CS 189: Autonomous Robot Systems

Spring 2018, Fridays 1-4pm, Pierce 301

Todo

- Notes:
  - This lecture took a full 1.5 hour – just the full trash collector took 45 min. But was good, history part was fun and went well.
  - Maybe for future, just take a break 5 minutes between two segments (note 2014, 5 min break worked well)
  - 2018: I should add a code examples [state machine]
- Brief but gripping history of robotics?
  - In the future hand out MMs chapter?
  - RM’s chapter not bad, but less people oriented? Maybe also hand out that too?
- 2018 Notes: BEAM instead (better match for now?)

Agenda

- Lecture: Autonomy III: Programming Complex Behaviors
  - Software Architectures for Robotics
  - Brief History of Robotics
- Demo Time:
  - Pset 3 (a) tests
  - Lab 3 (More Interactions: Text2Speech and ARTags)
- Upcoming:
  - Pset3 (b): Follower

- References: This lecture is based in part on
  "Introduction to AI Robotics", chapters 2-5, Robin Murphy, 2000 [trashcollector, architectures]
  "A Robotics Primer", chapter 2, Maja Mataric, 2007 [history]

What Does it Mean to be Autonomous?



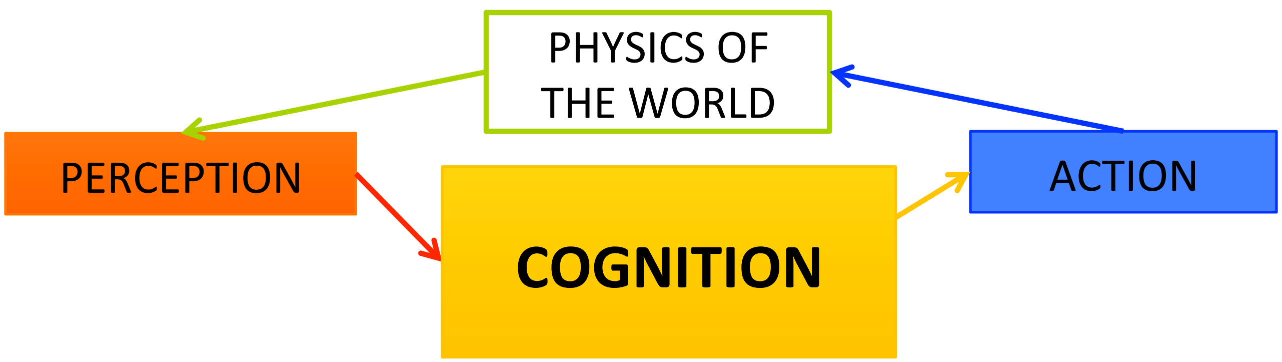

Basics of Autonomy

ACTION

- Action (Actuators)
  - Locomotion: Wheels (Differential Drives, Kinematics)

PERCEPTION

- Perception (Sensors)
  - Proprioception and Exteroception (Bump, Depth)

COGNITION

- Cognition (Control)
  - Reactive Behaviors + PID Control
  - **Software Architectures: Simple to Complex Tasks**

Example: Trash Collector

Murphy 2000

Example: Trash Collector

- Lets solve a bigger problem
  - AAAI Competition: Pick Up the Trash
- Problem Specs
  - Coca Cola cans, collect them and deposit them in the trashcan (changes year-to-year)
- What are some primitive behaviors we might need?

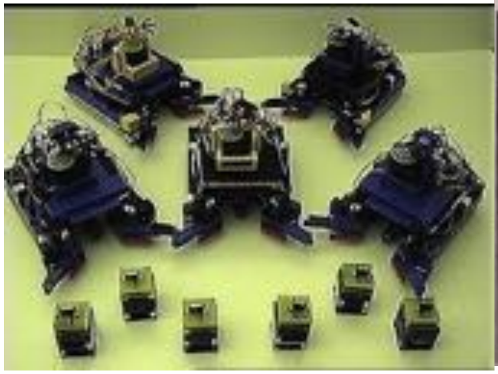

LOLA (AAAI 1995 Winner)
GaTech (AAAI 1994 Winner)
Example based on
Murphy 2000, section 5.5

