

Representing and Modeling Space

McGill COMP 765

Sept 7th, 2017



* The topology of quantum wormholes will be left for self-study.



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Goals today

- Introduce models of a robot's movements in space
- Introduce models to represent the space around the robot (environment)
- A few key robotics problems: forward and inverse kinematics, forward and inverse dynamics
- Case studies on a few of our favorite 5 robots
- Intro to estimating spatial representations under uncertainty

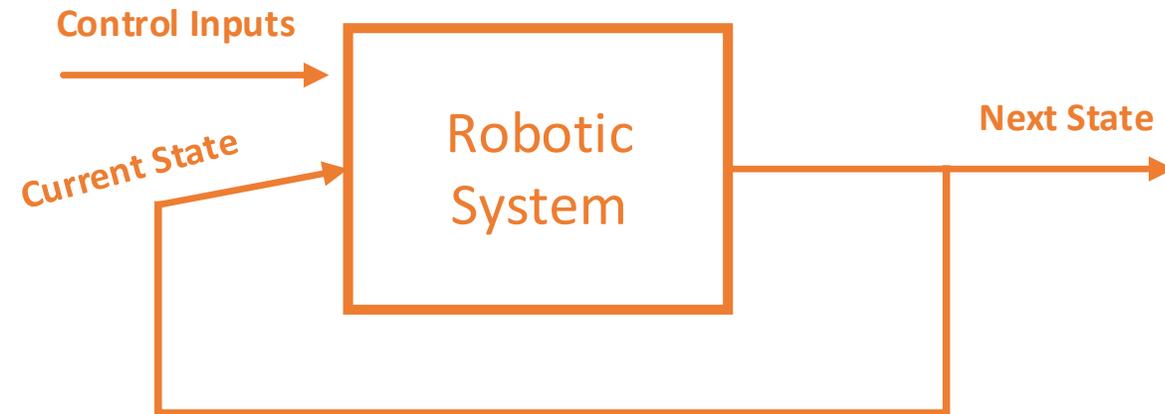


McGill

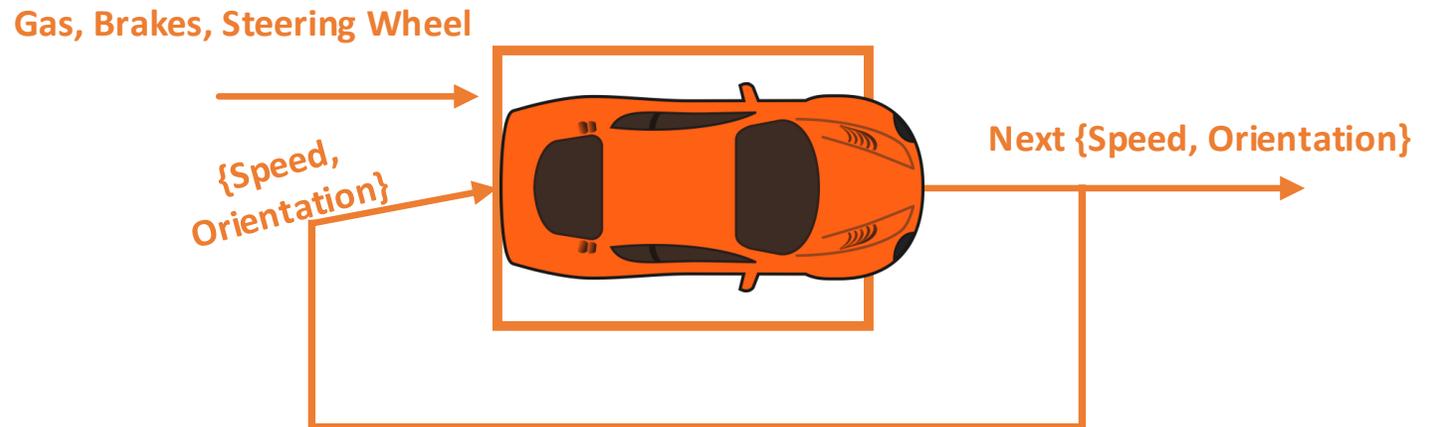
School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Physical models of how systems move



Kinematics & Dynamics:
Idealized models of
robotic motion



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Kinematics and Dynamics

- Kinematics considers models of locomotion independently of external forces and control:
 - For example, how the speed of a car's wheels affect the motion of its chassis.
- Dynamics considers models of locomotion as functions of their control inputs and state – considering all present forces:
 - For example, how a quadrotor will move when the rotor's fight gravity

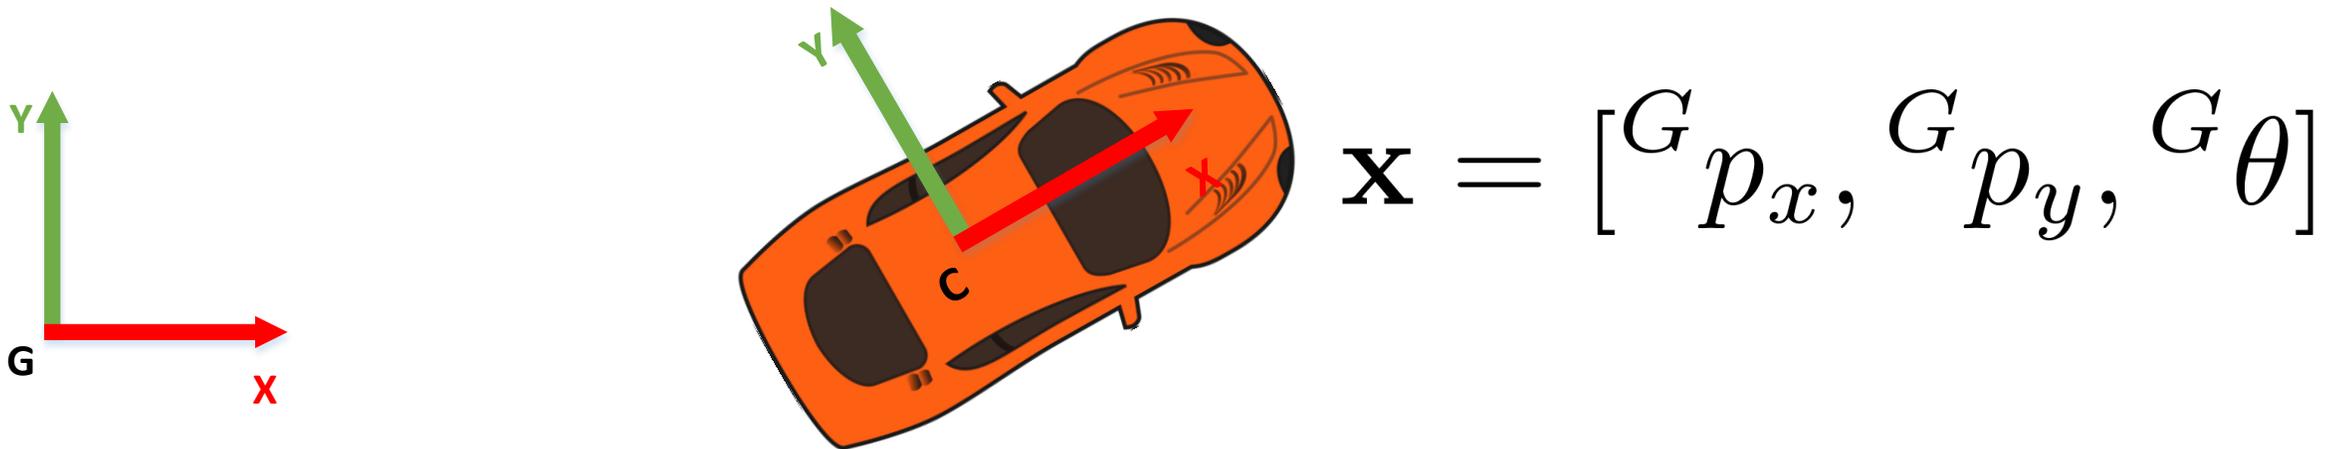


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Example: Kinematics of a simple car - state



State = [Position and orientation]

Position of the car's frame of reference C with respect to a fixed frame of reference G , expressed in frame G .

The angle is the orientation of frame C with respect to G .



McGill

School of Computer Science
Centre for Intelligent Machines

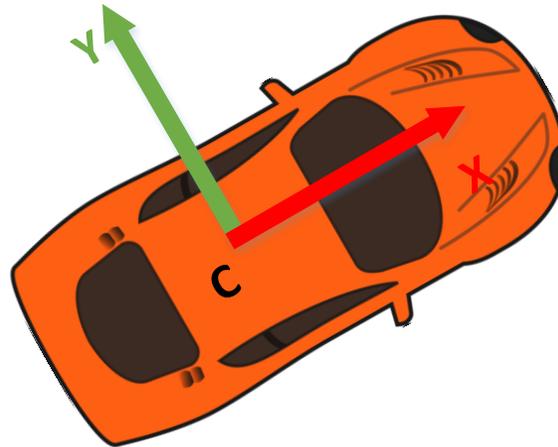
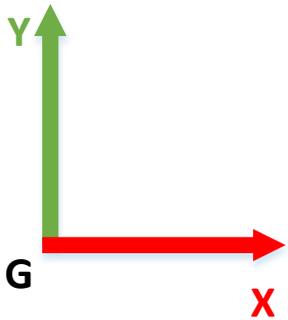
Intelligent Robotics

Note on Inertial frames of reference

- G, the global frame of reference is fixed, i.e. with zero velocity in our previous example.
- If a robot's pose is known in the global frame, life is good:
 - It can correctly report its position
 - We can compute the direction and distance to a goal point
 - We can avoid collisions with obstacles also known in global frame
- We may also know the pose in another frame: e.g. of a map, sensor, or relative to its starting point. These are often sub-steps for us.



Example: the dynamics of a simple car - controls

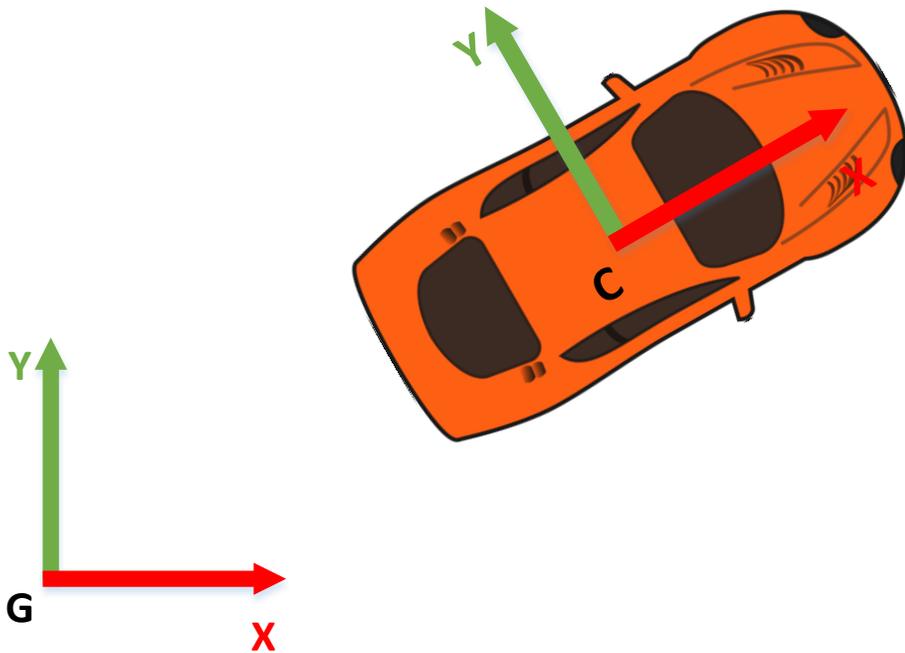


$$\mathbf{u} = [{}^C v_x, {}^C \omega_z]$$

Controls = [Forward speed and angular velocity]
Linear velocity and angular velocity of the car's frame of reference C with respect to a fixed frame of reference G, expressed in coordinates of C.



The dynamical system of a simple car



$$\dot{p}_x = v_x \cos(\theta)$$

$$\dot{p}_y = v_x \sin(\theta)$$

$$\dot{\theta} = \omega_z$$

Note: reference frames have been removed for readability.



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Forward and Inverse Kinematics and Dynamics

- Kinematics:
 - Forward – what robot position resulting from a given configuration? Example: given joint angles, output the pose of each joint in 3-space.
 - Inverse – which configuration will result in a desired robot position? Example: to grasp a point in 3-space, solve for the right joint angles.
- Dynamics:
 - Forward – what robot motion will result from given input forces? Example: predict the motion of a quadrotor spinning only one of its propellers.
 - Inverse – which input forces will result in a desired robot motion? Example: in order to stabilize gravity, solve for the force is needed at each propeller.

