
- New College dataset: stereo + omni-directional vision + laser + IMU [Smith et al., 2009]
- TUM RGB-D dataset [Sturm et al., 2011/12]
- ...



Visual Navigation for Flying Robots

25

Dr. Jürgen Sturm, Computer Vision Group, TUM

TUM RGB-D Dataset

[Sturm et al., RSS RGB-D 2011; Sturm et al., IROS 2012]

- RGB-D dataset with ground truth for SLAM evaluation
- Two error metrics proposed (relative and absolute error)
- Online + offline evaluation tools
- Training datasets (fully available)
- Validation datasets (ground truth not publicly available to avoid overfitting)


Visual Navigation for Flying Robots

26

Dr. Jürgen Sturm, Computer Vision Group, TUM

Recorded Scenes

- Various scenes (handheld/robot-mounted, office, industrial hall, dynamic objects, ...)
- Large variations in camera speed, camera motion, illumination, environment size, ...



Visual Navigation for Flying Robots

27

Dr. Jürgen Sturm, Computer Vision Group, TUM

Dataset Acquisition

- Motion capture system
 - Camera pose (100 Hz)
- Microsoft Kinect
 - Color images (30 Hz)
 - Depth maps (30 Hz)
 - IMU (500 Hz)
- External video camera (for documentation)

Visual Navigation for Flying Robots

28

Dr. Jürgen Sturm, Computer Vision Group, TUM

Motion Capture System

- 9 high-speed cameras mounted in room
- Cameras have active illumination and pre-process image (thresholding)
- Cameras track positions of retro-reflective markers



Visual Navigation for Flying Robots

29

Dr. Jürgen Sturm, Computer Vision Group, TUM

Calibration

Calibration of the overall system is not trivial:

1. Mocap calibration
2. Kinect-mocap calibration
3. Time synchronization

Visual Navigation for Flying Robots

30

Dr. Jürgen Sturm, Computer Vision Group, TUM

Calibration Step 1: Mocap

- Need at least 2 cameras for position fix
- Need at least 3 markers on object for full pose
- Calibration stick for extrinsic calibration



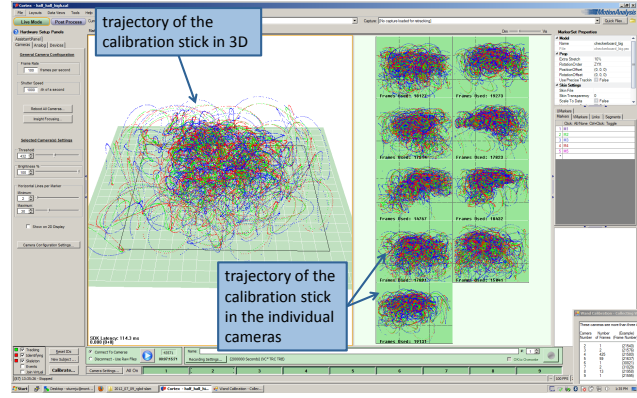
Visual Navigation for Flying Robots




31

Dr. Jürgen Sturm, Computer Vision Group, TUM

Calibration Step 1: Mocap



Example: Raw Image from Mocap




Visual Navigation for Flying Robots

33

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Position Triangulation of a Single Marker




Visual Navigation for Flying Robots

34

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Tracked Object (4 Markers)




Visual Navigation for Flying Robots

35

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example: Recorded Trajectory




Visual Navigation for Flying Robots

36

Dr. Jürgen Sturm, Computer Vision Group, TUM

Calibration Step 2: Mocap-Kinect

- Need to find transformation between the markers on the Kinect and the optical center
- Special calibration board visible both by Kinect and mocap system (manually gauged)



Visual Navigation for Flying Robots



37


Dr. Jürgen Sturm, Computer Vision Group, TUM

Calibration Step 3: Time Synchronization

- Assume a constant time delay between mocap and Kinect messages
- Choose time delay that minimizes reprojection error during checkerboard calibration