Solution 1: Finite State Machine

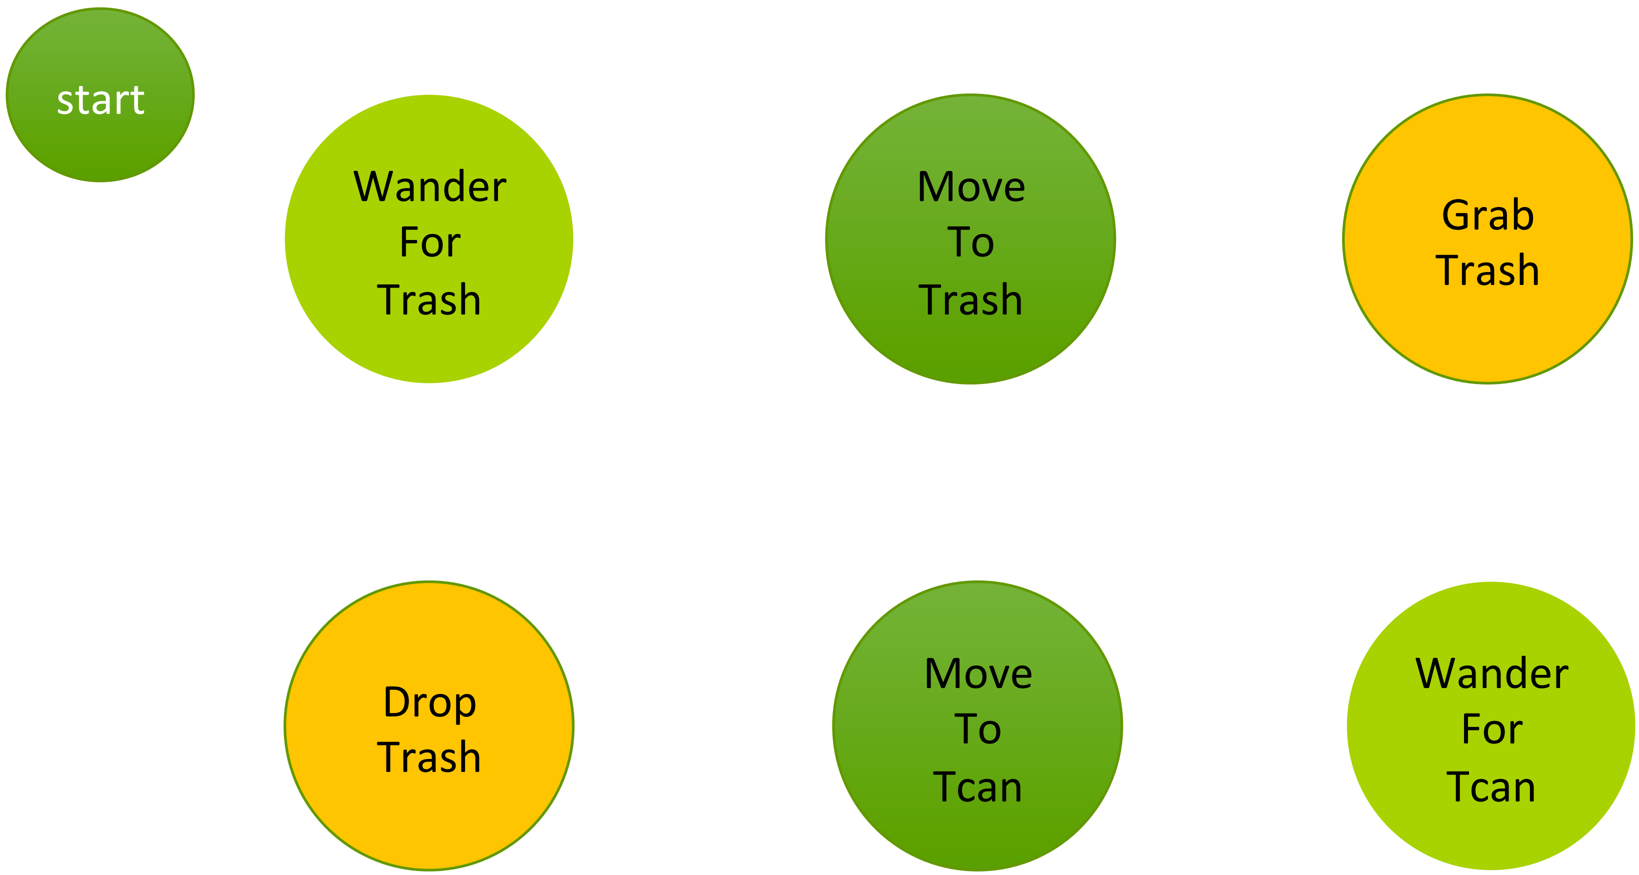

Solution 1: Finite State Machine

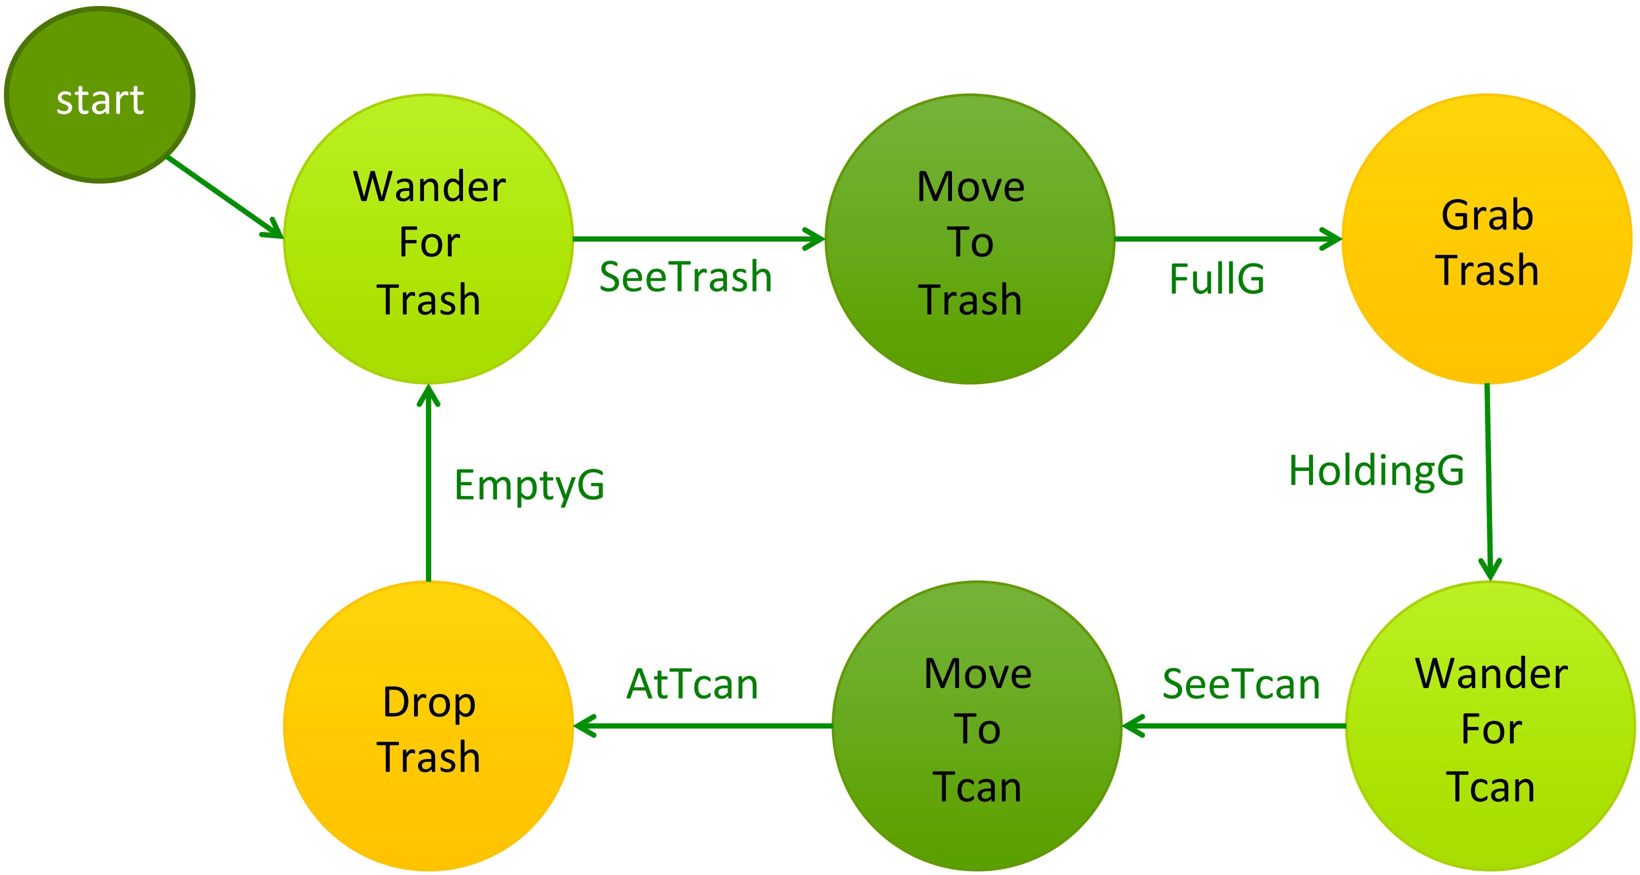

Solution 1: Finite State Machine

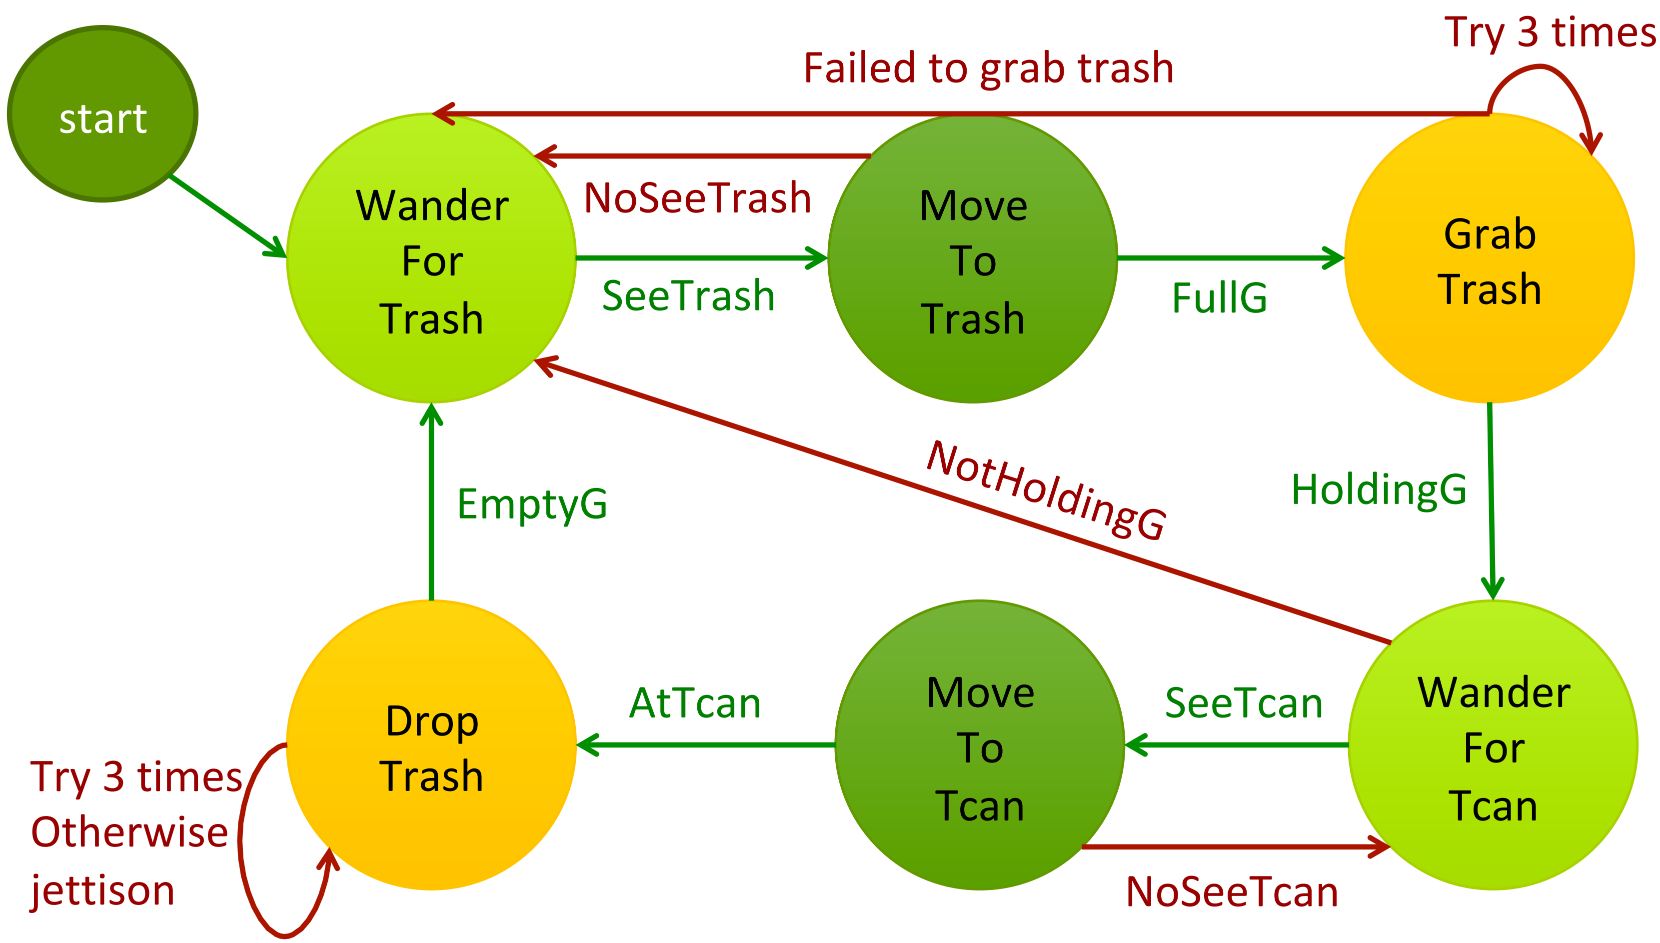

Solution 1: Finite State Machine

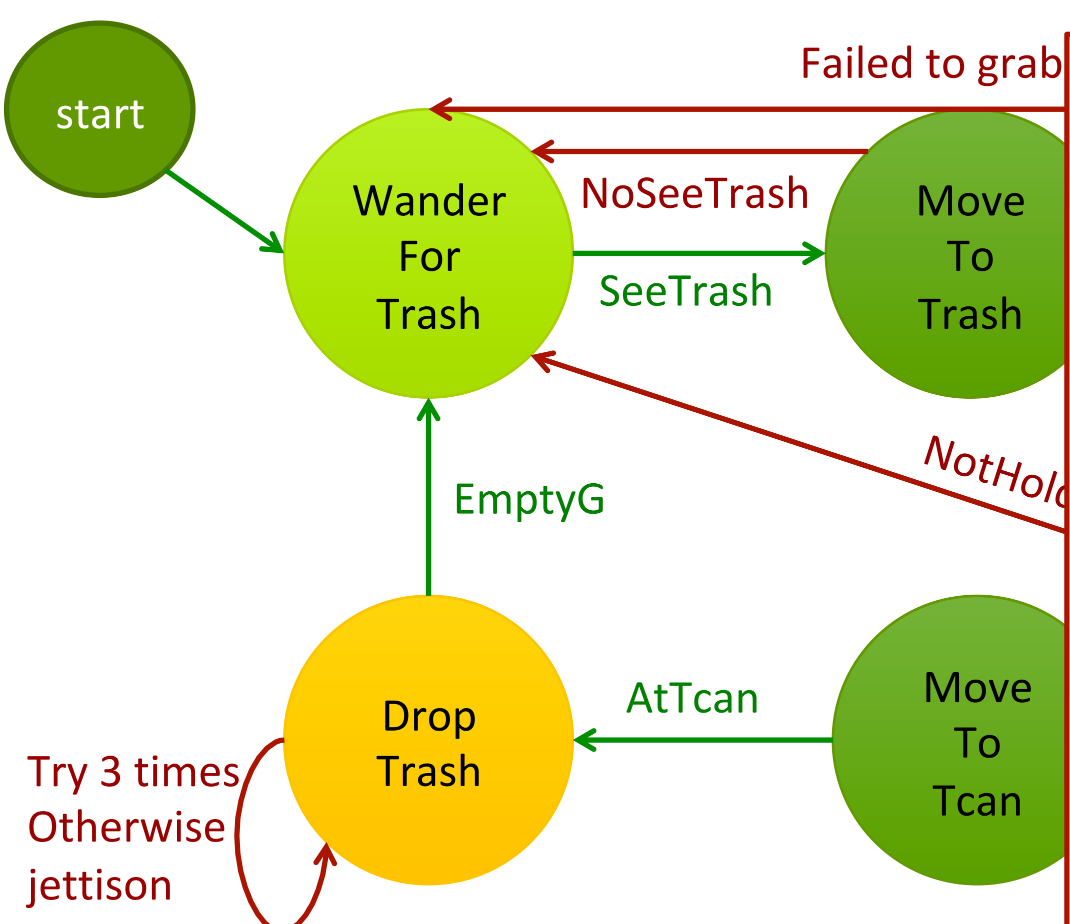

Things to Note:

1: Uses Internal State