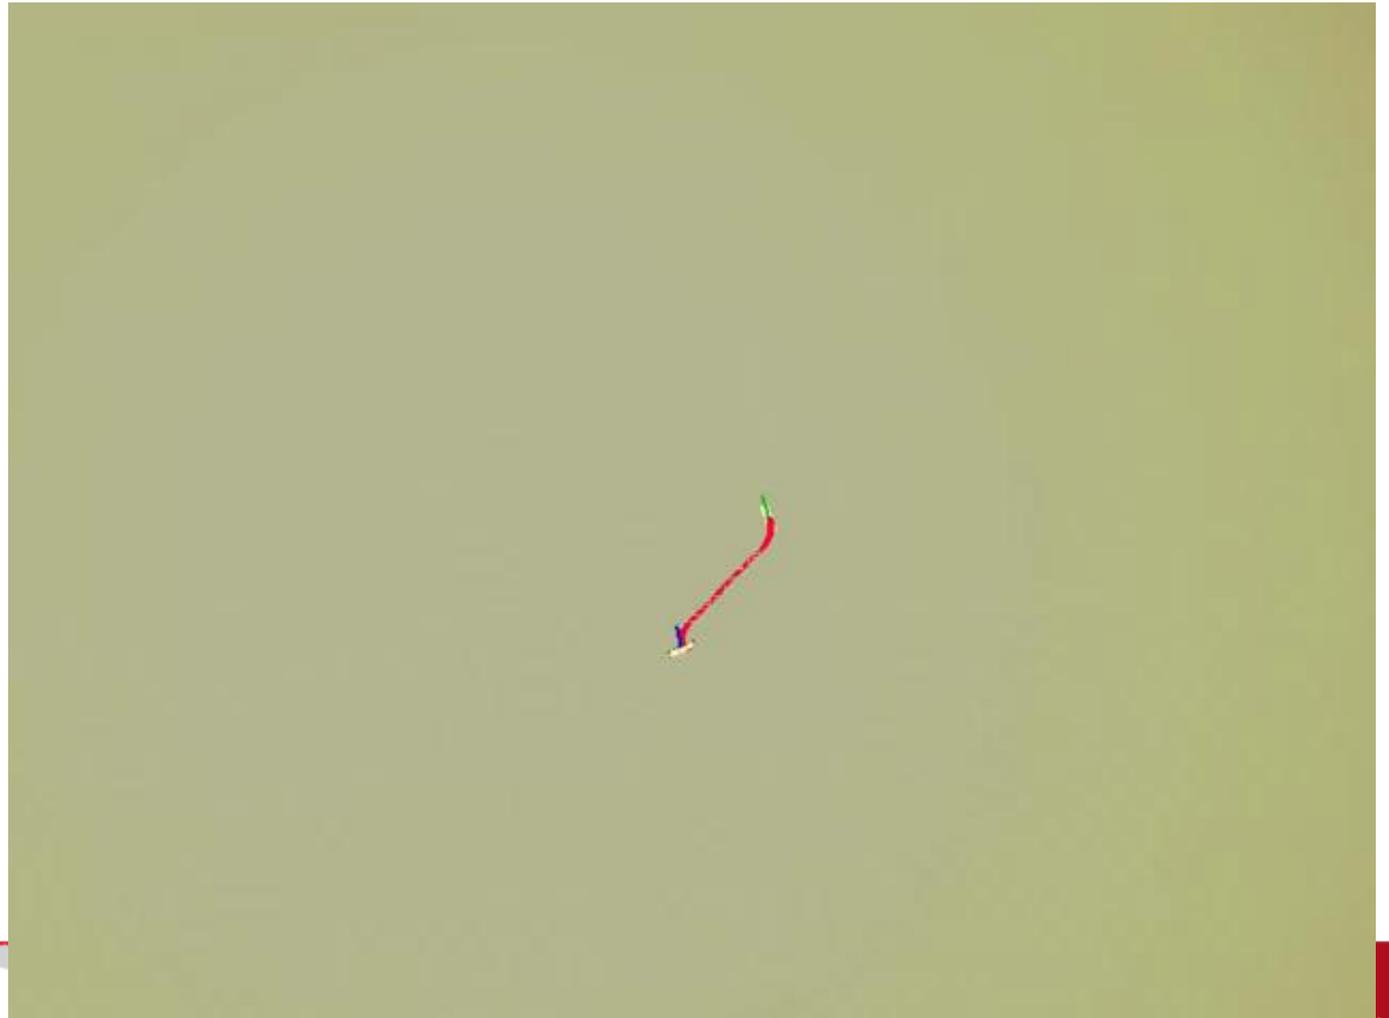


Special case of simple car: Dubins car

- Can only go forward
- Constant speed

$$C v_x = \text{const} > 0$$

- You only control the angular velocity

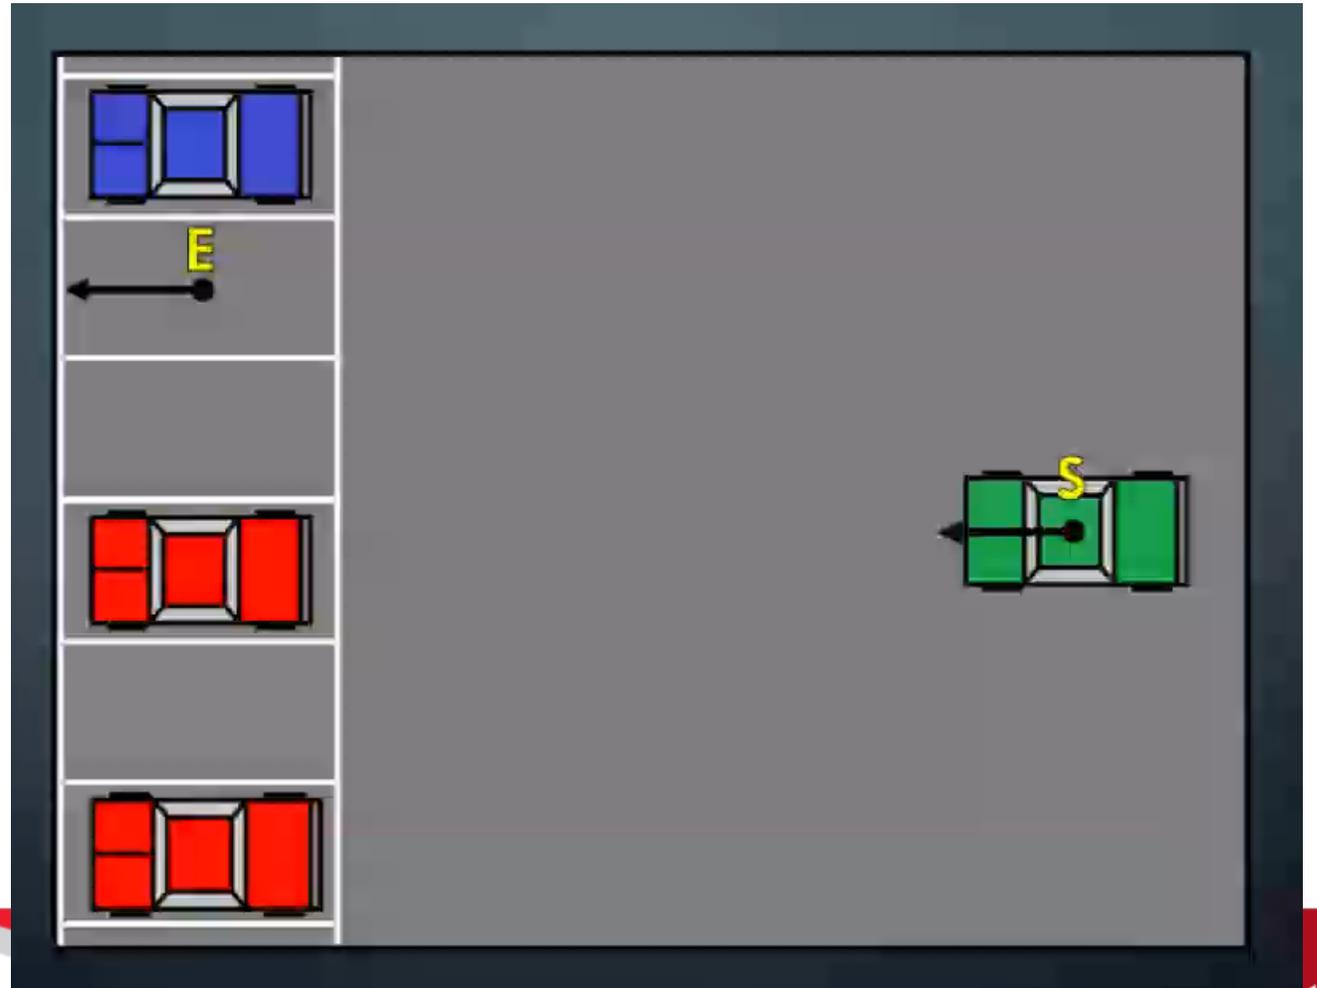


Special case of simple car: Dubins car

- Can only go forward
- Constant speed

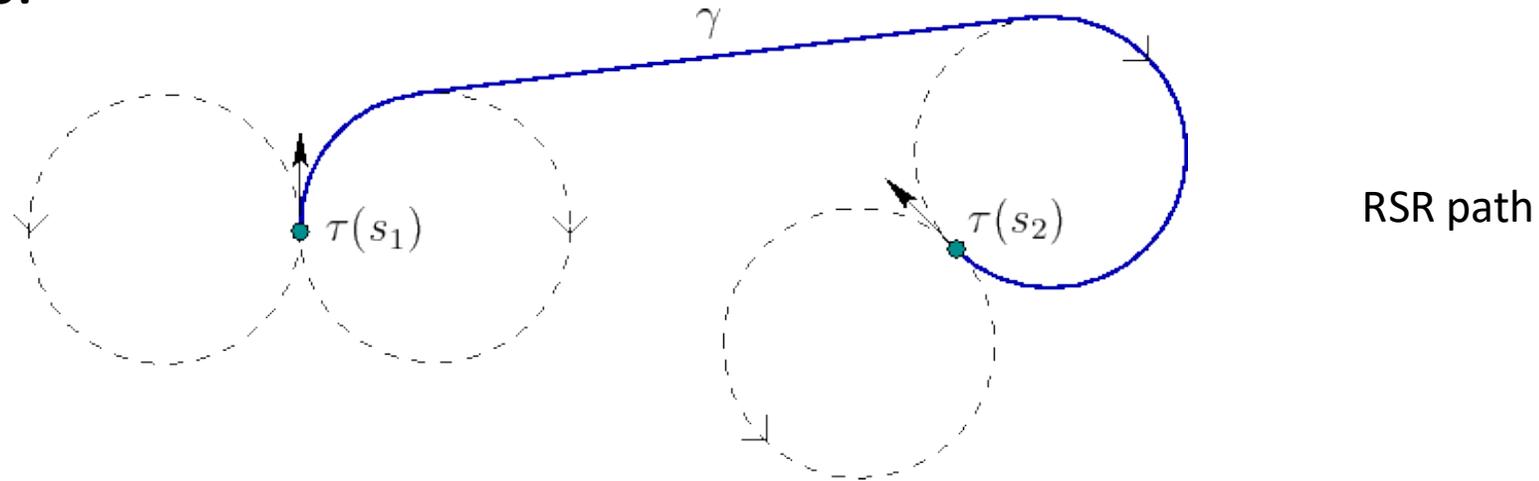
$$C v_x = \text{const} > 0$$

- You only control the angular velocity

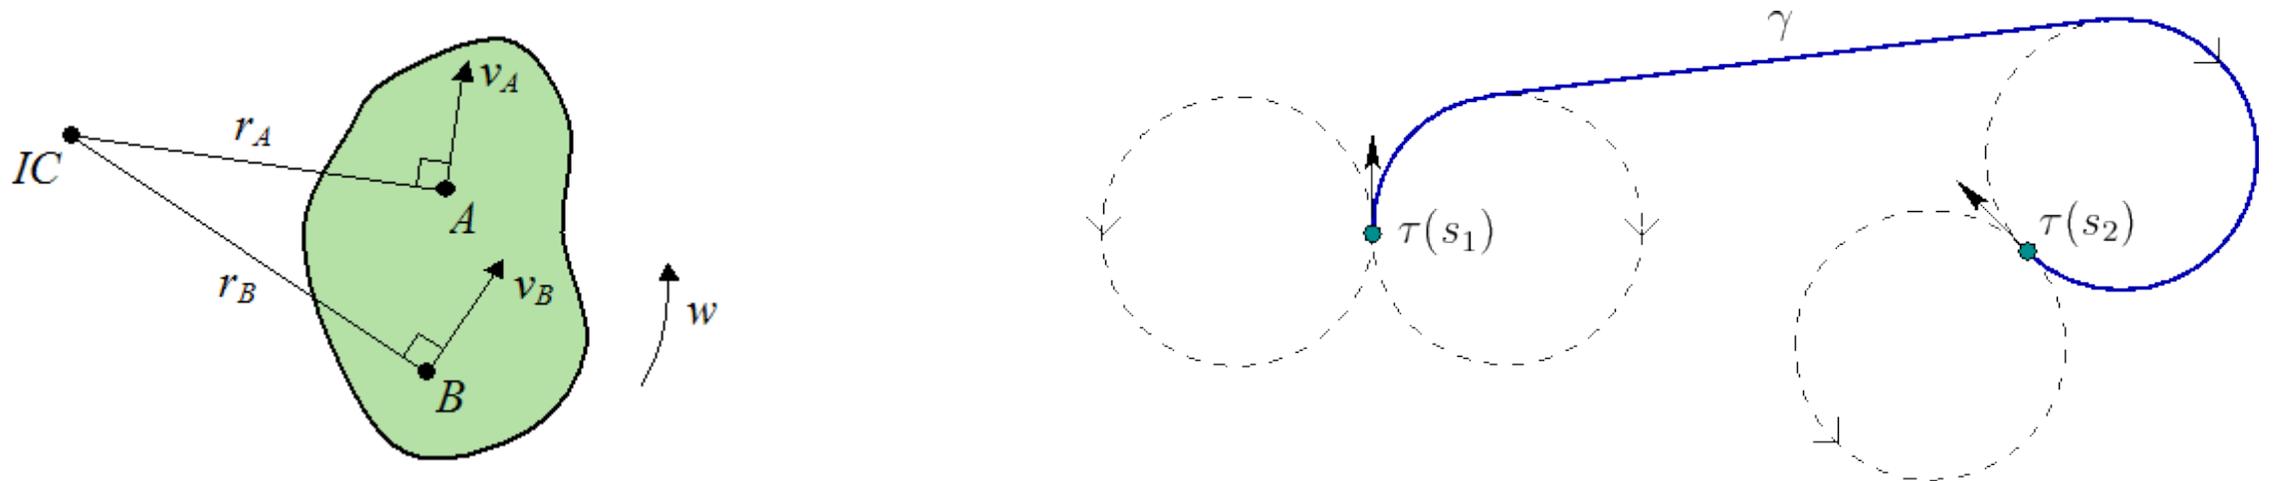


Dubins car: motion primitives

- The path of the car can be decomposed to L(left), R(right), S(traight) segments.



Instantaneous Center of Rotation



IC = Instantaneous Center of Rotation
The center of the circle circumscribed by the turning path.
Undefined for straight path segments.



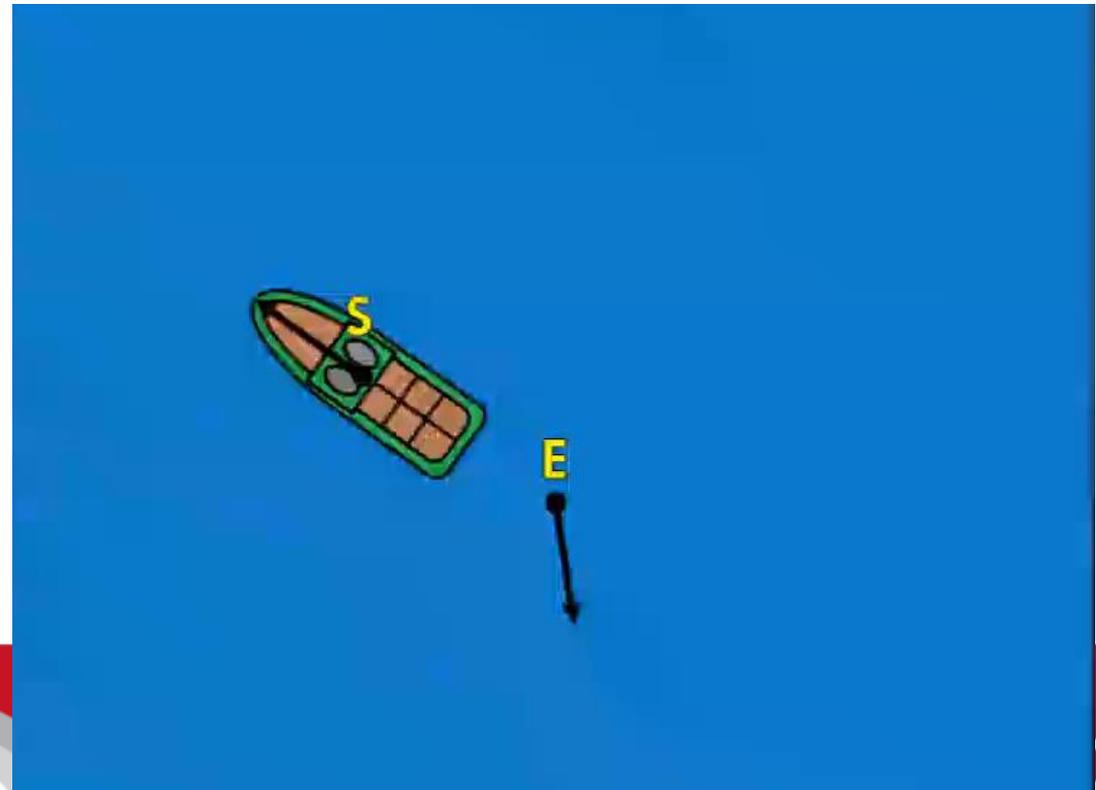
McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Dubins car \longrightarrow Dubins boat

- Why do we care about a car that can only go forward?
- Because we can also model idealized airplanes and boats
- Dubins boat = Dubins car



Holonomic constraints

- Equality constraints on the state of the system, but not on the higher-order derivatives:

$$\mathbf{f}(\mathbf{x}, t) = 0$$

- For example, if you want to constrain the state to lie on a circle:

$$\|\mathbf{x}\|^2 - 1 = 0$$

- Another example: train tracks are a holonomic constraint.