Visual Navigation for Flying Robots



38

Dr. Jürgen Sturm, Computer Vision Group, TUM

Calibration - Validation


- Intrinsic calibration
- Extrinsic calibration color + depth
- Time synchronization color + depth
- Mocap system slowly drifts (need re-calibration every hour)
- Validation experiments to check the quality of calibration
 - 2mm length error on 2m rod across mocap volume
 - 4mm RMSE on checkerboard sequence

Visual Navigation for Flying Robots

39

Dr. Jürgen Sturm, Computer Vision Group, TUM

Example Sequence: Freiburg1/XYZ



Sequence description (on the website):

“For this sequence, the Kinect was pointed at a typical desk in an office environment. This sequence contains only translatory motions along the principal axes of the Kinect, while the orientation was kept (mostly) fixed. This sequence is well suited for debugging purposes, as it is very simple.”

Visual Navigation for Flying Robots

40

Dr. Jürgen Sturm, Computer Vision Group, TUM

Computer Vision Group
Technische Universität München

Home • Datasets and Software • Datasets • RGB-D SLAM Dataset and Benchmark

RGB-D SLAM Dataset and Benchmark
Contact: Jürgen Sturm

We provide a large dataset containing RGB-D data and ground-truth data with the goal to establish a novel benchmark for the evaluation of visual odometry and visual SLAM systems. Our dataset contains the color and depth images of a Microsoft Kinect sensor along the ground-truth trajectory of the sensor. The data was recorded at full frame rate (30 Hz) and sensor resolution (640x480). The ground-truth trajectory was obtained from a high-accuracy motion-capture system with eight high-speed tracking cameras (100 Hz). Further, we provide the accelerometer data from the Kinect. Finally, we propose an evaluation criterion for measuring the quality of the estimated camera trajectory of visual SLAM systems.

How can I use the RGB-D Benchmark to evaluate my SLAM system?

1. Download one or more of the RGB-D benchmark sequences (file formats, useful tools)
2. Run your favorite visual odometry/visual SLAM algorithm (for example, [RGB-D SLAM](#))
3. Save the estimated camera trajectory to a file (file formats: [example trajectory](#))
4. Evaluate your algorithm by comparing the estimated trajectory with the ground truth trajectory. We provide an automated evaluation tool to help you with the evaluation. There is also an online version of the tool.

Computer Vision Group
Technische Universität München

Home • Datasets and Software • Datasets • RGB-D SLAM Dataset and Benchmark • download

Dataset Download

We recommend that you use the 'xyz' series for your first experiments. The motion is relatively small, and only a small volume on an office desk is covered. Once this works, you might want to try the 'desk' dataset, which covers four tables and contains several loop closures.

We are happy to share our data with other researchers. Please refer to the respective publication when using this data.

Remarks:

- The file formats are described here.
- The intrinsic camera parameters are here.
- We provide a set of useful tools for working with the dataset.
- The `_validation` sequences do not contain ground truth. They can only be evaluated using the online tool.

Sequence name	Duration	Length	Download
Category: Testing and Debugging			
freiburg1_xyz	30.09s	7.112m	tgz (0.47GB) more info
freiburg1_rny	27.67s	1.664m	tgz (0.42GB) more info
freiburg2_xyz	122.74s	7.029m	tgz (2.39GB) more info



Dataset Website

- In total: 39 sequences (19 with ground truth)
- One ZIP archive per sequence, containing
 - Color and depth images (PNG)
 - Accelerometer data (timestamp ax ay az)
 - Trajectory file (timestamp tx ty qx qy qz qw)
- Sequences also available as ROS bag and MRPT rawlog

<http://vision.in.tum.de/data/datasets/rgbd-dataset>


Visual Navigation for Flying Robots

44

Dr. Jürgen Sturm, Computer Vision Group, TUM

What Is a Good Evaluation Metric?