2: Defines a Sequence

Issues:

1: Things may not happen in expected order (explicitly anticipate problems)

2: Behaviors like *Avoid Obstacles* and *Reset* underlie everything

Solution 2: Parallel Behaviors

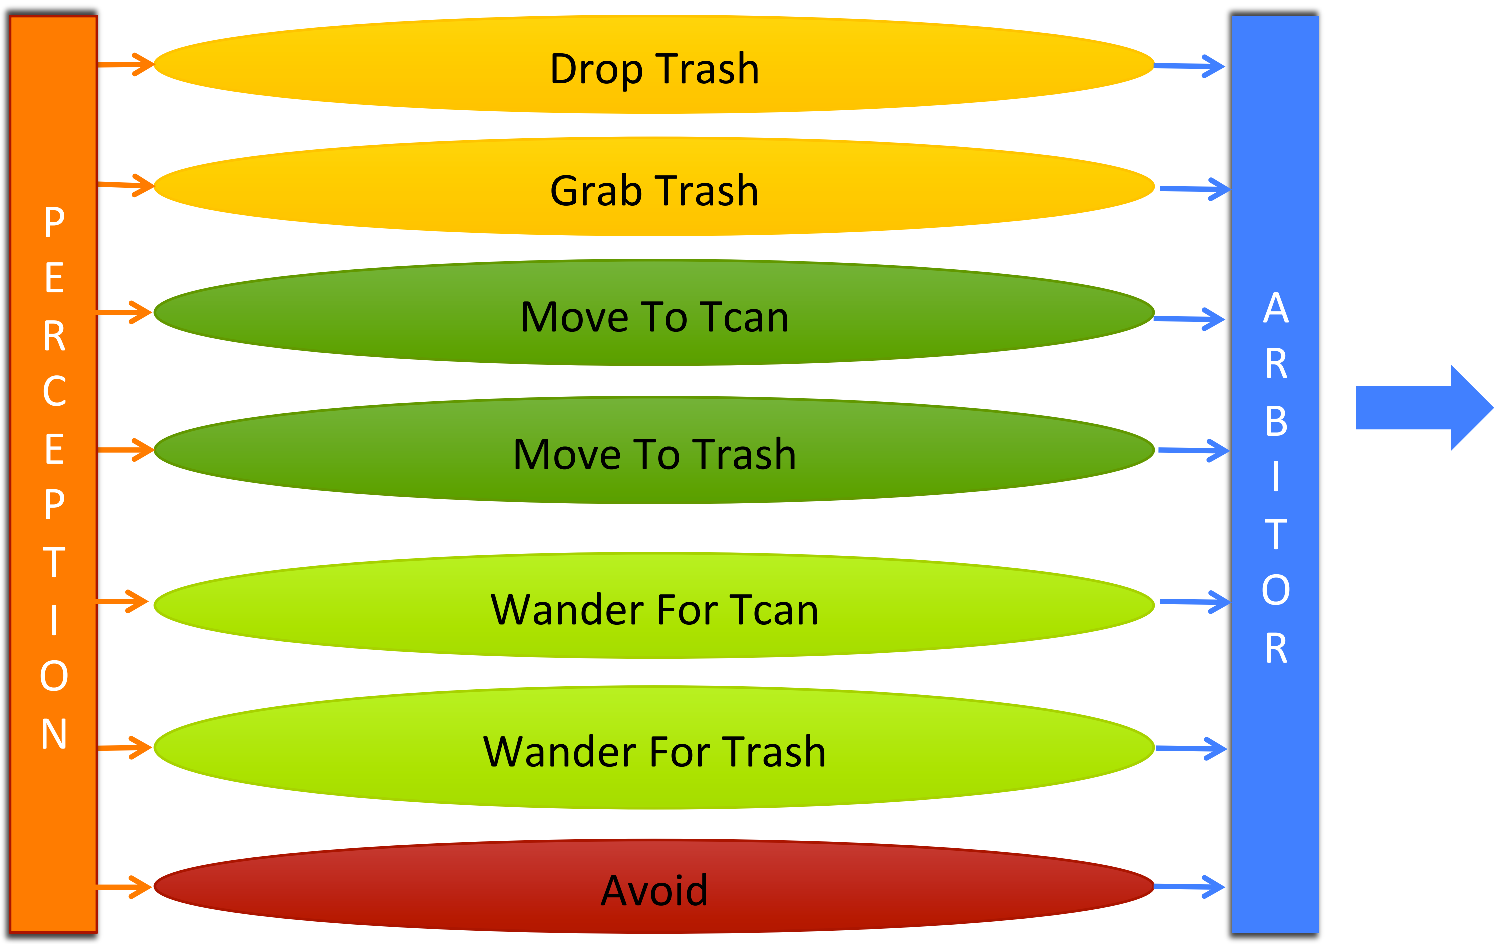

Solution 2: Parallel Behaviors

Solution 2: Parallel Behaviors

Things to Note:
Concurrency rules!