Non-holonomic constraints

- Equality constraints that involve the derivatives of the state (e.g. velocity) in a way that it cannot be integrated out into holonomic constraints, i.e.

$$\mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, t) = 0$$

but not

$$\mathbf{f}(\mathbf{x}, t) = 0$$



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

The Dubins car is non-holonomic

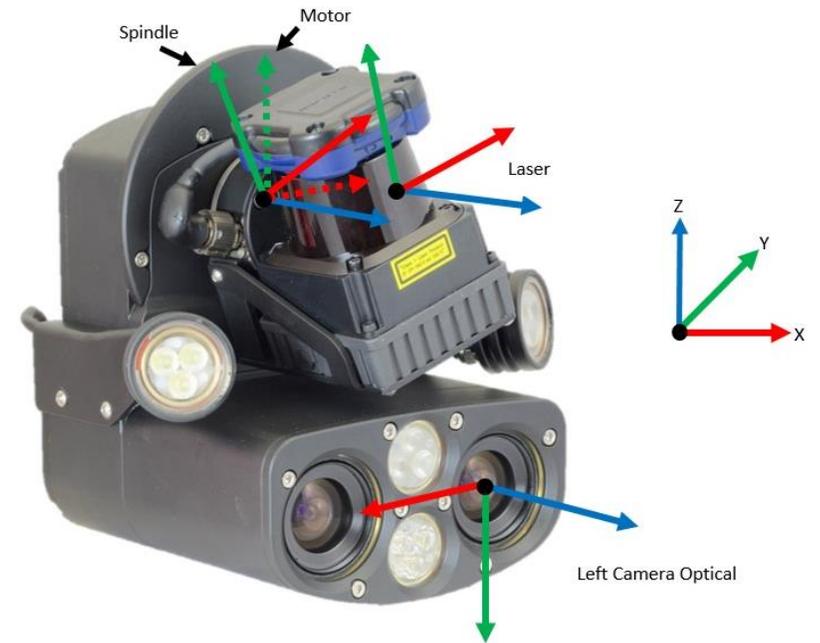
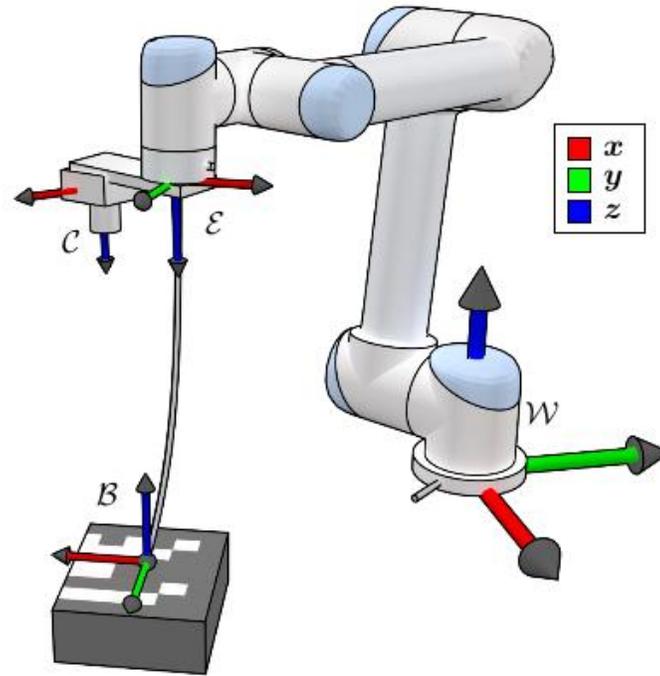
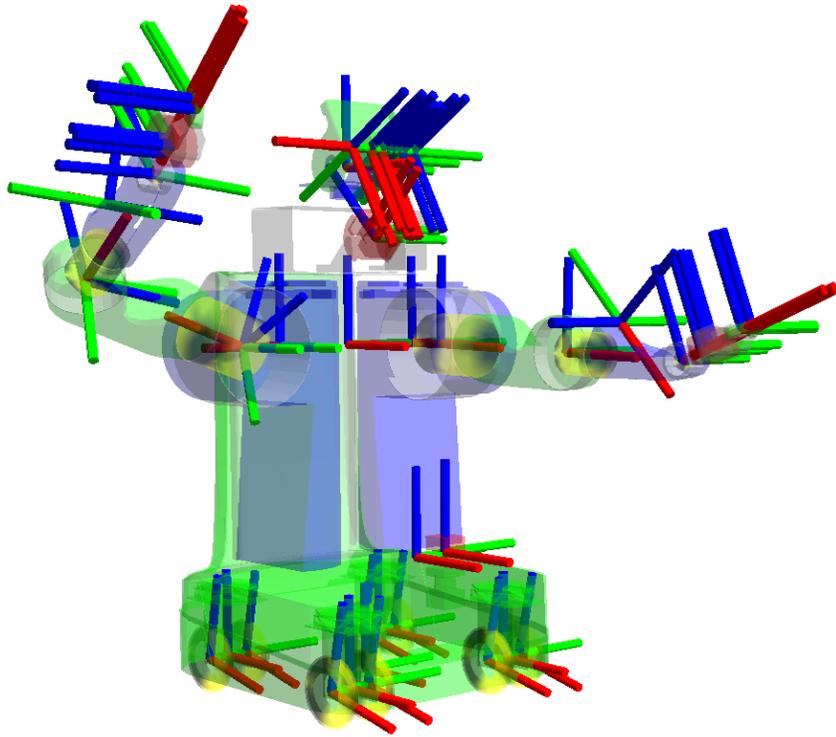
- Dubins car is constrained to move straight towards the direction it is currently heading. It cannot move sideways. It needs to “parallel park” to move laterally.
- In a small time interval dt the vehicle is going to move by δp_x and δp_y in the global frame of reference. Then from the dynamical system:

$$\begin{array}{l} \delta p_x \sin(\theta) = v_x \cos(\theta) \sin(\theta) \\ \delta p_y \cos(\theta) = v_x \sin(\theta) \cos(\theta) \end{array} \quad \longrightarrow \quad \delta p_x \sin(\theta) = \delta p_y \cos(\theta) \quad \longrightarrow \quad v_x \sin(\theta) = v_y \cos(\theta)$$

Car is constrained to move along the line of current heading,
i.e. non-holonomic



3D frames of reference are everywhere in robotics



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

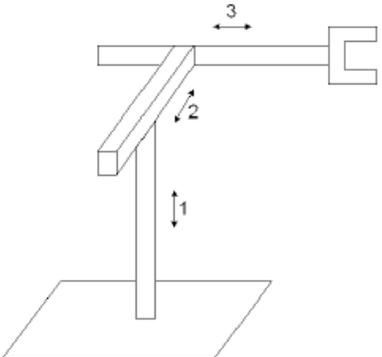
Manipulators

- Robot arms, industrial robot
 - Rigid bodies (links) connected by joints
 - Joints: revolute or prismatic
 - Drive: electric or hydraulic
 - End-effector (tool) mounted on a flange or plate secured to the wrist joint of robot

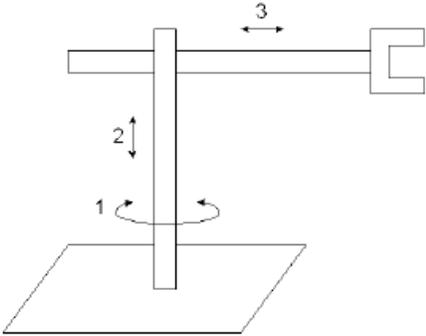


Manipulators

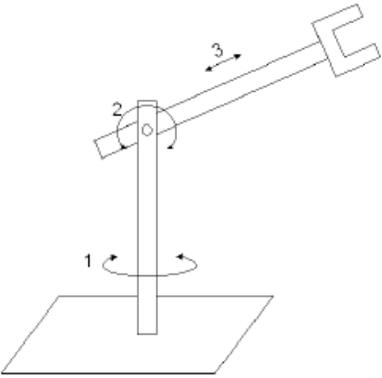
- Robot Configuration:



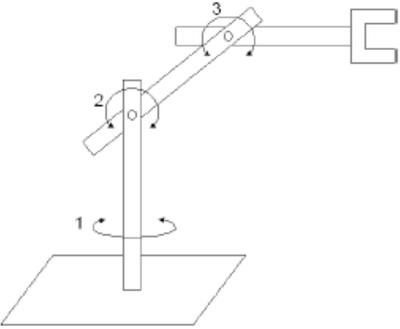
Cartesian: PPP



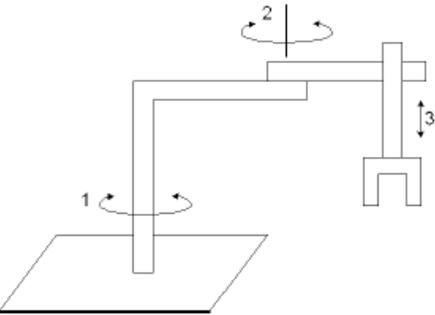
Cylindrical: RPP



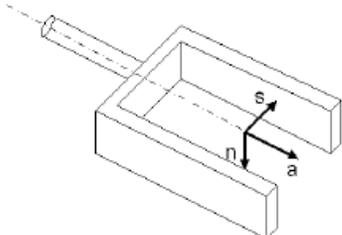
Spherical: RRP



Articulated: RRR



SCARA: RRP
(Selective Compliance
Assembly Robot Arm)



Hand coordinate:
n: normal vector; **s**: sliding vector;
a: approach vector, normal to the
tool mounting plate

Manipulators

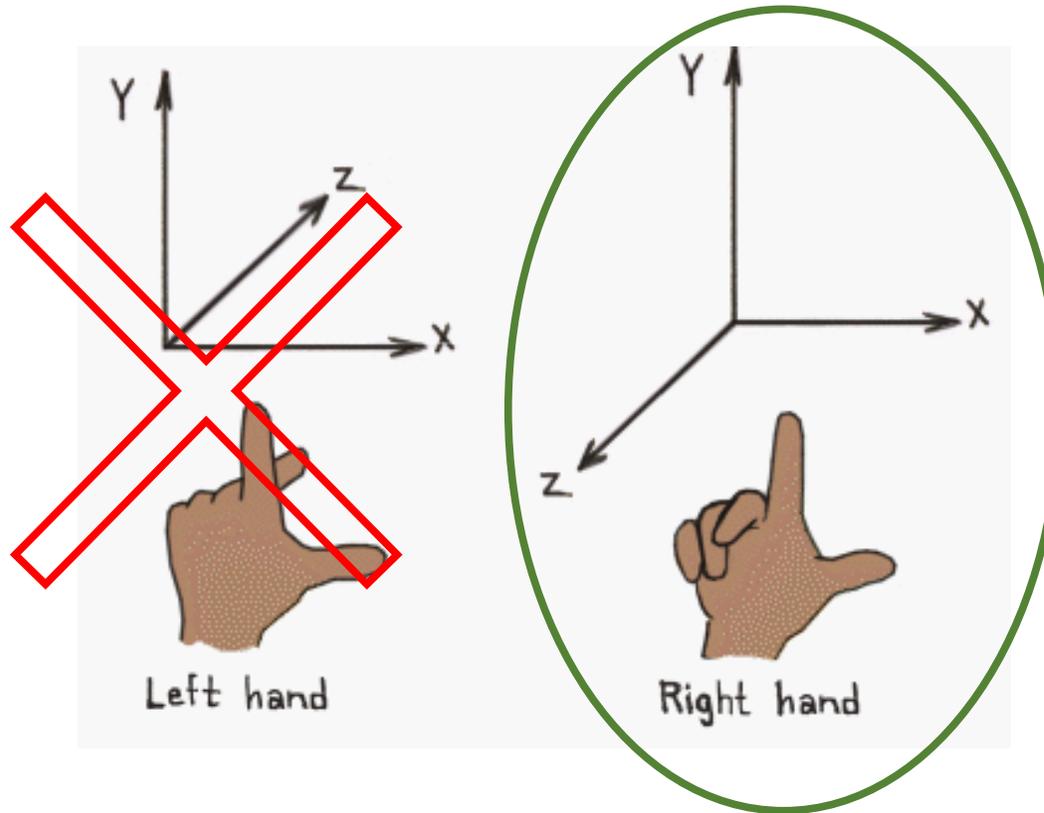
- Robot Specifications
 - Number of Axes
 - Major axes, (1-3) => Position the wrist
 - Minor axes, (4-6) => Orient the tool
 - Redundant, (7-n) => reaching around obstacles, avoiding undesirable configuration
 - Degree of Freedom (DOF)
 - Workspace
 - Payload (load capacity)
 - Precision v.s. Repeatability



Which one is more important?



Right-handed vs left-handed frames



Unless otherwise specified,
we use right-handed
frames in robotics



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Why do we need to use so many frames?

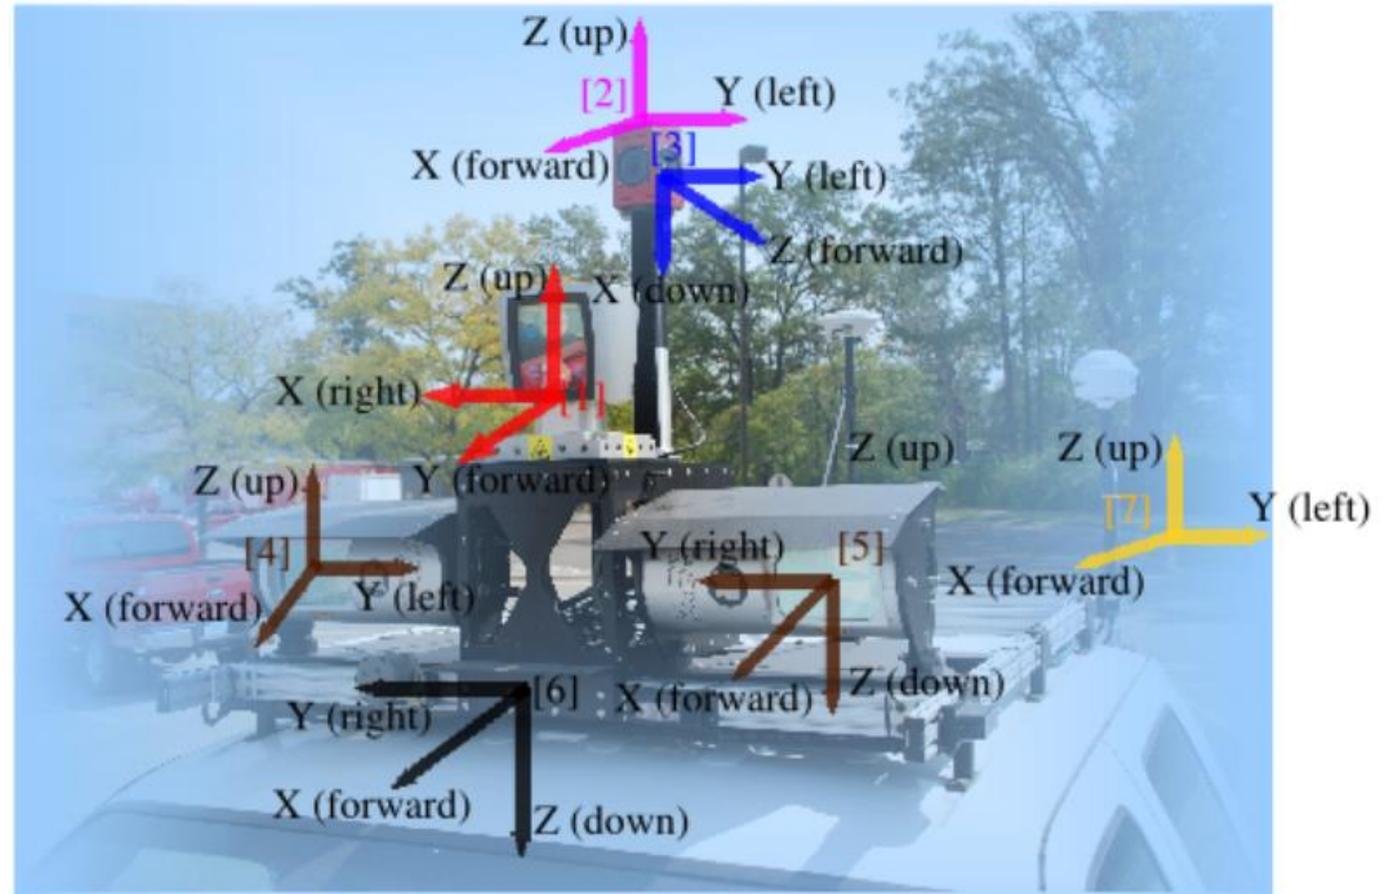
- Because we want to reason and express quantities relative to their local configuration.
- For example: “grab the bottle behind the cereal bowl”
- Many algorithms in this class are mostly about representing frames of reference and reasoning about how to express quantities in one frame to quantities in the other.