- Compare camera trajectories
 - Ground truth trajectory $Q_1, \dots, Q_n \in SE(3)$
 - Estimate camera trajectory $P_1, \dots, P_n \in SE(3)$
- Two common evaluation metrics
 - Relative pose error (drift per second)
 - Absolute trajectory error (global consistency)



Visual Navigation for Flying Robots

45


Dr. Jürgen Sturm, Computer Vision Group, TUM

Relative Pose Error (RPE)

- Measures the (relative) **drift**
- Recommended for the evaluation of visual odometry approaches

$$E_i := \left(Q_i^{-1} Q_{i+\Delta} \right)^{-1} \left(P_i^{-1} P_{i+\Delta} \right)$$

Relative error
True motion
Estimated motion



Visual Navigation for Flying Robots


46

Dr. Jürgen Sturm, Computer Vision Group, TUM

Absolute Trajectory Error (ATE)

- Measures the **global error**
- Requires pre-aligned trajectories
- Recommended for SLAM evaluation

$$E_i := Q_i^{-1} S P_i$$



Visual Navigation for Flying Robots

47

Dr. Jürgen Sturm, Computer Vision Group, TUM

Evaluation metrics

- Average over all time steps


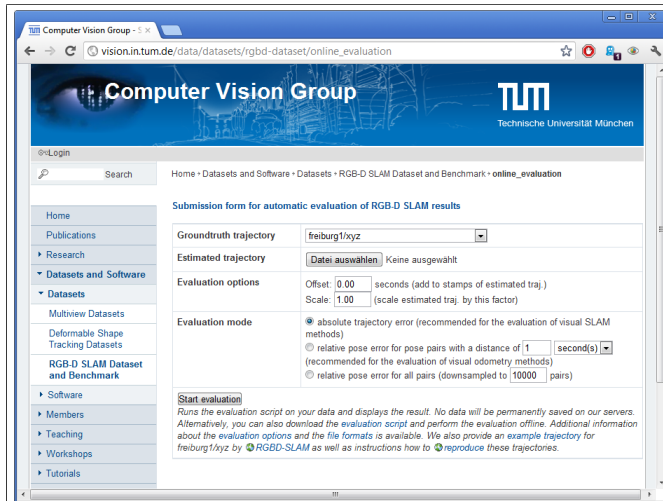
$$RMSE(E_{1:n}) := \left(\frac{1}{n} \sum_{i=1}^n \|trans(E_i)\|^2 \right)^{1/2}$$

- Reference implementations for both evaluation metrics available
- Output: RMSE, Mean, Median (as text)
- Plot (png/pdf, optional)

Visual Navigation for Flying Robots

48

Dr. Jürgen Sturm, Computer Vision Group, TUM



Summary – TUM RGB-D Benchmark

- Dataset for the evaluation of RGB-D SLAM systems
- Ground-truth camera poses
- Evaluation metrics + tools available

Visual Navigation for Flying Robots

51

Dr. Jürgen Sturm, Computer Vision Group, TUM

Discussion on Benchmarks

Pro:

- Provide objective measure
- Simplify empirical evaluation
- Stimulate comparison

Con:

- Introduce bias towards approaches that perform well on the benchmark (overfitting)
- Evaluation metrics are not unique (many alternative metrics exist, choice is subjective)

Visual Navigation for Flying Robots

52

Dr. Jürgen Sturm, Computer Vision Group, TUM

Three Phases of Evolution in Research

- Novel research problem appears (e.g., market launch of Kinect, quadcopters, ...)
 - Is it possible to do something at all?
 - Proof-of-concept, qualitative evaluation
- Consolidation
 - Problem is formalized
 - Alternative approaches appear
 - Need for quantitative evaluation and comparison
- Settled
 - Benchmarks appear
 - Solid scientific analysis, text books, ...

Visual Navigation for Flying Robots

53

Dr. Jürgen Sturm, Computer Vision Group, TUM

Final Exam

- Oral exam **in teams** (2-3 students)
- At least 15 minutes per student
→ individual grades
- Questions will address
 - Your project
 - Material from the exercise sheets
 - Material from the lecture

Visual Navigation for Flying Robots

54

Dr. Jürgen Sturm, Computer Vision Group, TUM

Exercise Sheet 6

- Prepare final presentation
- Proposed structure: 4-5 slides
 1. Title slide with names + motivating picture
 2. Approach
 3. Results (video is a plus)
 4. Conclusions (what did you learn in the project?)
 5. Optional: Future work, possible extensions
- Hand in slides before Tue, July 17, 10am (!)

Time for Questions