1. No “State”
2. Multiple behaviors may be active at once
3. Arbitrator can be of many kinds (priority, summation)

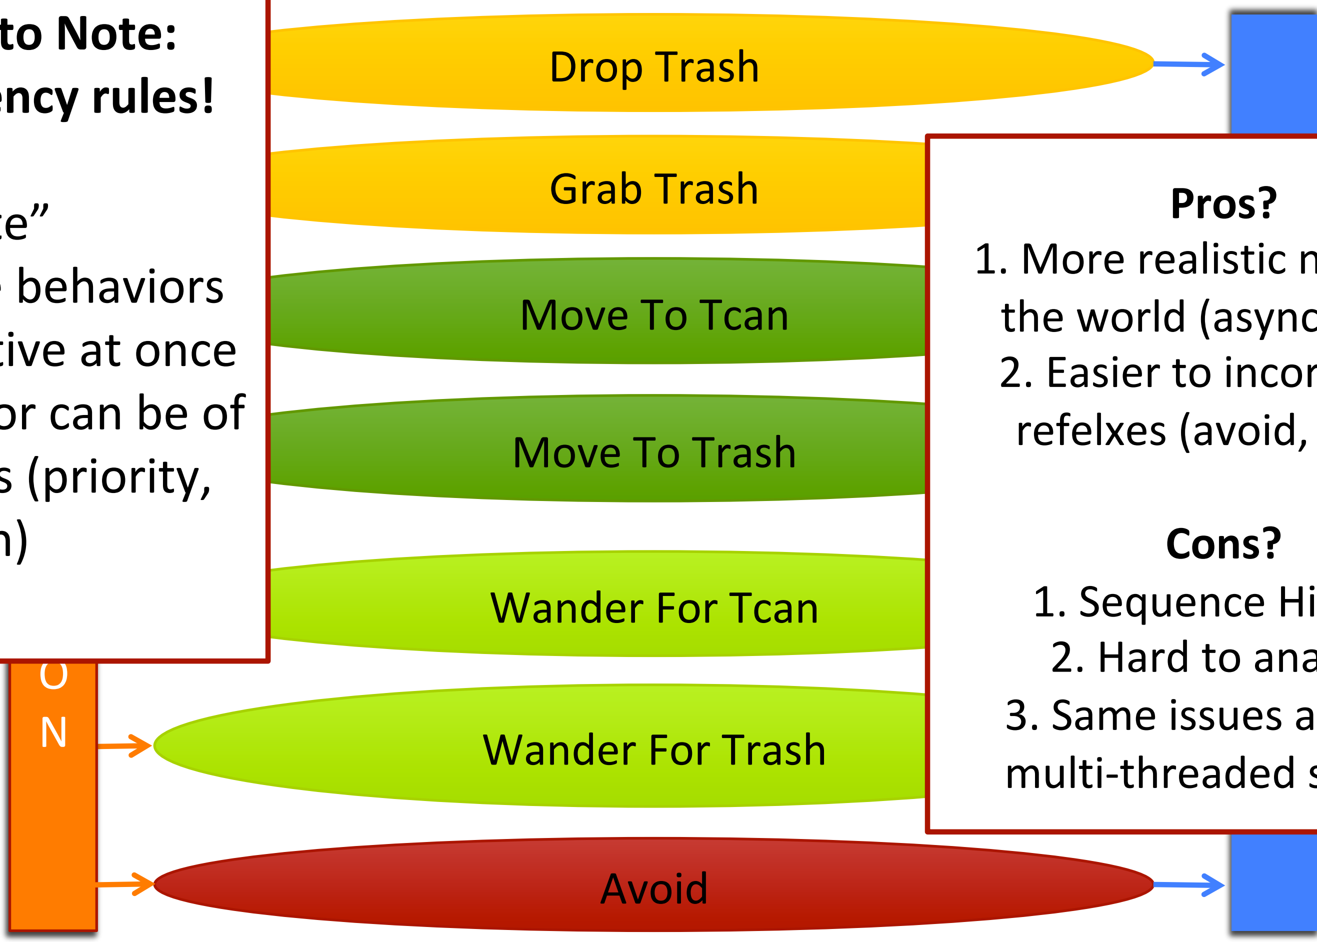

Pros?
1. More realistic notion of the world (asynchrony)
2. Easier to incorporate refelxes (avoid, reset)

Cons?
1. Sequence Hidden
2. Hard to analyze
3. Same issues as most multi-threaded system

Solution 3: Sense-Plan-Act

Sense	Plan	Act
Generate "world model"	Generate plan	Execute each step of plan

Sense and Construct Model of the World

- Camera: rotate in a circle to get 360 view
- Identify all visible cans, and trashcan, and esti
- `World Model: CANS ((x1,y1) (x2,y2) (x3 y3)`

Plan a Sequence: get each can and take them

- `Plan: Move-to-location(x1,y1), Grab Can, M location(x4,y4), DropCan, Move-to-locatio`

Act: each action corresponds to a "behavior" that

- *Example* `Move-to-location:` plan a path to a specific location

Pros?
Efficiency by being smart and looking ahead

Cons?
Assumption of perfect (error-free) perception and actuation

Roomba vs Neato

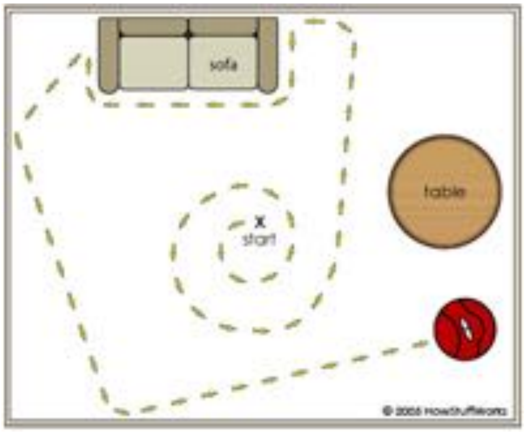

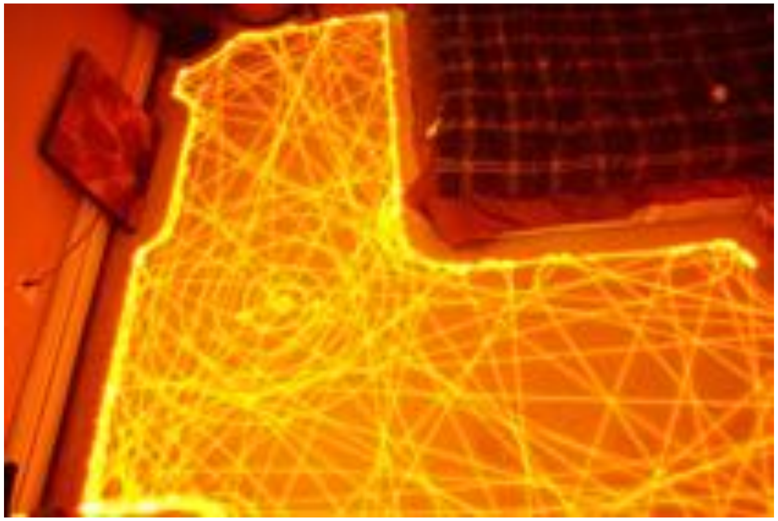

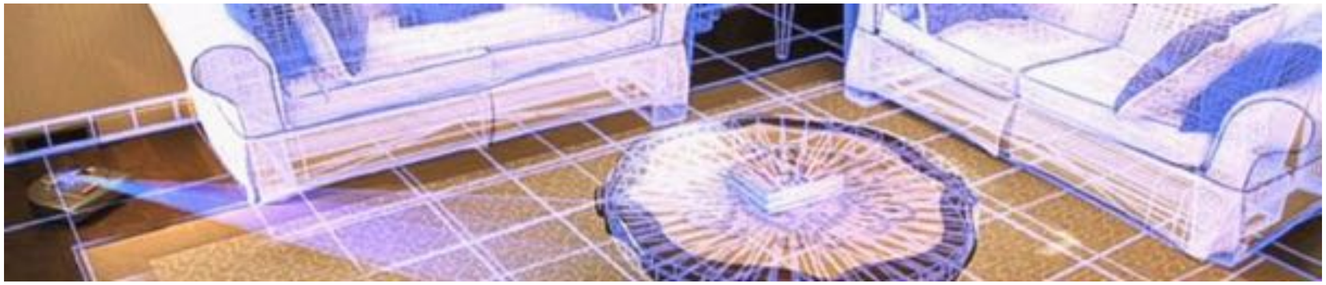

Example: Trash Collector

- Three Different Ways to Assemble Behaviors
  - Finite-State-Machine (Hybrid Architecture)
  - Parallel Behaviors (Reactive/Behavior-based/Subsumption)
  - Sense-Plan-Act (Deliberative Architecture)