McGill

School of Computer Sci
Centre for Intelligent M

The frames of self-driving



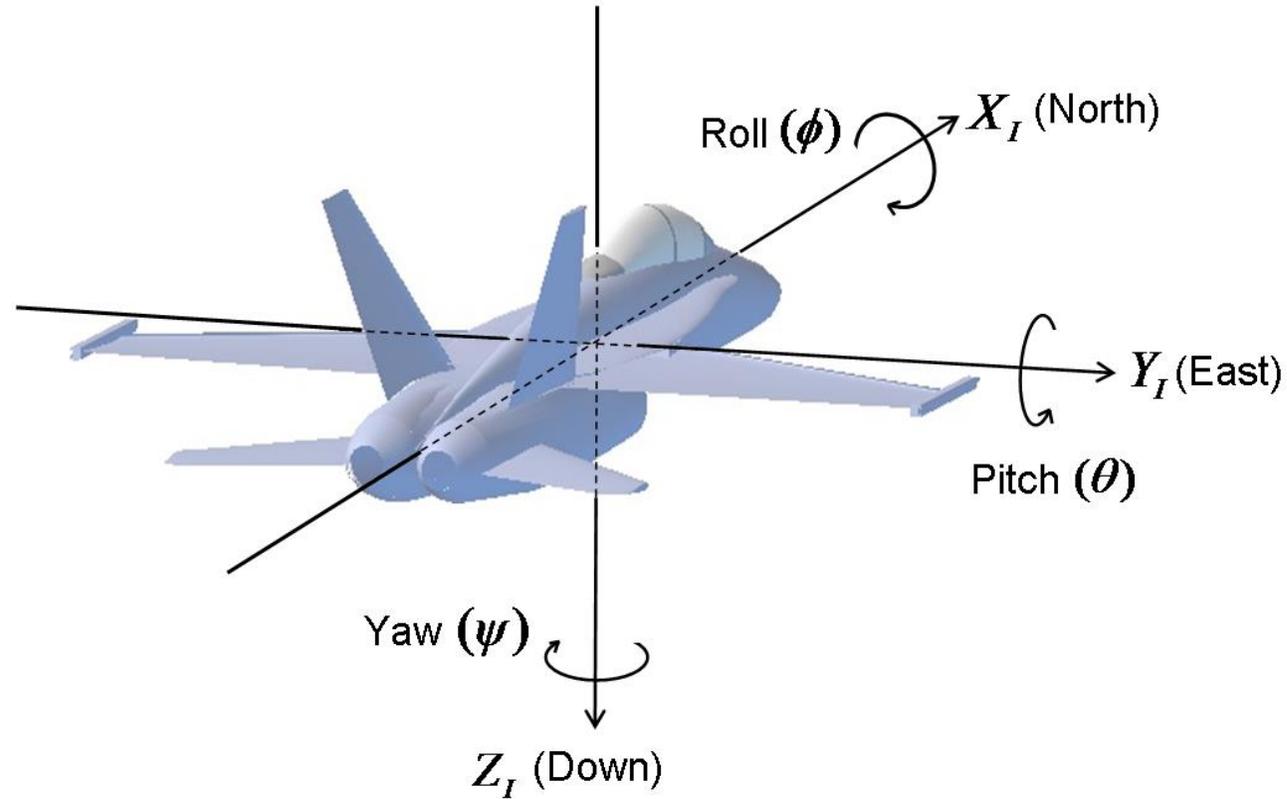
- [1] Velodyne, [2] Ladybug3 (actual location: center of camera system),
[3] Ladybug3 Camera 5, [4] Right Riegl, [5] Left Riegl,
[6] Body Frame (actual location: center of rear axle)
[7] Local Frame (Angle between the X-axis and East is known)



McGill

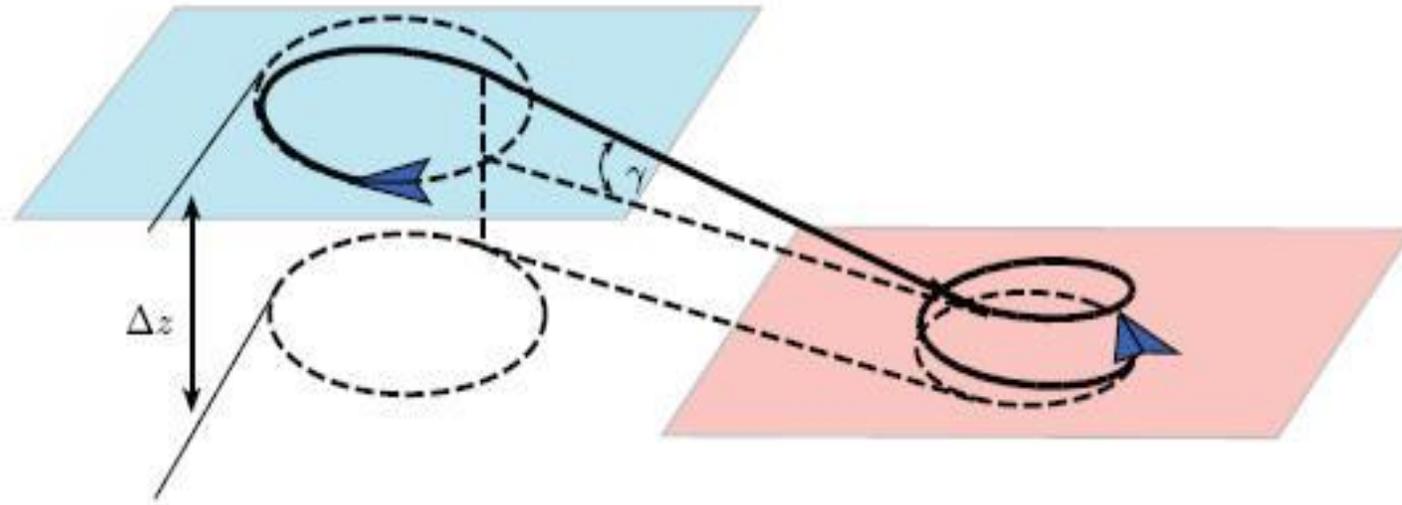
School of Computer Science
Centre for Intelligent Machines

Representing Rotations in 3D: Euler Angles



Dubins car \longrightarrow Dubins airplane in 3D

- Pitch angle ϕ and forward velocity determine descent rate
- Yaw angle θ and forward velocity determine turning rate



(a) The 3D view of the path.

$$\dot{p}_x = v_x \cos(\theta) \sin(\phi)$$

$$\dot{p}_y = v_x \sin(\theta) \sin(\phi)$$

$$\dot{p}_z = v_x \cos(\phi)$$

$$\dot{\theta} = \omega_z$$

$$\dot{\phi} = \omega_y$$

θ is yaw

ϕ is pitch



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Specification ambiguities in Euler Angles

- Need to specify the axes which each angle refers to.
- There are **12 different valid combinations** of fundamental rotations.
Here are the possible axes:
 - z-x-z, x-y-x, y-z-y, z-y-z, x-z-x, y-x-y
 - x-y-z, y-z-x, z-y-x, x-z-y, z-y-x, y-x-z



Specification ambiguities in Euler Angles

- Need to specify the axes which each angle refers to.
- There are **12 different valid combinations** of fundamental rotations. Here are the possible axes:
 - z-x-z, x-y-x, y-z-y, z-y-z, x-z-x, y-x-y
 - x-y-z, y-z-x, z-y-x, x-z-y, z-y-x, y-x-z
- E.g.: x-y-z rotation with Euler angles (θ, ϕ, ψ) means a rotation: expressed as a sequence of simple rotations $R_x(\theta)R_y(\phi)R_z(\psi)$



Specification ambiguities in Euler Angles

Simple rotations can be counter-clockwise or clockwise.
This gives **another 2 possibilities**.

$$\mathbf{R}_z(\alpha) := \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{C}_z(\alpha) := \begin{bmatrix} \cos \alpha & \sin \alpha & 0 \\ -\sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix}$$



How to handle this horrible situation?

- Use accepted conventions. For example, we will use ROS and Gazebo. Make sure your code fits with their visualizers, simulators etc.
- When inventing something new, use extreme documentation.
- When starting with someone else's implementation, do not assume anything and start by visualizing a wide range of motions.

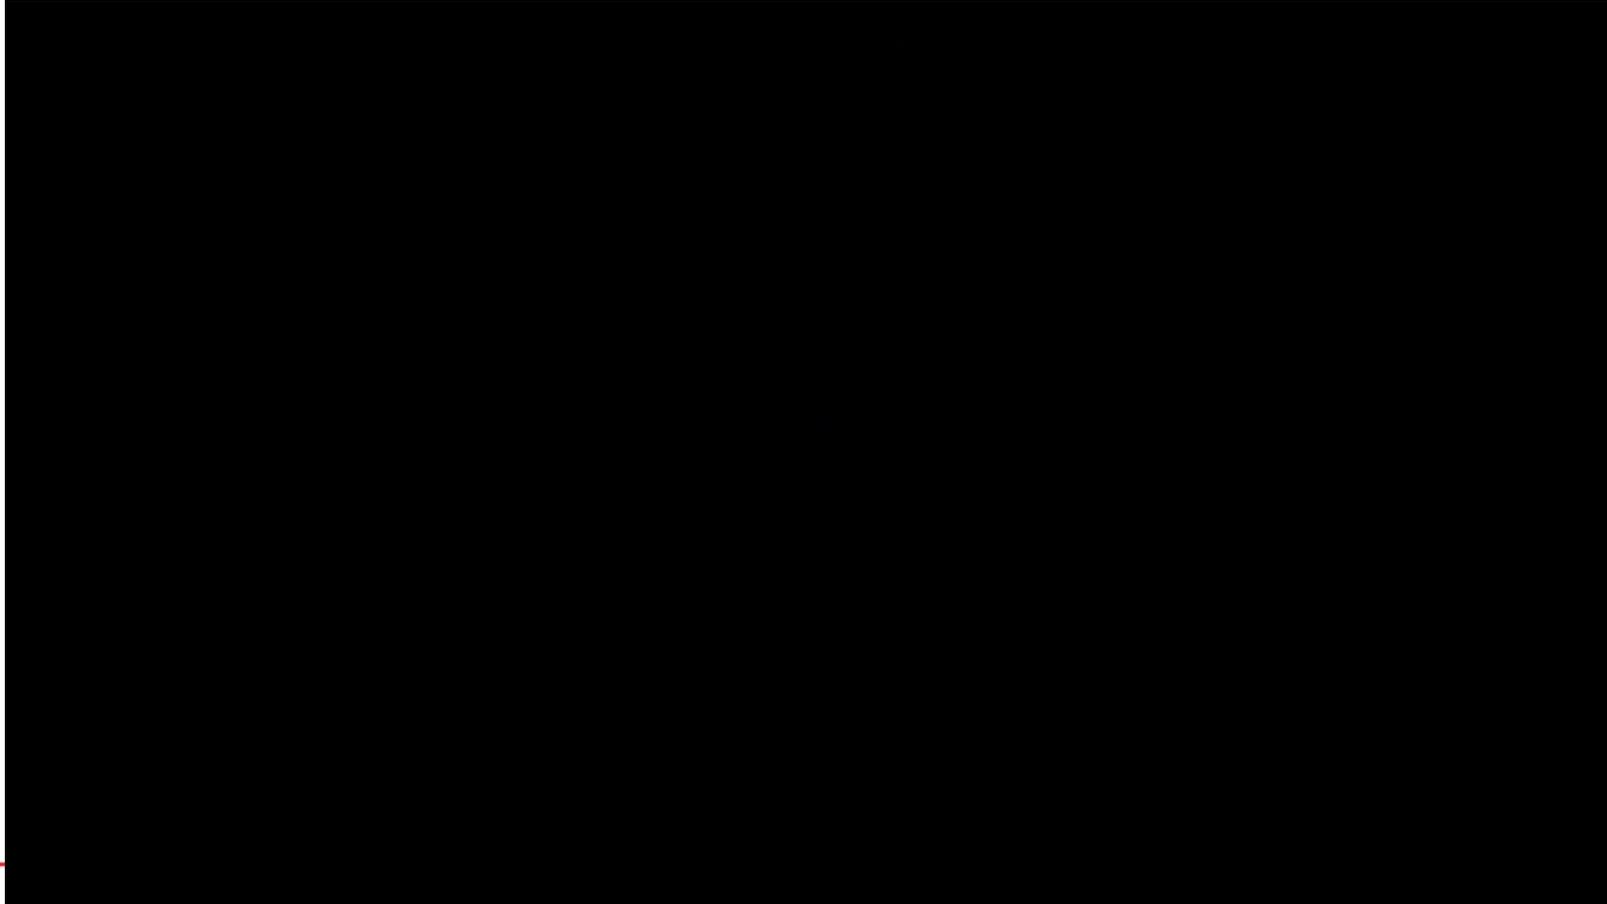


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Another problem with Euler angles: Gimbal Lock



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Beyond Euler

- Research papers will often use more complex angle representations:
 - Rotation matrices
 - Rodrigues' axis-angle
 - Quaternions
- This to avoid gimbal lock, imprecise definition and to exploit some particular properties.
- We will not go deep into this in the interest of time

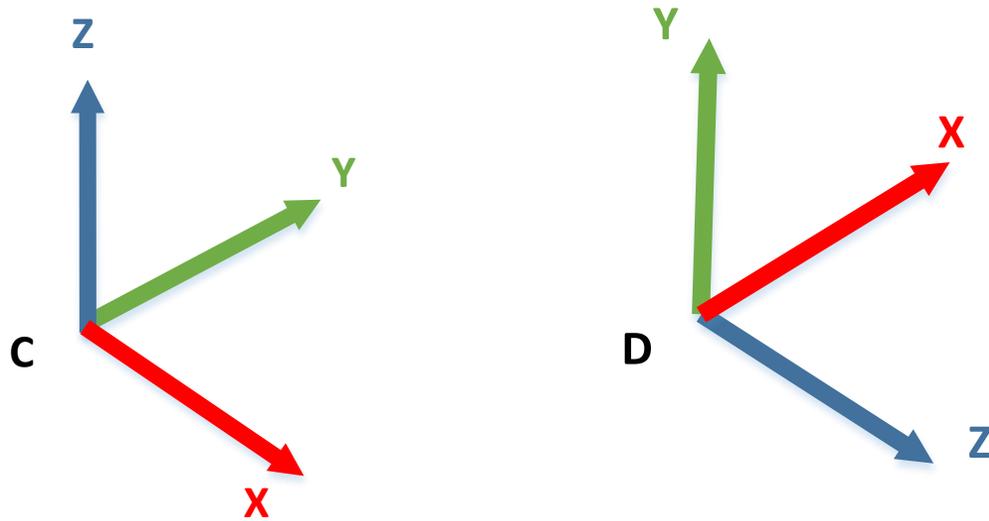


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Example: finding a rotation matrix that rotates one vector to another



$${}^D_C \mathbf{R} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$$

This matrix transforms the x-axis of frame C to the z-axis of frame D. Same for y and z axes.