History of AI-robotics is intimately tied with the question of how to program complex behaviors

Brief But Gripping History of Robotics
(based on Mataric 2007)

Brief History

Brief History

- The Word Robot
  - Czech playwright Karel Capek in his 1921 play "Rossum's Universal Robots"

    Rabota = menial work; Robotnik = serf

  - The word "Robotics" believed to be coined by Isaac Asimov

- Fields of Study Prior to Robotics and AI
  - **Control Theory** (Ancient!)
  - **Industrial Revolution** (Automation of human tasks)
  - **Cybernetics** (1940s, neuroscience and engineering)

Grey Walter's Tortoises (1940s)

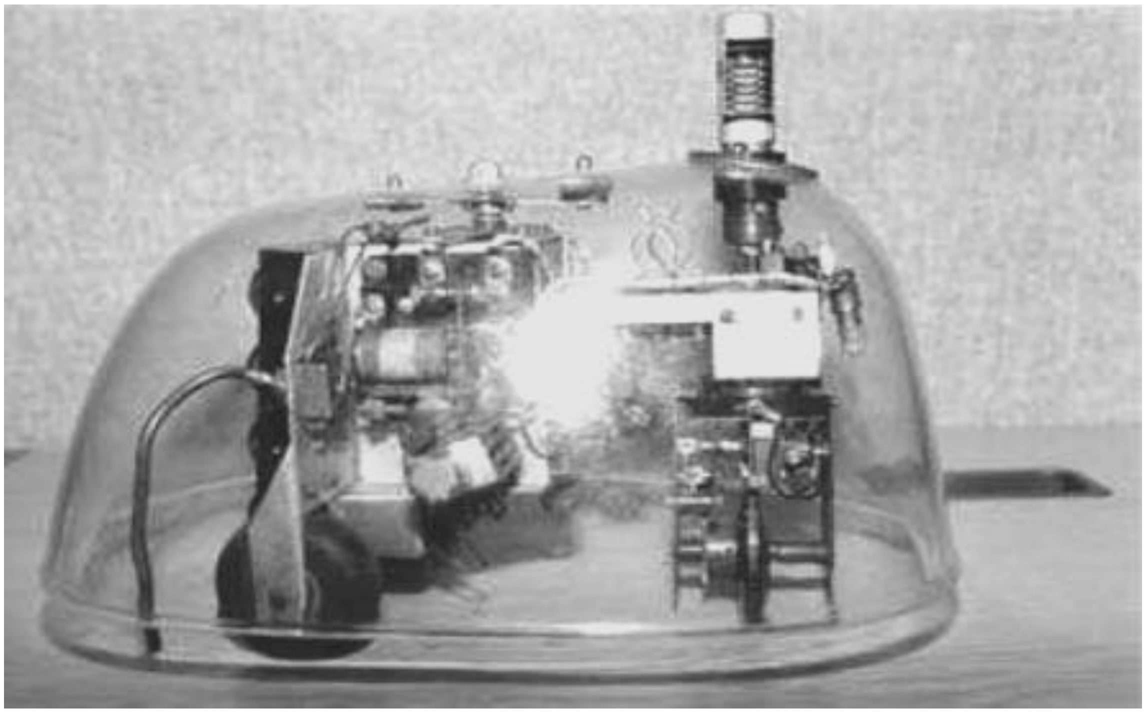

Elmer (Electro Mechanical Robots)
Elsie (Light Sensitive)
Machina Docilis (Tameable)

Grey Walter
(1910-1977; 1940s)
Neurophysiologist,
Cyberneticist

Simple robots based on reflexes

Before AI existed
Inspiration for Braitenberg

Birth of AI (1956) and Shakey (1960s)

- **Birth of Artificial Intelligence 1956**
  - Dartmouth Conference: Minsky, McCarthy, Newell, Simon etc)
  - *Theme: Intelligent Agents*
    - Internal Models of the world, Reason to solve problems, Symbolic representation of information
    - Intelligence == Playing chess, solving theorems
- **Shakey Robot (1960s-1970s)**
  - Stanford Research Institute (SRI)
  - Reason about its own actions

    (given a goal, construct *its own* plan)
  - Ambitious! included vision, natural language, A* search, STRIPs planning language
- **Stanford CART (1977, Hans Moravec)**
  - Birth of mobile robotics as we know it today

Shakey

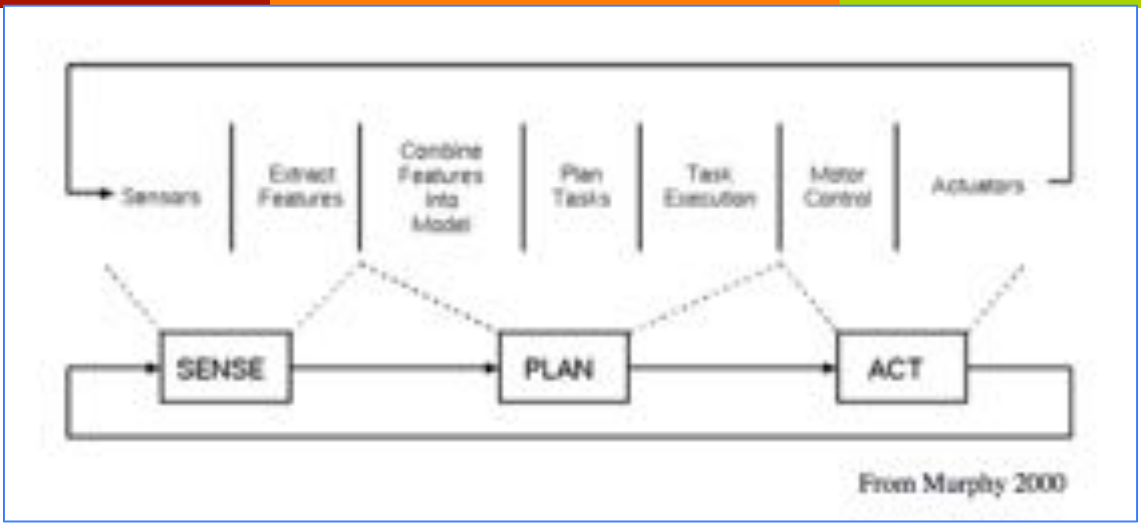

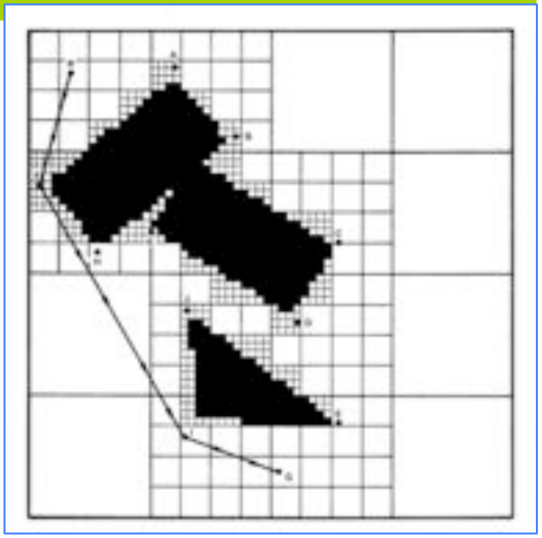

- World Model: Planar Grid Map with space, obstacles, and unknowns (Perception!!)
- Plans: generated using a STRIPs Planner ("theorem-proving")
  - Goal: Move three objects to a common location

    Exists (p,s), s.t. pos(OB1,s) and pos(OB2, s) and pos(OB3,s)

    (aka There exists a situation s and place p, such that OBI, OB2, and OB3 are all at place p in situation)
  - The task for the planner is to "prove" that this conjecture follows from axioms that describe present position of objects & effects of actions.

Shakey

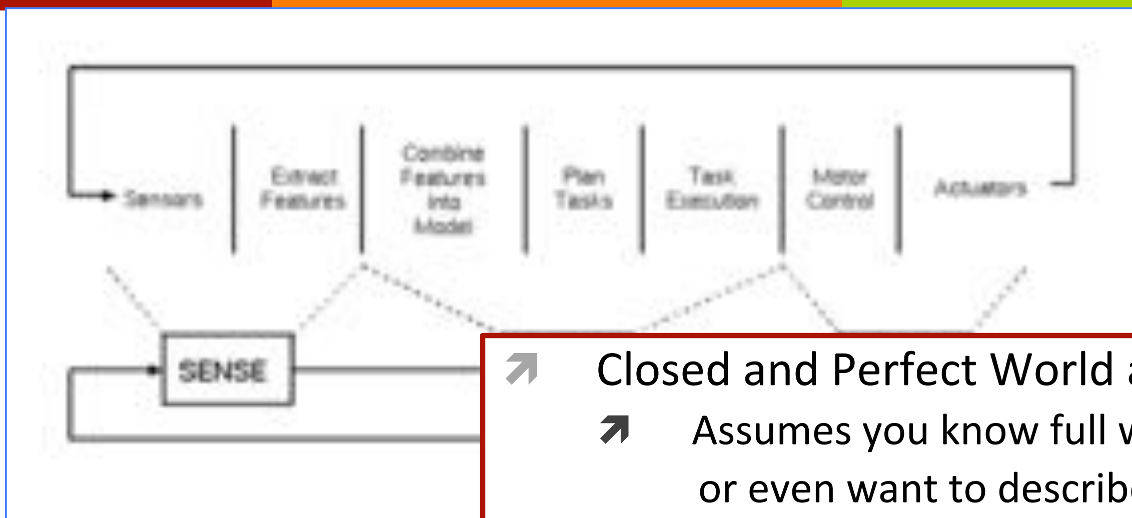

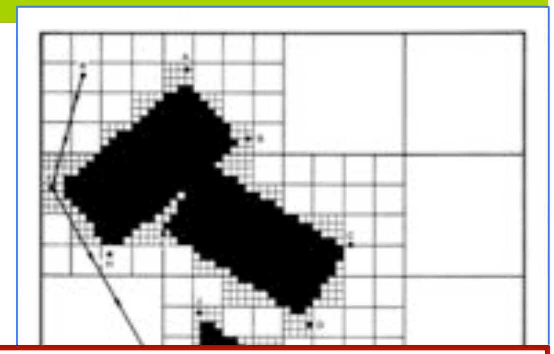

- World Model: Planar G
- Plans: generated using
  - Goal: Move three ob

    Exists (p,s), s.t. pos(OB1

    There exists a situation s
  - The task for the plan describe present po

- Closed and Perfect World assumption
  - Assumes you know full world (partially observable) or even want to describe it… (tedious)
  - Significant burden on perception
  - But also Planner relies on a model of how the world works [(put A on B) => AonB]
  - Deriving a reasonable model of world v. hard!
- Too much planning only to fail on the first step… (replanning all the time without getting anything done)

Even a simple world is complex

The initial environment of the Automaton was real, but contrived. It has been sufficiently simple to allow current visual capabilities to be useful to the Automaton, and sufficiently complex to indicate the weaknesses of current methods and to suggest areas of further research. Perhaps the most important result of our vision-research effort on the Automaton project is an appreciation of the potential complexity of the problem of vision when the real world is the subject matter, and a strong notion that the first step we have taken towards a general capability is very small indeed.

Nilsson, 1975, SRI Technical report

Moravec's Paradox

it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility

An Alternative View (1980s)

The Classical View
(1960s-70s)

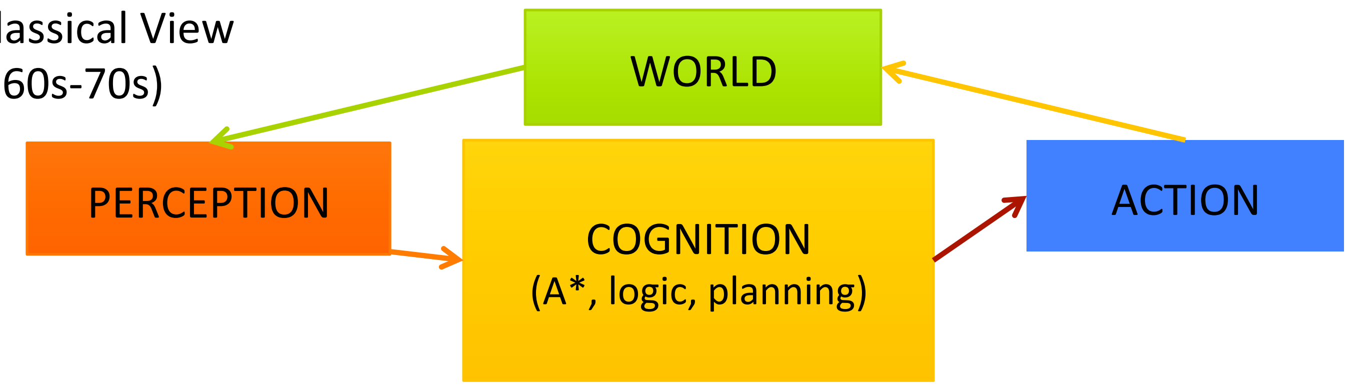

A Radical Idea
(Brooks, 1983)