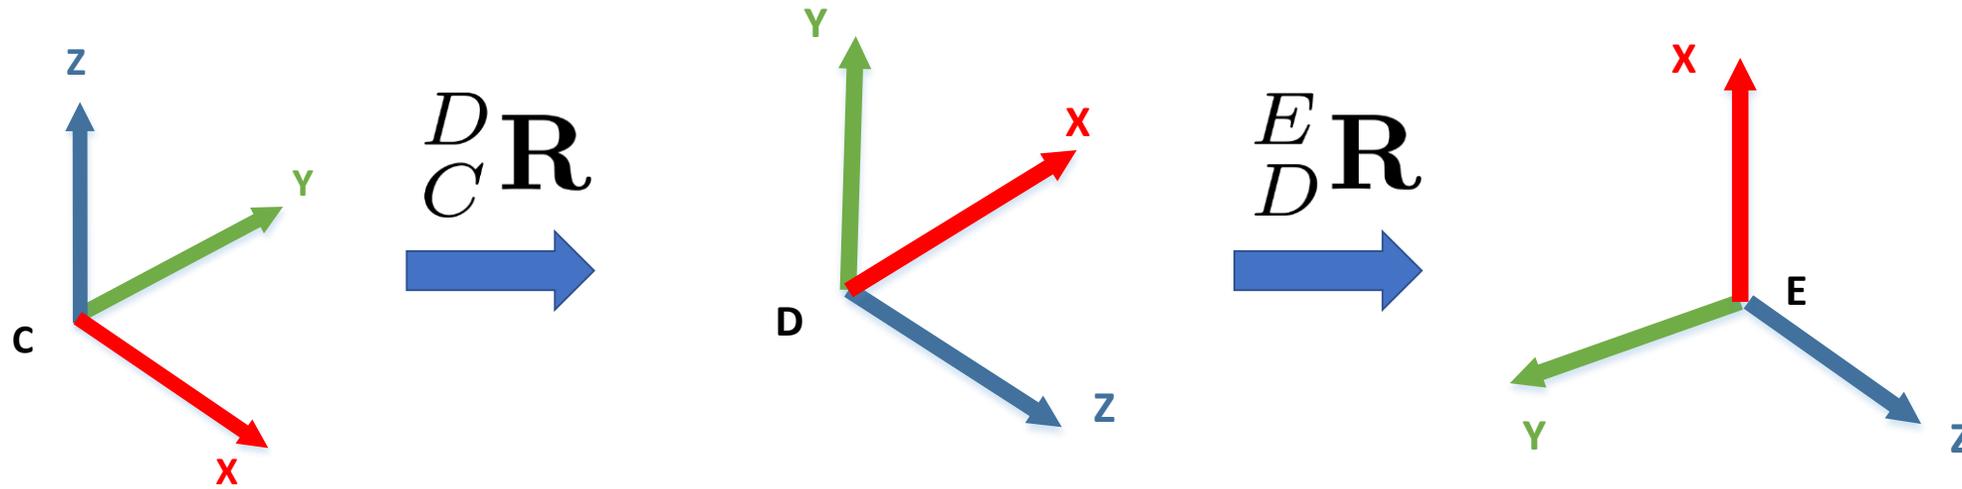


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Compound rotations



$$\frac{E}{C}\mathbf{R} = \frac{E}{D}\mathbf{R}\frac{D}{C}\mathbf{R}$$



Passive Dynamics

- Dynamics of systems that operate without drawing (a lot of) energy from a power supply.
- Interesting because biological locomotion systems are more efficient than current robotic systems.



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Passive Dynamics

- Dynamics of systems that operate without drawing (a lot of) energy from a power supply.
- Usually propelled by their own weight.
- Interesting because biological locomotion systems are more efficient than current robotic systems.

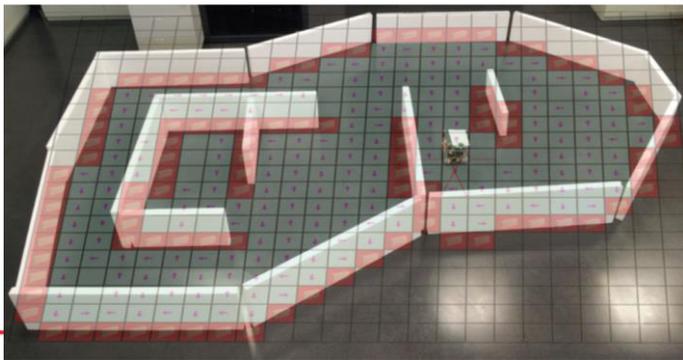
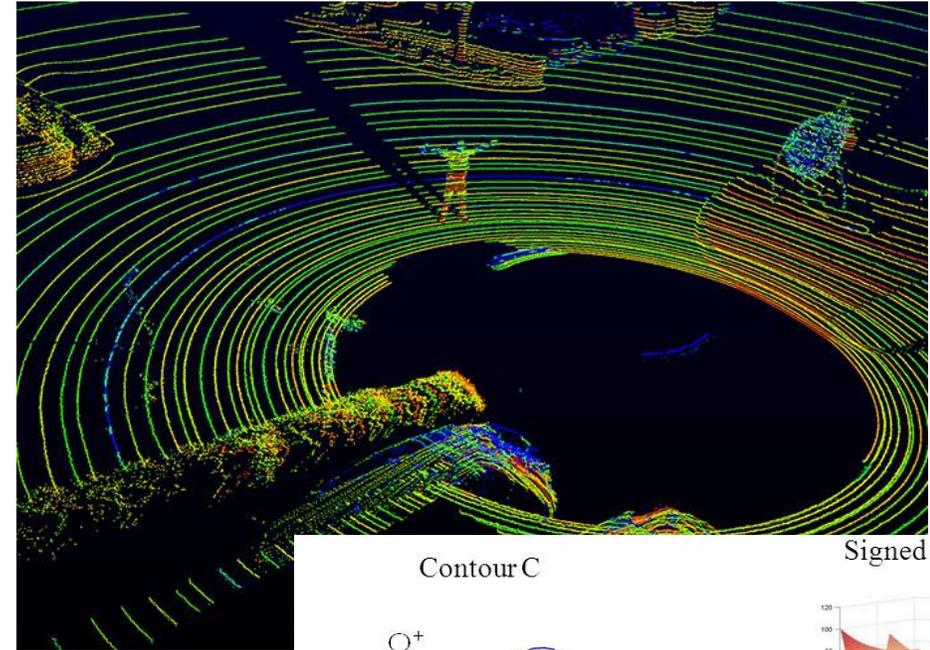
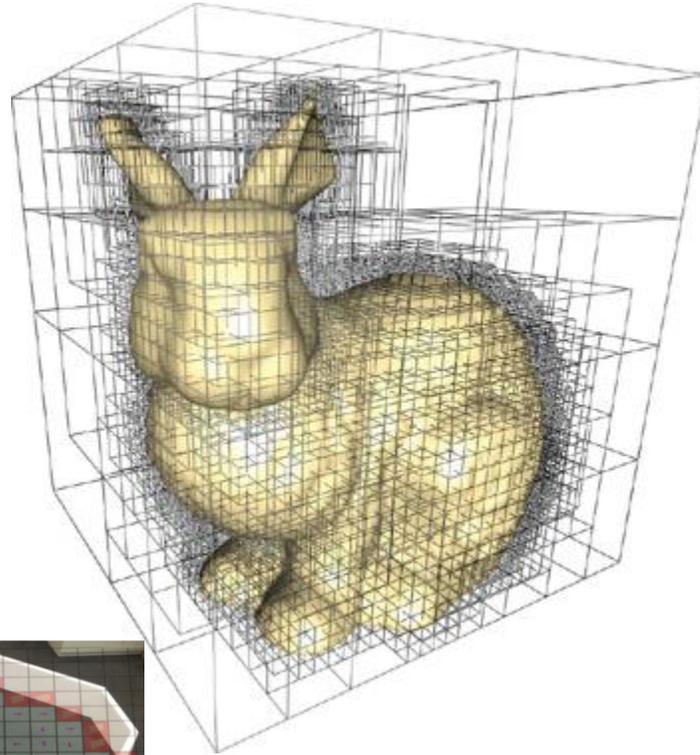


McGill

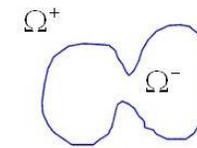
School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

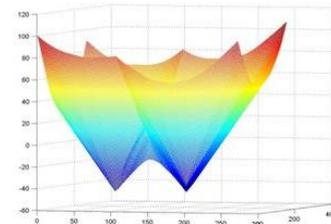
Representing the space around a robot: maps



Contour C



Signed distance function



A signed distance function is defined by :

$$\phi(x) = \begin{cases} \text{dist}(x, C) & \text{if } x \text{ is outside } C \\ 0 & x \in C \\ -\text{dist}(x, C) & \text{if } x \text{ is inside } C \end{cases}$$



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Categories of maps

- Metric
 - Map accurately represents lengths and angles
- Topological
 - Map is reduced to a graph representation of the structure of free space
- Topometric
 - Atlas: a combination of local metric maps (nodes) connected via edges
- Sequence of raw time-series observations (e.g. video)
 - No metric or topological information directly represented by the map



Typical operations on maps

- Distance and direction to closest obstacle
- Collision detection: is a given robot configuration in free space?
- Map merging / alignment
- Occupancy updates
- Raytracing



Metric Maps

(all we will do right now)

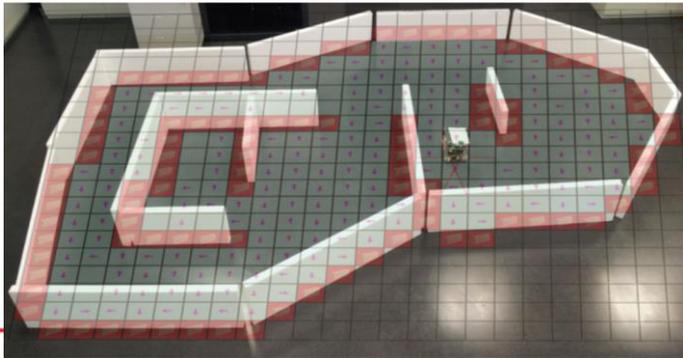
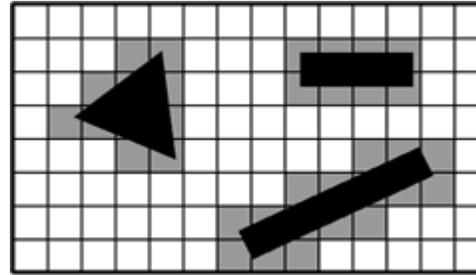
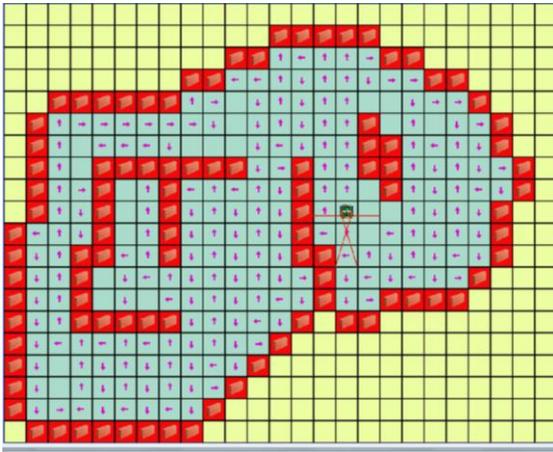


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Occupancy Grids



- Each cell contains either:
- unknown/unexplored (grey)
 - probability of occupation

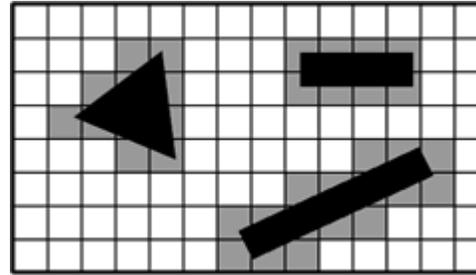
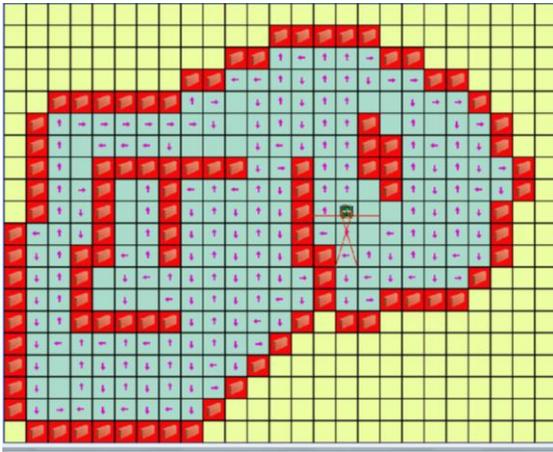


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Occupancy Grids

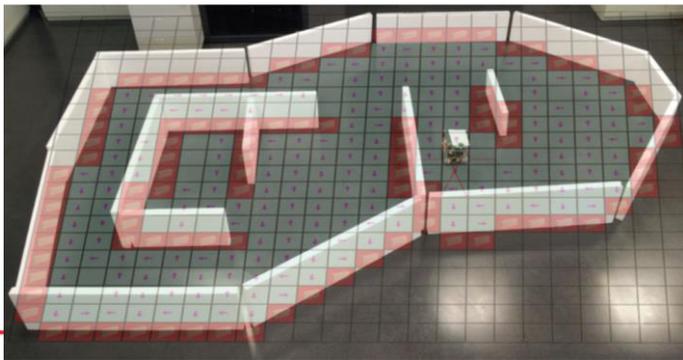


Advantages:

- $O(1)$ occupancy lookup and update
- Supports image operations

Disadvantages:

- Doesn't scale well in higher dimensions



Each cell contains either:

- unknown/unexplored (grey)
- probability of occupation

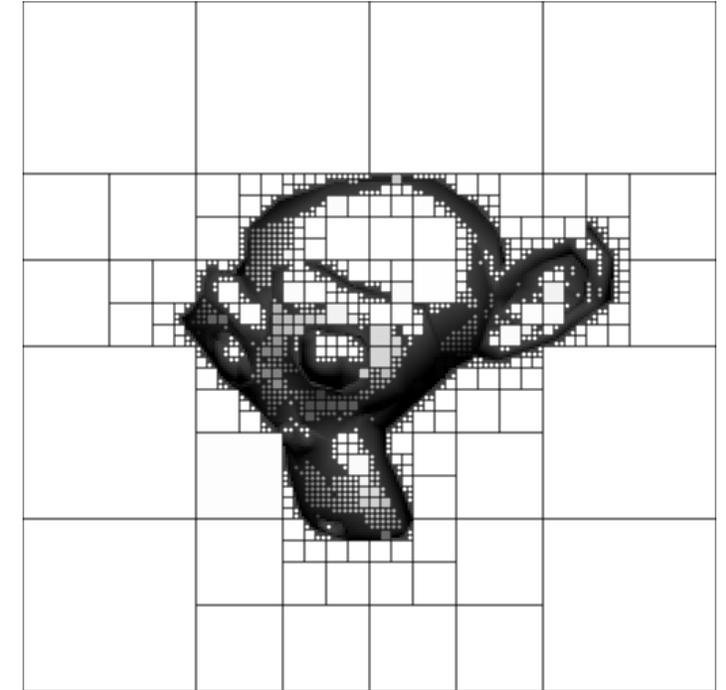
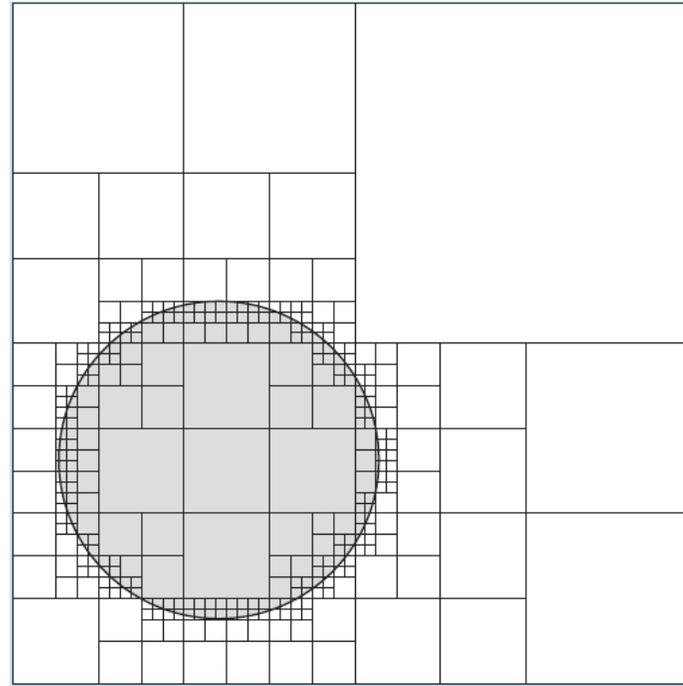
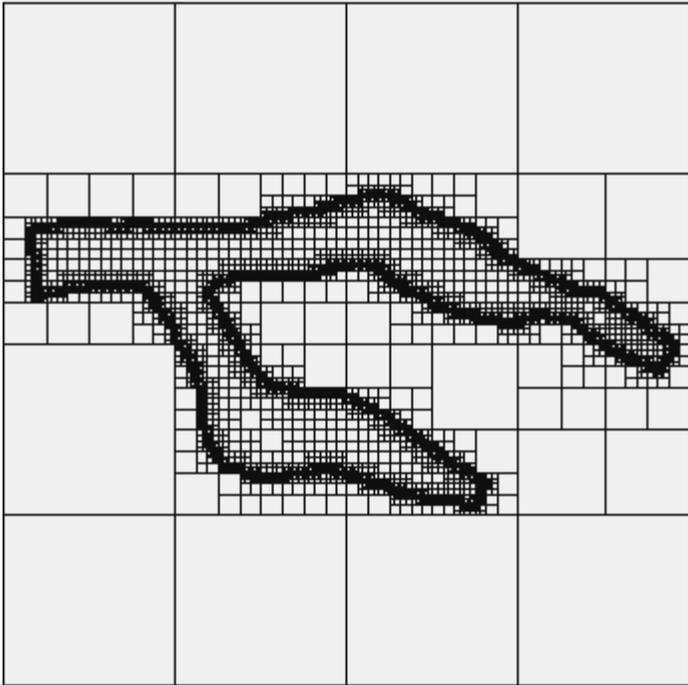


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Quadrees



Each node represents a square. If the node is fully empty or fully occupied it has no children.
If it is partially occupied it has four children. Subdivision stops after some minimal square size.

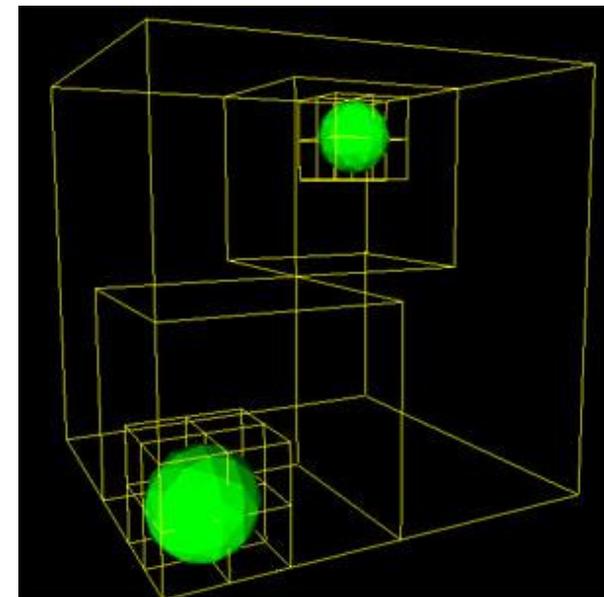
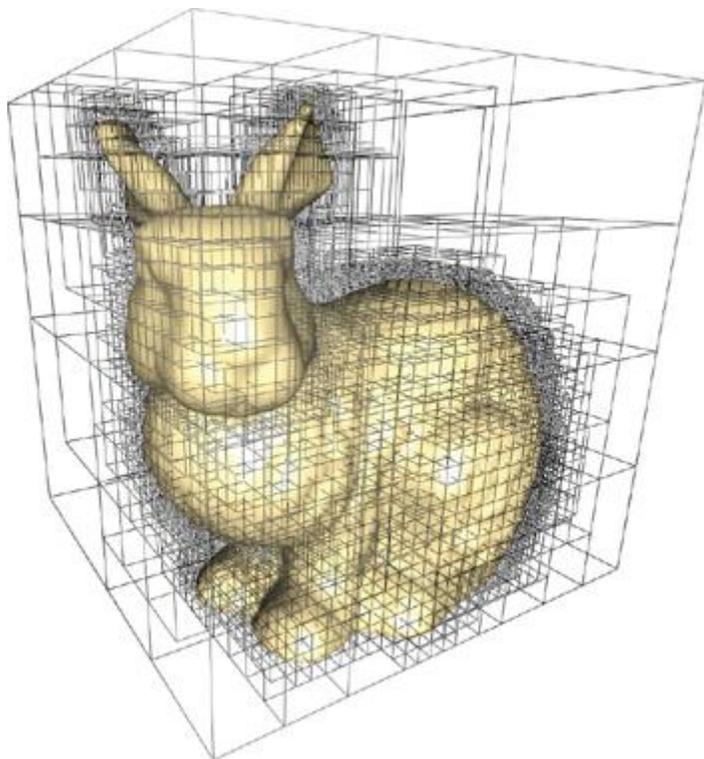


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Octrees



Each node represents a cube. If the node is fully empty or fully occupied it has no children.

If it is partially occupied it has eight children. Subdivision stops after some minimal cube size.

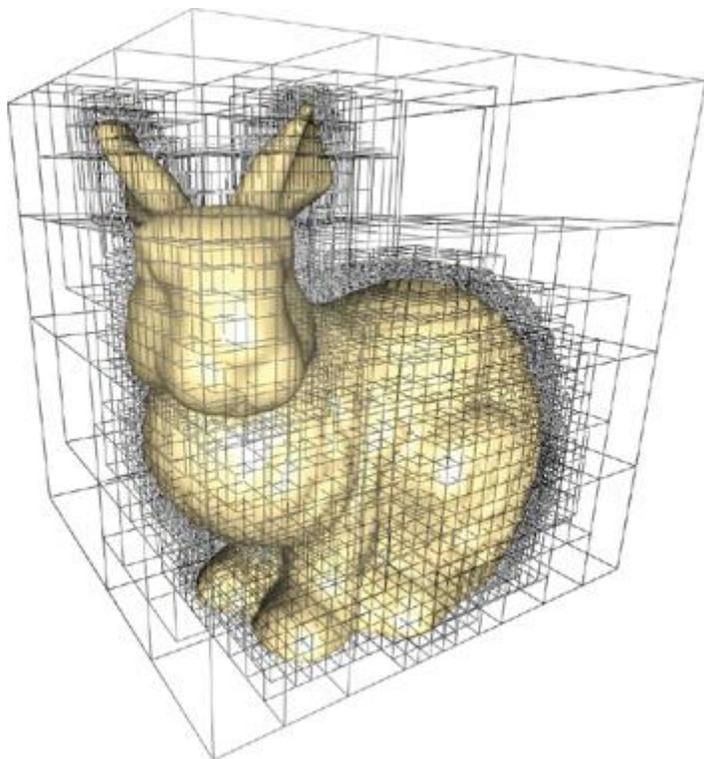


McGill

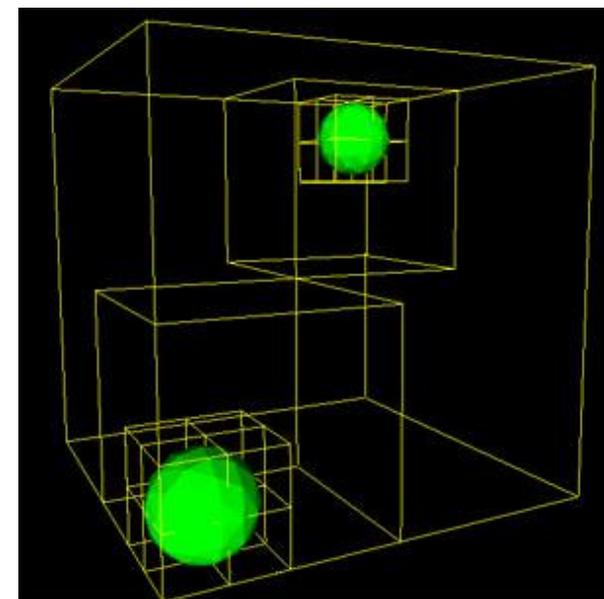
School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Octrees



Problem 1: quadtrees and octrees are not balanced trees. So, in the worst case an occupancy query could be $O(n)$ in the number of nodes.



Each node represents a cube. If the node is fully empty or fully occupied it has no children.
If it is partially occupied it has eight children. Subdivision stops after some minimal cube size.

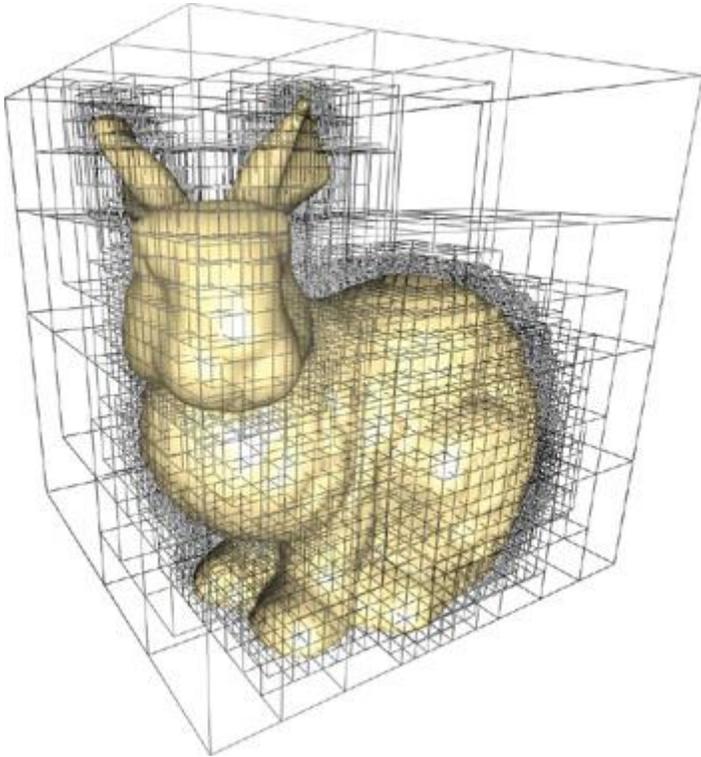


McGill

School of Computer Science
Centre for Intelligent Machines

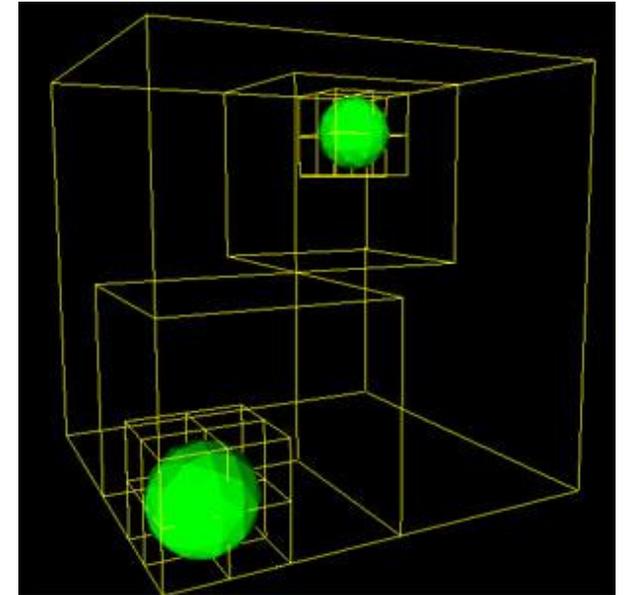
Intelligent Robotics

Octrees



Problem 1: quadtrees and octrees are not balanced trees. So, in the worst case an occupancy query could be $O(n)$ in the number of nodes.

Problem 2: quadtrees and octrees are sensitive to small changes in the location of obstacles.



Each node represents a cube. If the node is fully empty or fully occupied it has no children.

If it is partially occupied it has eight children. Subdivision stops after some minimal cube size.



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Octree Example: Octomap

Open source
as a ROS package



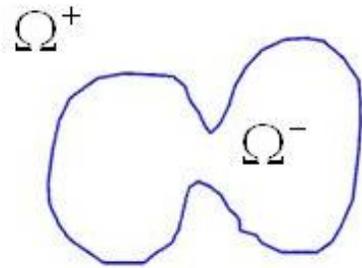
McGill

School of Computer Science
Centre for Intelligent Machines

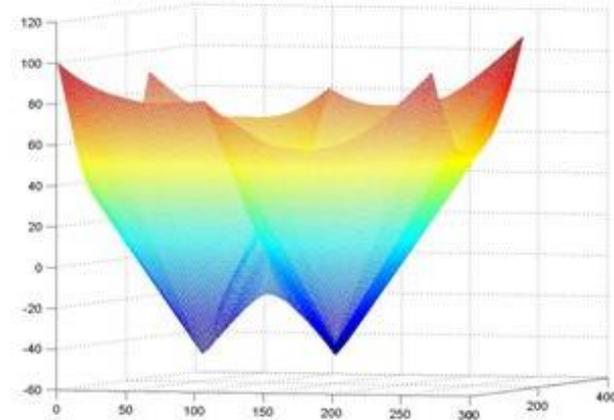
Intelligent Robotics

Implicit Surface Definitions: Signed Distance Function

Contour C



Signed distance function



This distance function is defined over any point in 3D space.