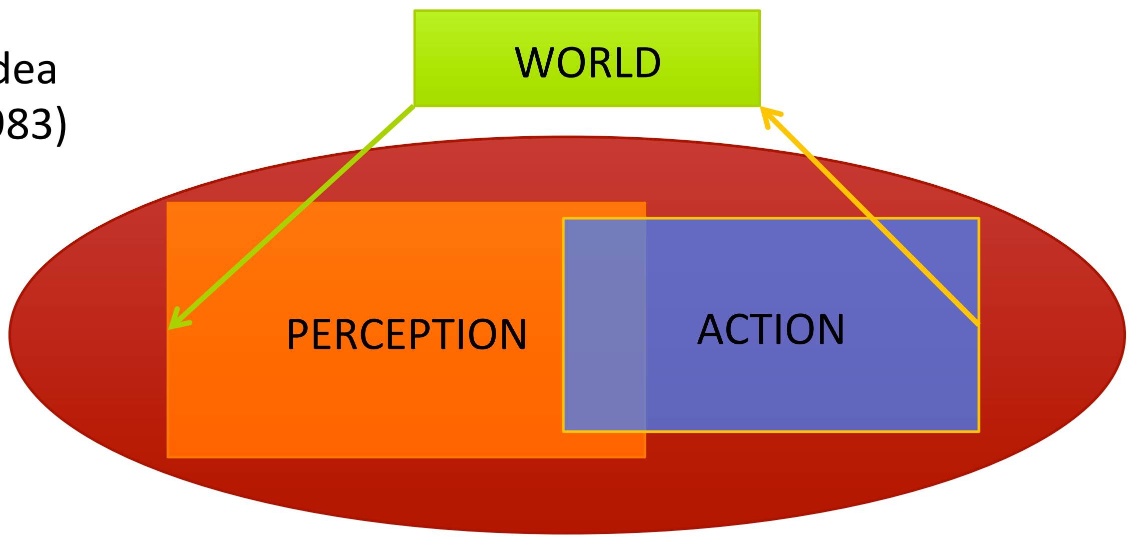

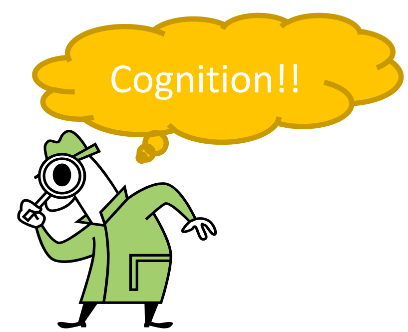

Related Influences in 1980s

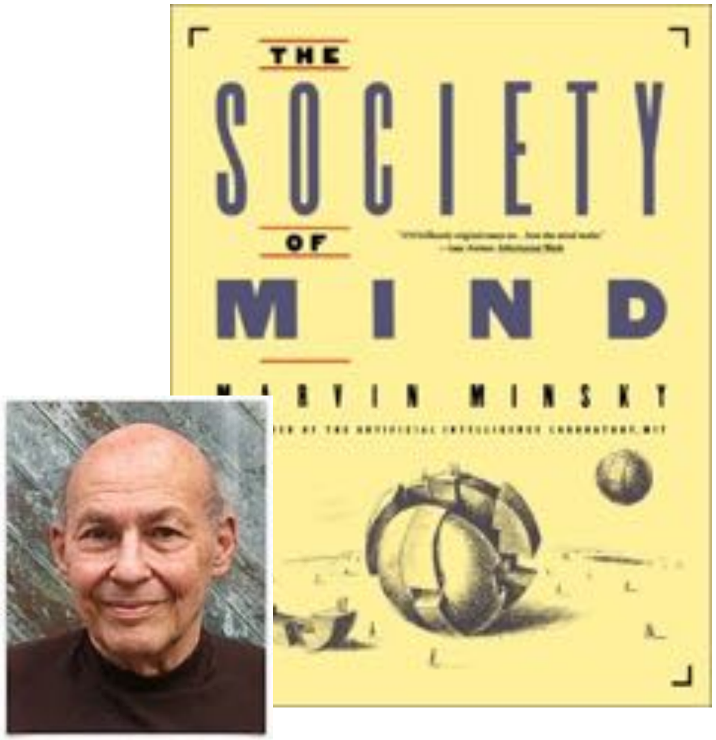

Society of Mind (Minsky, 1986)
Idea that human intelligence is not a centrally controlled system, but rather a *confederate of competing interests*

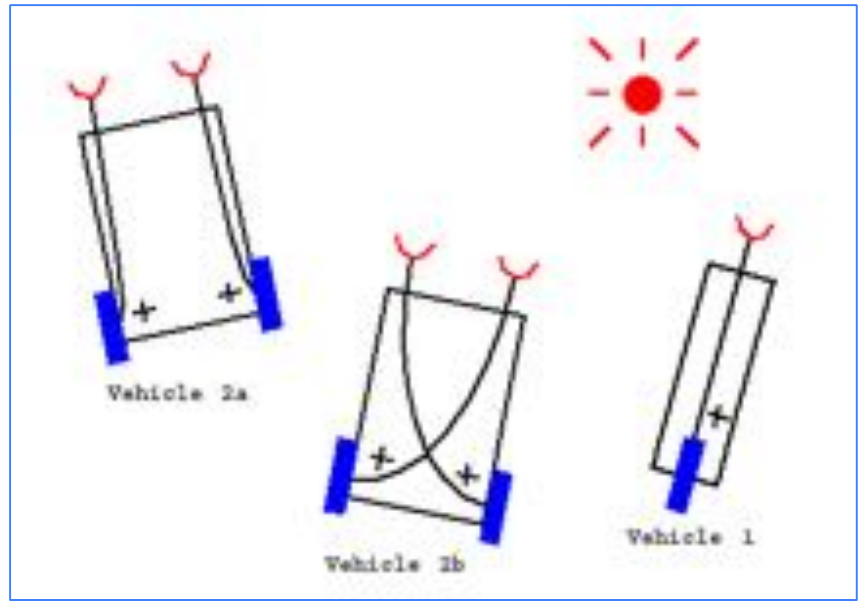

Braitenberg Vehicles:
Experiments in Synthetic Psychology, MIT Press [1984]

Behavior-based Robotics (1980s-90s)

- Rodney Brooks
  - *Started Behavior-based Robotics movement*
  - PhD at Stanford, 1981; Joined MIT 1984
  - MIT CSAIL director, 1997-2008
  - IJCAI Computers and Thought award (1991)
  - Co-Founder of iRobot (1991), Rethink Robotics (2009)
- Wrote Papers with Provocative titles!
  - *Elephants don't play chess*
  - *Intelligence without representation*
  - *Intelligence without without reason*
  - *Planning is a just a way to avoid figuring out what to do next…*
- Robots
  - Incredible number of robots to prove his point
  - And produced a generation of famous students

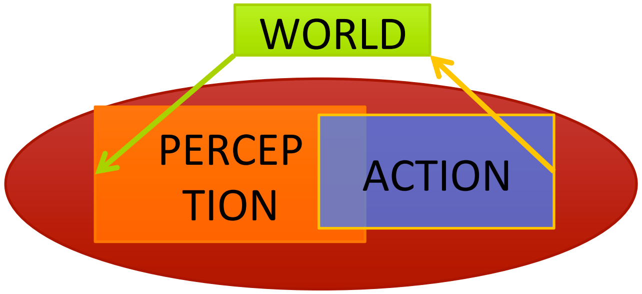

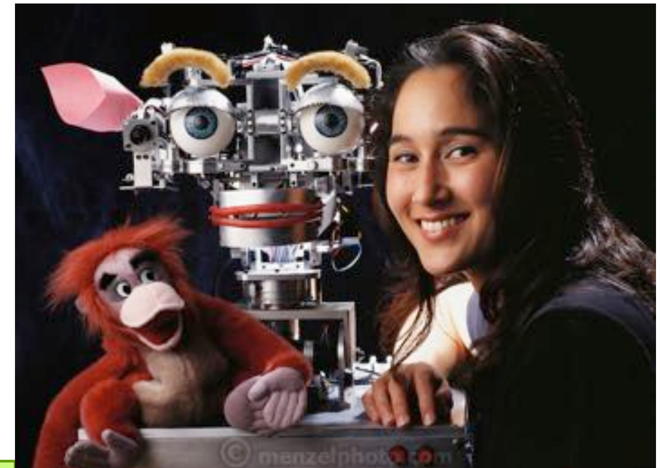

Generation of Robotics Faculty and Entrepreneurs (to name a few)

Maja Mataric (USC)
Yoky Matsuok (Uwash, NEST, MacArthur Genuis)
Brian Scasellati (Yale)
Cynthia Brazael (MIT, Founder Jibo)
Aaron Edsinger (Founder Meka Robotics)
James McLurkin (Rice, Google)
Colin Angle (iRobot)
Helen Greiner (Founder iRobot, CyPhyWorks)
Holly Yanco (Umass)
[clockwise]

Basic Idea: Reactive Architectures

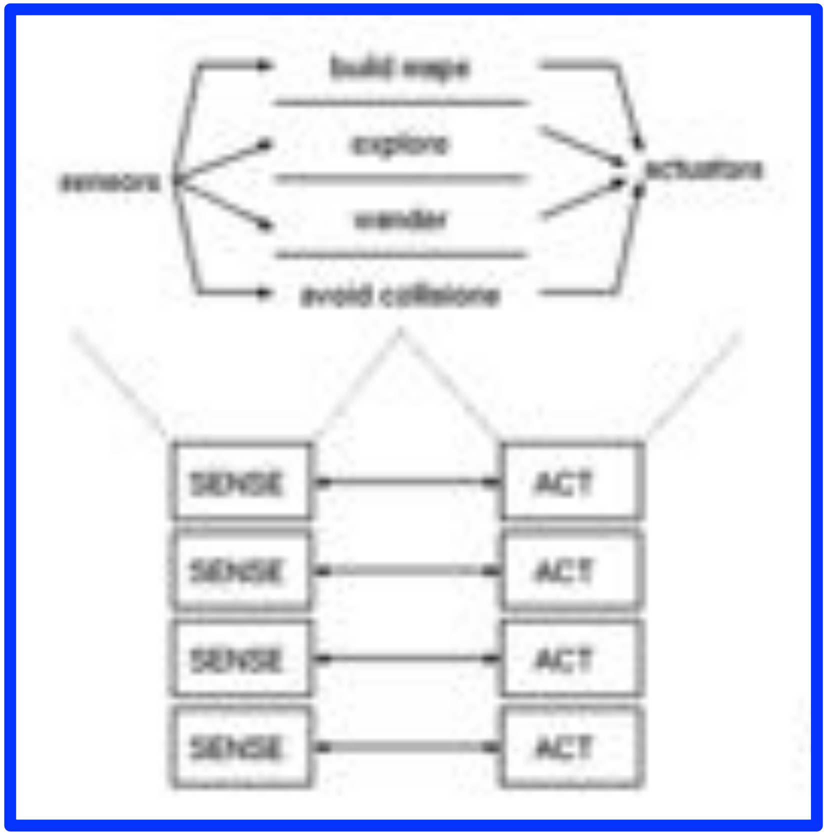

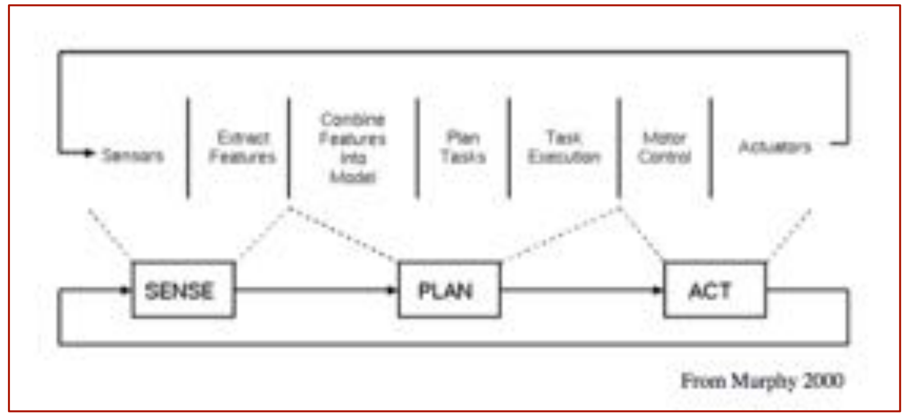

NEW Model

Concurrent "Behaviors"