A signed distance function is defined by :

$$\phi(x) = \begin{cases} \text{dist}(x, C) & \text{if } x \text{ is outside } C \\ 0 & x \in C \\ -\text{dist}(x, C) & \text{if } x \text{ is inside } C \end{cases}$$



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

SDF Example

Real-Time Camera Tracking and 3D Reconstruction Using Signed Distance Functions

Erik Bylow, Jürgen Sturm, Christian Kerl,
Fredrik Kahl, Daniel Cremers

**Robotics: Science and Systems (RSS)
June 2013**



Computer Vision Group
Department of Computer Science
Technical University of Munich

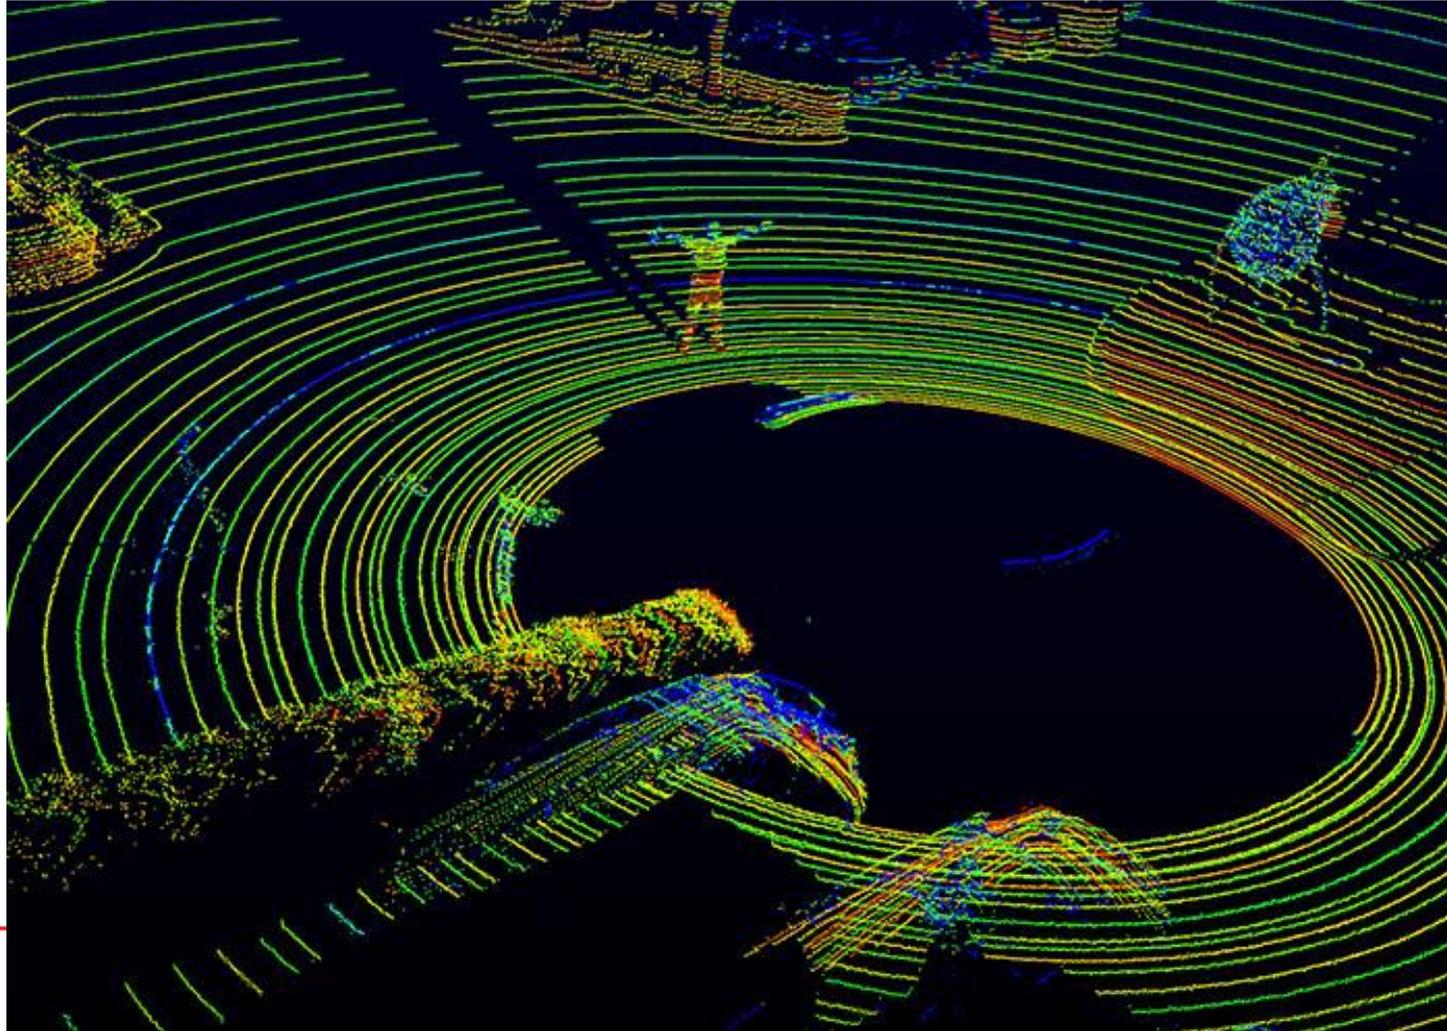


McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Pointclouds



McGill

Centre for Intelligent Machines

Intelligent Robotics

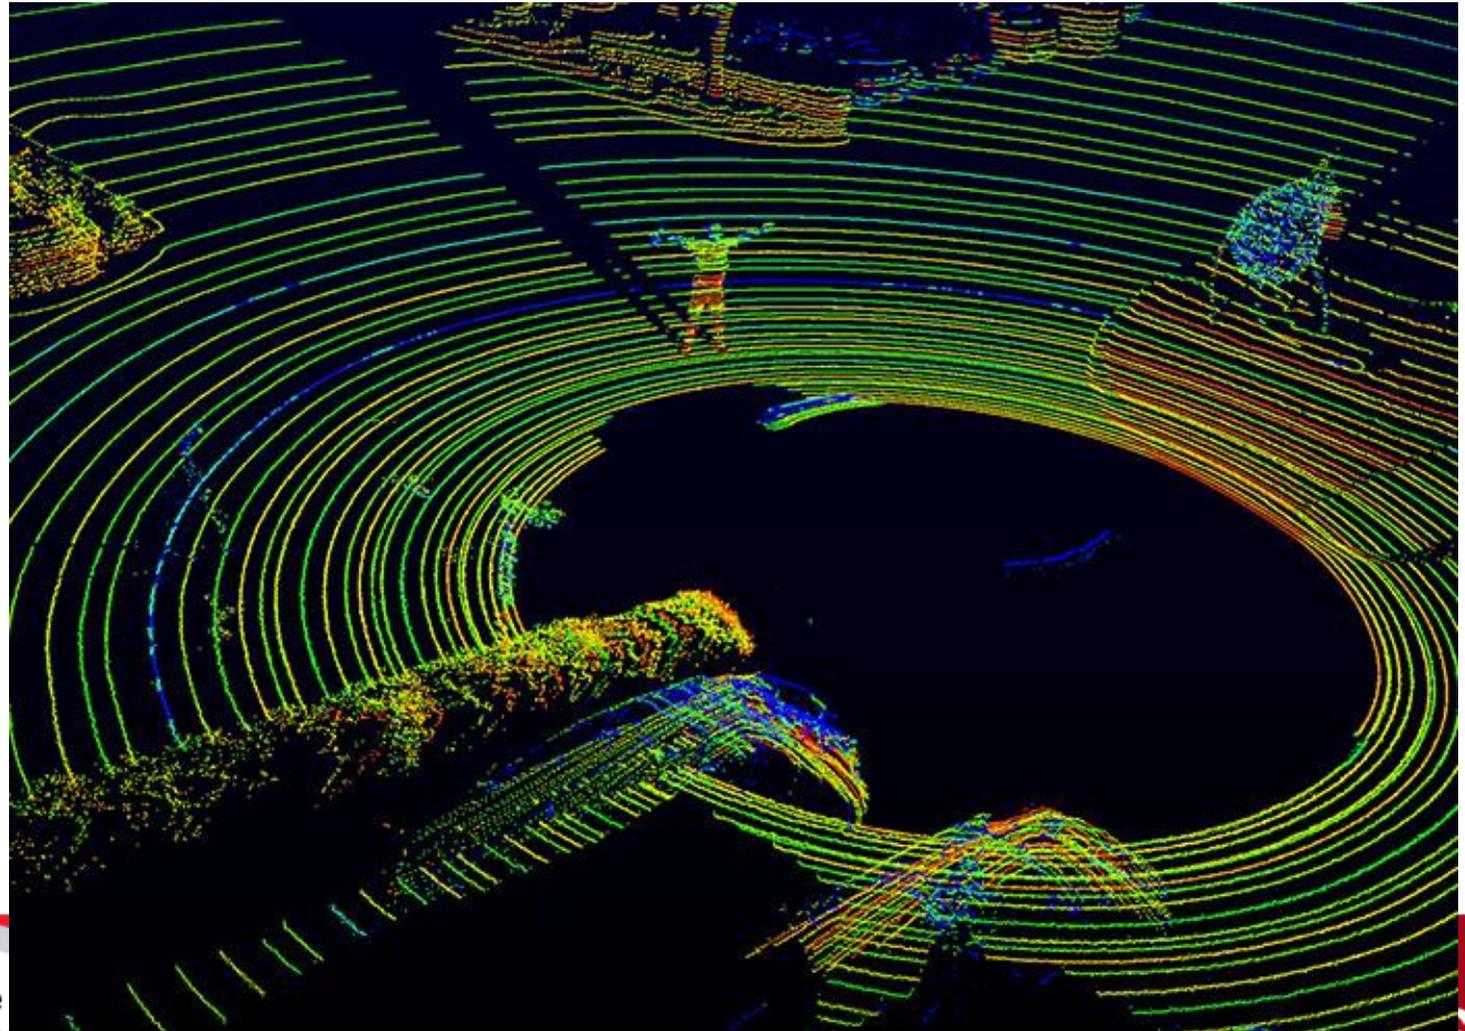
Pointclouds

Advantages:

- can make local changes to the map without affecting the pointcloud globally
- can align pointclouds
- nearest neighbor queries are easy with kd-trees or locality-sensitive hashing

Disadvantages:

- need to segment objects in the map
- raytracing is approximate and nontrivial



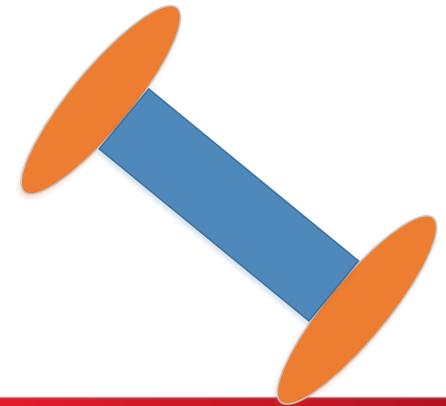
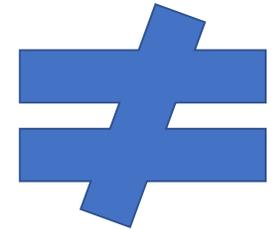
McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

The Utility of Models

- “All models are wrong, but some are useful” – George Box (statistician)
- Model: a function that describes a physical phenomenon or a system, i.e. how a set of input variables cause a set of output variables.
- Models are useful if they can predict reality up to some degree.
- Mismatch between model prediction and reality = **error / noise**
- We need to use these models carefully, and mix them with data that can improve performance!



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

So, where does the data come from?

- A robot's sensors give it feedback on its internal state (proprioception), and allow it to observe the external world (exteroception)
- Like models, no sensor is perfect, but much work goes into making sensor data optimally useful for robotics applications:
 - What a pain... thank goodness we'll see how probability helps here next week!



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Types of sensors

- General classification:
 - contact vs. non-contact
 - active vs. passive
 - sampling rate: fast vs. slow
 - local vs. non-local
- General examples:
 - vision
 - laser
 - radar
 - sonar
 - compass, gyroscope, accelerometer
 - touch (tactile)
 - infrared



Sensors

- Devices that can sense and measure physical properties of the environment.
- Key phenomenon is **transduction** (conversion of energy from one form to another). E.g.:
 - Imaging sensors: light to pixel voltages
 - Depth sensors: mechanical pressure to voltage
- Measurements are **noisy**, and difficult to interpret



Sensors: general characteristics

- Sensitivity: $(\text{change of output}) \div (\text{change of input})$
- Linearity: constancy of $(\text{output} \div \text{input})$
- Measurement range: $[\text{min}, \text{max}]$ or $\{\text{min}, \text{max}\}$
- Response time: time required for input change to cause output change
- Accuracy: difference between measurement and actual
- Repeatability/Drift: difference between repeated measures
- Resolution: smallest observable increment
- Bandwidth: required rate of data transfer
- SNR: signal-to-noise ratio

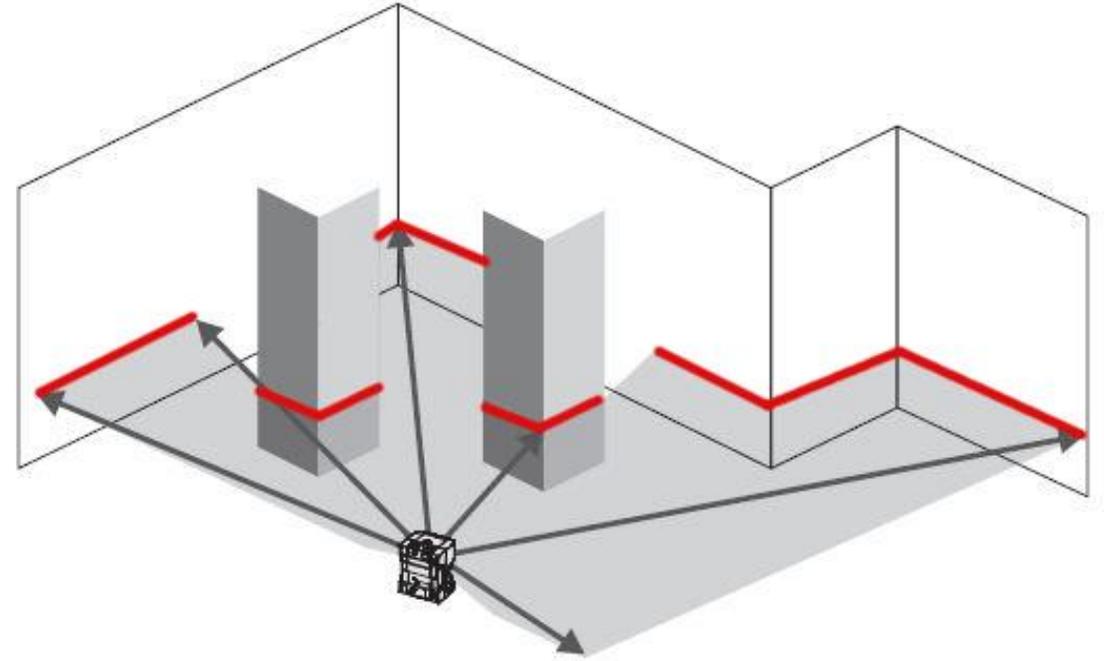


2D LIDAR (Light detection and ranging)

Produces a scan of 2D points and intensities

- (x,y) in the laser's frame of reference
- Intensity is related to the material of the object that reflects the light

Certain surfaces are problematic for LIDAR: e.g. glass



McGill

School of Computer Science
Centre for Intelligent Machines

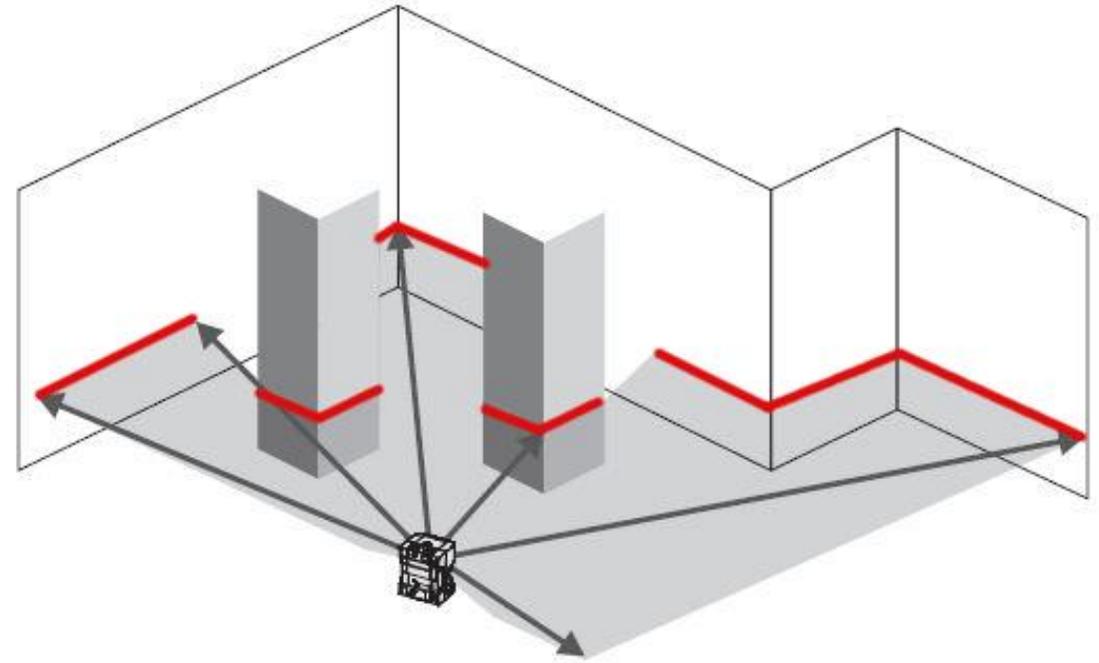
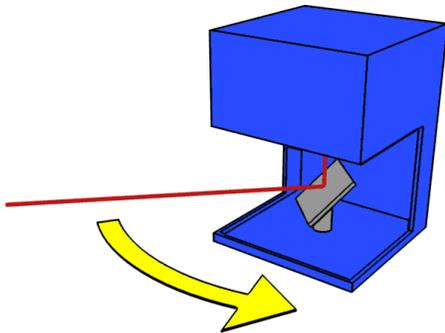
Intelligent Robotics

2D LIDAR (Light detection and ranging)

Produces a scan of 2D points and intensities

- (x,y) in the laser's frame of reference
- Intensity is related to the material of the object that reflects the light

Certain surfaces are problematic for LIDAR: e.g. glass



Usually around 1024 points in a single scan.



McGill

School of Computer Science
Centre for Intelligent Machines

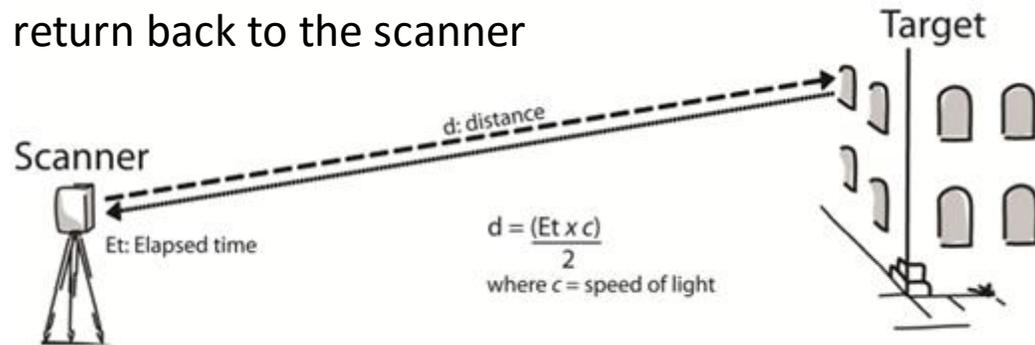
Intelligent Robotics

3D LIDAR (Light detection and ranging)

Produces a pointcloud of 3D points and intensities

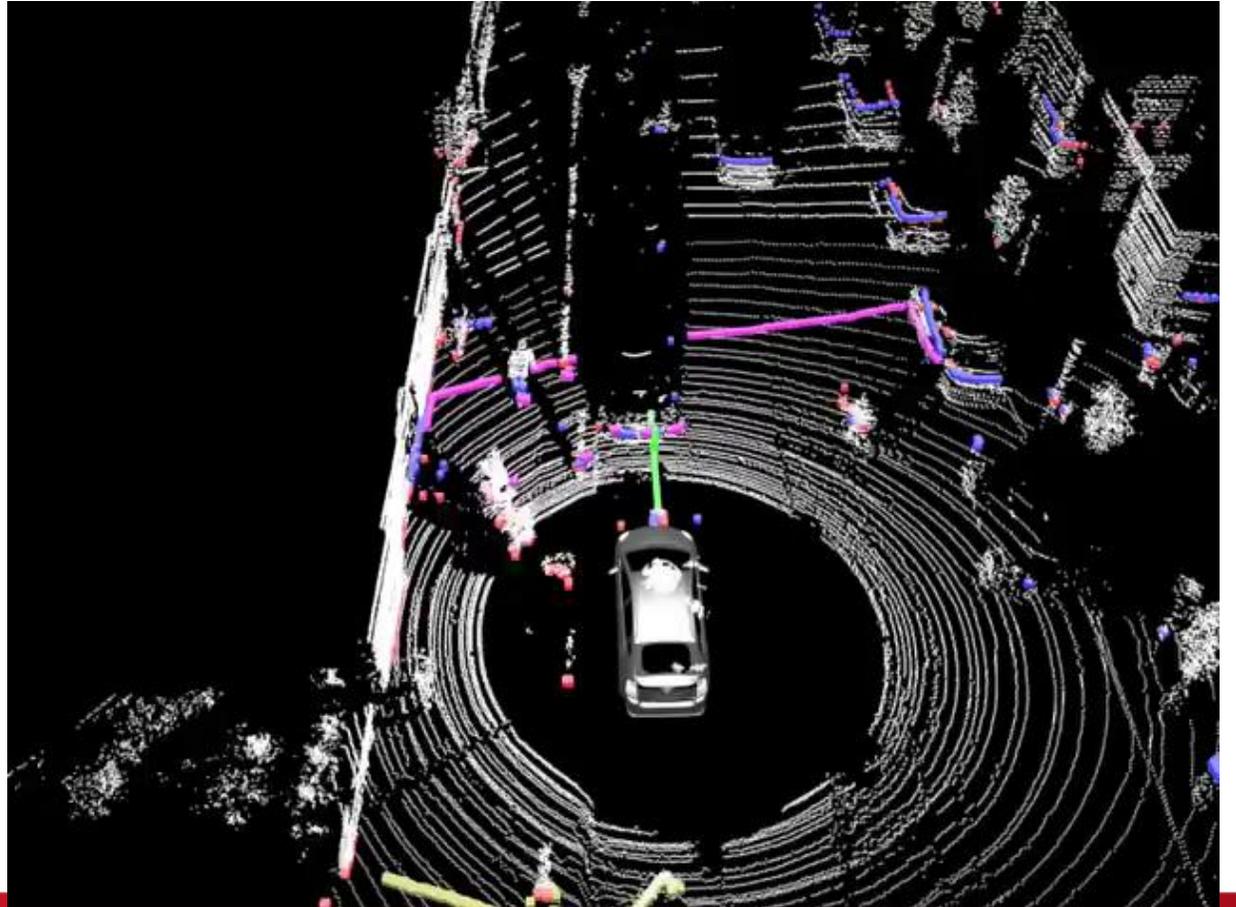
- (x,y,z) in the laser's frame of reference
- Intensity is related to the material of the object that reflects the light

Works based on time-of-flight for each beam to return back to the scanner



Not very robust to adverse weather conditions:
rain, snow, smoke, fog etc.

Usually around 1million points in a single pointcloud



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Inertial Sensors

- Gyroscopes, Accelerometers, Magnetometers
- Inertial Measurement Unit (IMU)
- Perhaps the most important sensor for 3D navigation, along with the GPS
- Without IMUs, plane autopilots would be much harder, if not impossible, to build



Beyond the visible spectrum: RGBD cameras



Main ideas:

- Active sensing
- Projector emits infrared light in the scene
- Infrared sensor reads the infrared light
- Deformation of the expected pattern allows computation of the depth



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Beyond the visible spectrum: RGBD cameras

Drawbacks:

- Does not work outdoors, sunlight saturates its measurements
- Maximum range is [0.5, 8] meters

Advantages:

- Real-time depth estimation at 30Hz
- Cheap



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Beyond the visible spectrum: RGBD cameras

Enabled a wave of research, applications,
and video games, based on real-time
skeleton tracking



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Beyond the visible spectrum: RGBD cameras

Despite their drawbacks RGBD sensors have been extensively used in robotics.



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

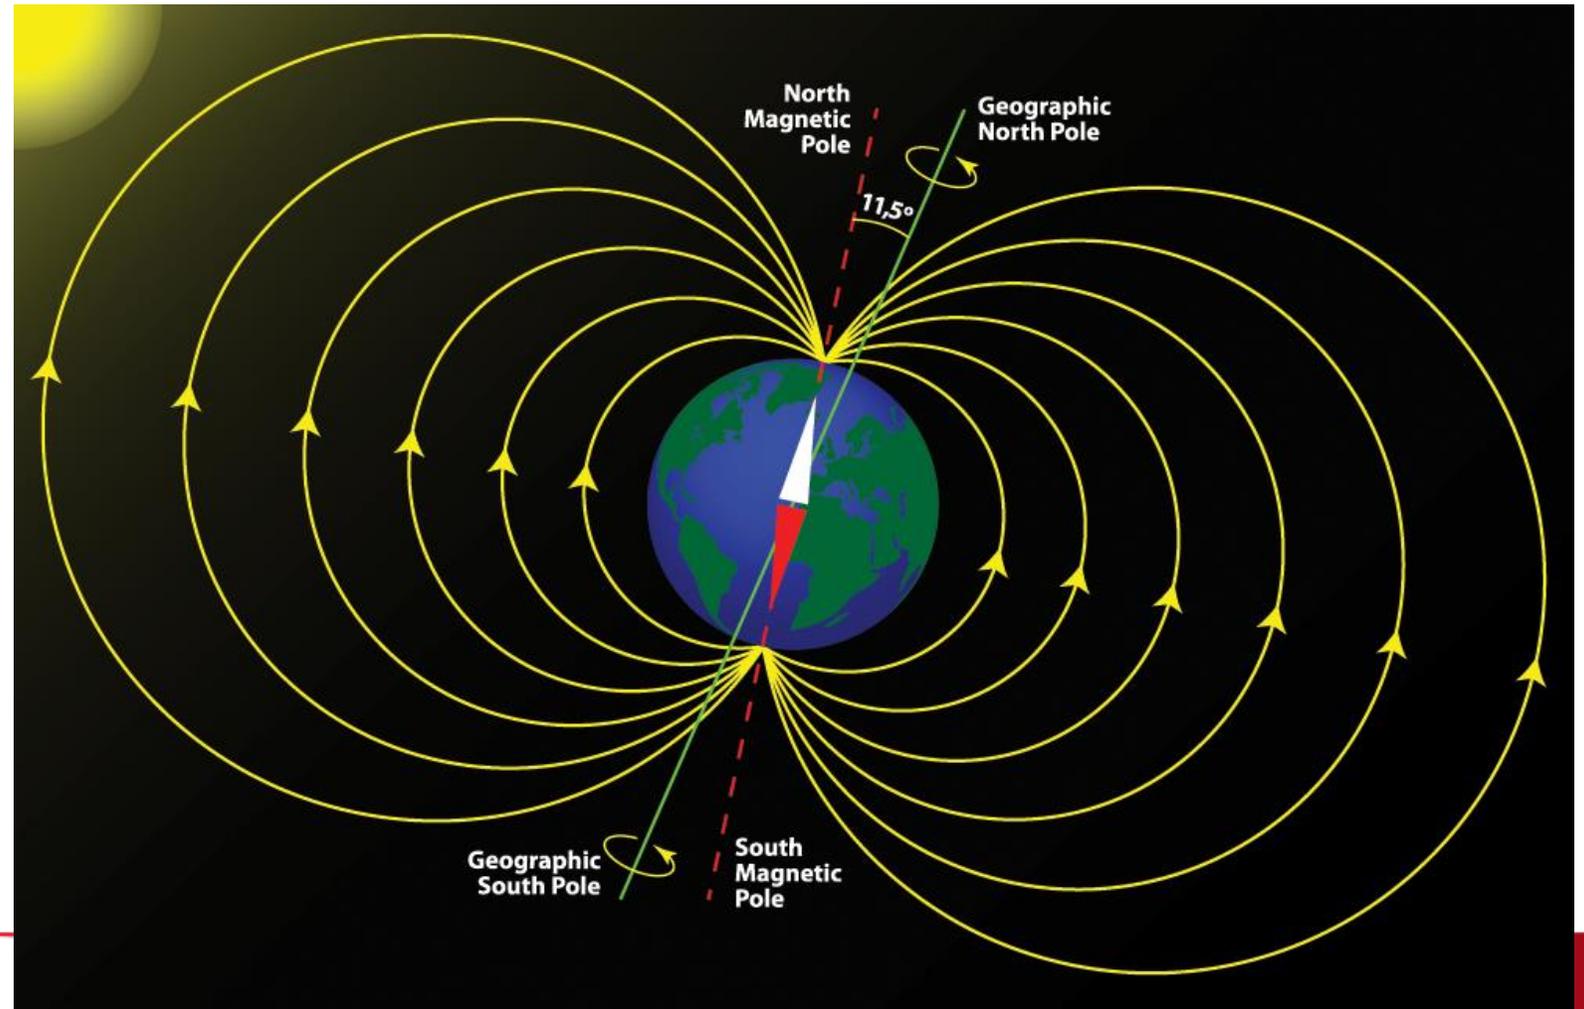
Magnetometers

Drawbacks:

- Needs careful calibration
- Needs to be placed away from moving metal parts, motors

Advantages:

- Can be used as a compass for absolute heading



McGill

School of Computer Science
Centre for Intelligent Machines

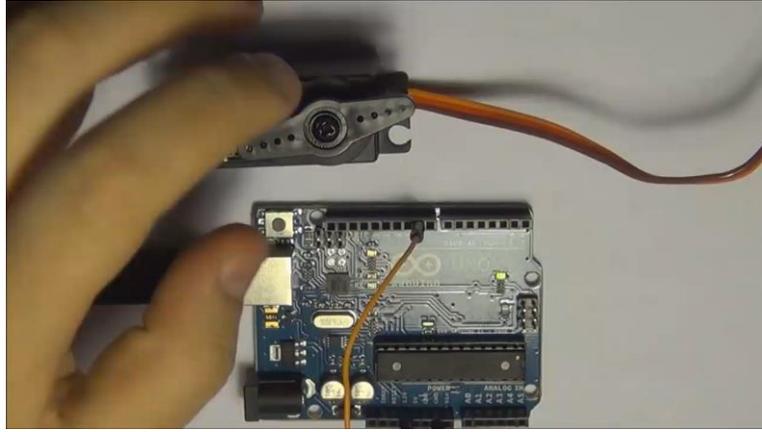
Intelligent Robotics

Actuators are also a form of sensors...



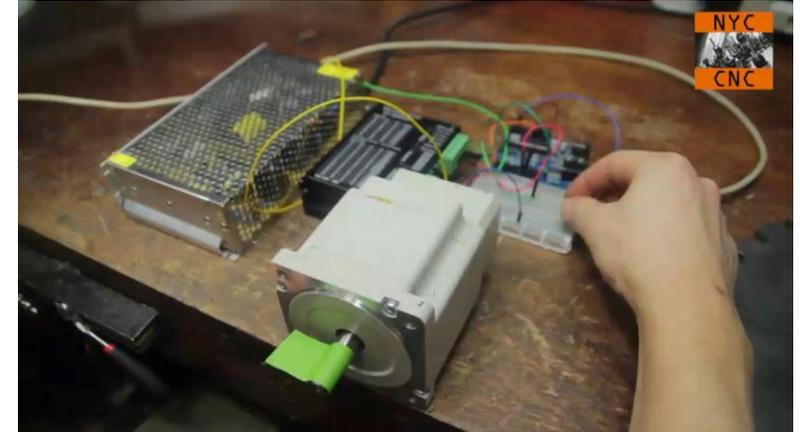
DC (direct current) motor

They turn continuously at high RPM (revolutions per minute) when voltage is applied. Used in quadrotors and planes, model cars etc.



Servo motor

Usually includes: DC motor, gears, control circuit, position feedback
Precise control without free rotation (e.g. robot arms, boat rudders)
Limited turning range: 180 degrees



Stepper motor

Positioning feedback and no positioning errors.
Rotates by a predefined step angle.
Requires external control circuit.
Precise control without free rotation.
Constant holding torque without powering the motor (good for robot arms or weight-carrying systems).



McGill

School of Computer Science
Centre for Intelligent Machines

Intelligent Robotics

Cliffhanger

- We will use models of motion and perception to estimate a robot's position with uncertainty:
 - $p(x_t | x_{t-1}, u_{t-1})$ is a generative model of a robot's one-step motion
 - $p(z_t | x_t, m)$ is a generative model of a robot's current sensing
- We can use these models in countless places throughout robotics, but the simplest and most well-studied is localization: "Where am I now?"
 - Infer $p(x_t | z_{1...t}, m)$
- We'll start on how to do this next time.