Every behavior can access raw inputs, and affect raw outputs…)

"Subsumption" Architecture

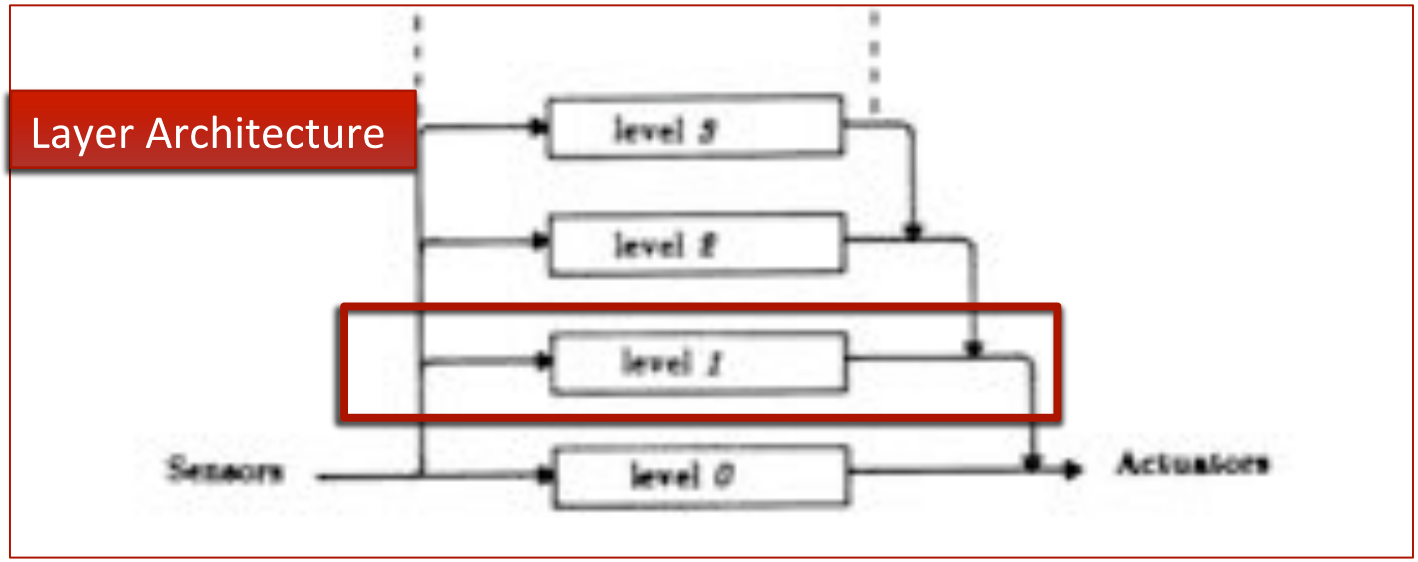

- **Rod Brooks defined the "Subsumption" Architecture**
  - Layers of Competency; Inhibit each other
- **Behavior-based Control** (Ron Arkin, Maja Mataric, etc)
  - E.g. Adding an "arbitrator" box between behaviors and actuators

Famous Examples

Layered Control for a Mobile Robot (1984)

- Reason about objects
- Plan changes to the world
- Identify objects
- Monitor changes
- Build maps
- Explore
- Wander
- Avoid objects

Decentralized Control for a Six-legged Robot (1989)

Key Features, and Bugs

- Ideas
  - Task is composed of Behaviors
  - Layers of Competence
  - Sensor->Action relationship in all behaviors
  - Concurrent Execution, Avoid memory and internal state
- Problems
  - Creativity in designing systems, but hard to analyze
  - Not all complex tasks easy to think of this way
  - Pre-planning is helpful in certain cases
- But Many Lasting Ideas: *“The world is its own best model”*
  - Avoid world->symbolic transformation
  - Assume highly dynamic world (less reliance on knowing and memory)
  - Perception is direct, distributed, and ego-centric

Where is behavior-based robotics today?

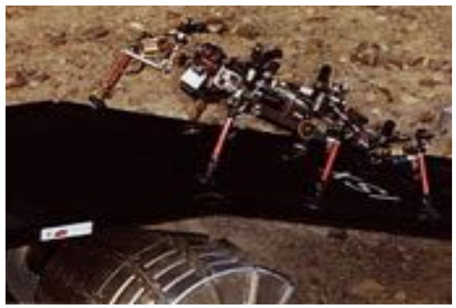

Modern Robotics (2000-now)

- **Competitions moved the field forward**
  - AAAI Scavenger hunts and “Hor D’ Oeuvres Anyone?”
  - Robot Soccer (goal: beat human champions in 2050)
  - Darpa Urban Challenge (self-driving cars)

- **Many new architectural ideas…**
  - *Probabilistic Reasoning for Navigation*
  - *Coordination Architectures (Role-based, Play-based, etc)*
  - *Human-Robot Interaction Architectures*

Probabilistic Robotics

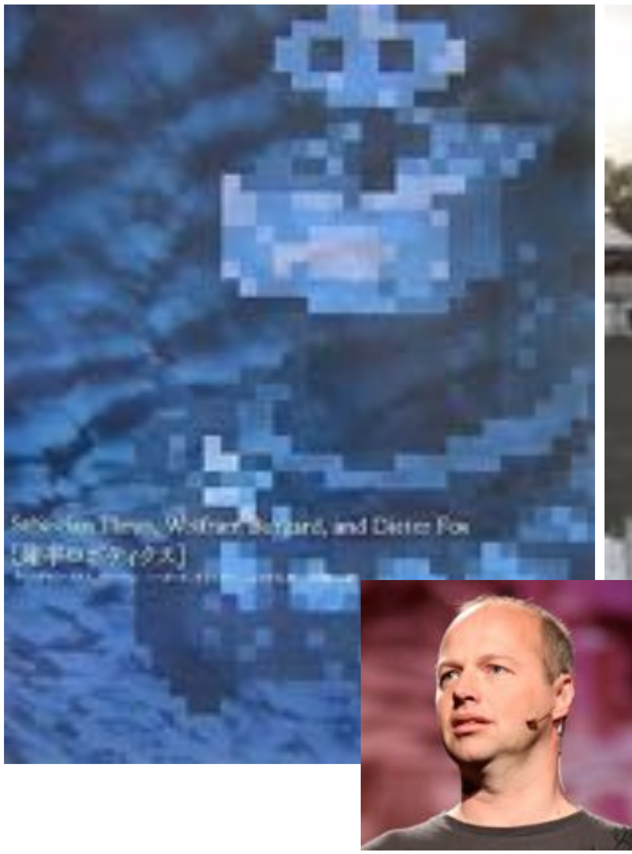

- Prof. Sebastian Thrun, Stanford
  - E.g. FastSLAM = Simultaneous Localization and mapping

Personal Robots

GRACE, 2002

KISMET(MIT) to Jibo

Diligent Robotics

Multi-Robot Systems

Grand Challenge for robotics and AI: *By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team.*
(Prof Manuela Veloso, CMU, co-founder)

Modern Robotics (2000-now)

- Competitions moved the field forward
  - AAAI Scavenger hunts and “Hor D’ Oeuvres Anyone?”
  - Robot Soccer (beat human champions in 2050)
  - Darpa Urban Challenge (self-driving cars)
- Many new architectural ideas…
  - Probablistic Reasoning for Navigation
  - Coordination Architectures (Role-based, Play-based, etc)
  - Human-Robot Interaction Architectures
- **History of Robotics is far from “written”**
  - **Most debates on autonomy remain open**

Pset 3: Two things

- Wandering: Rough Space Map
  - Similar corridor widths to last time, but some “side arm” areas.
  - Many possible wandering strategies, focus on trying to “see” all areas.
  - Include Bump+cliff+wheeldrop

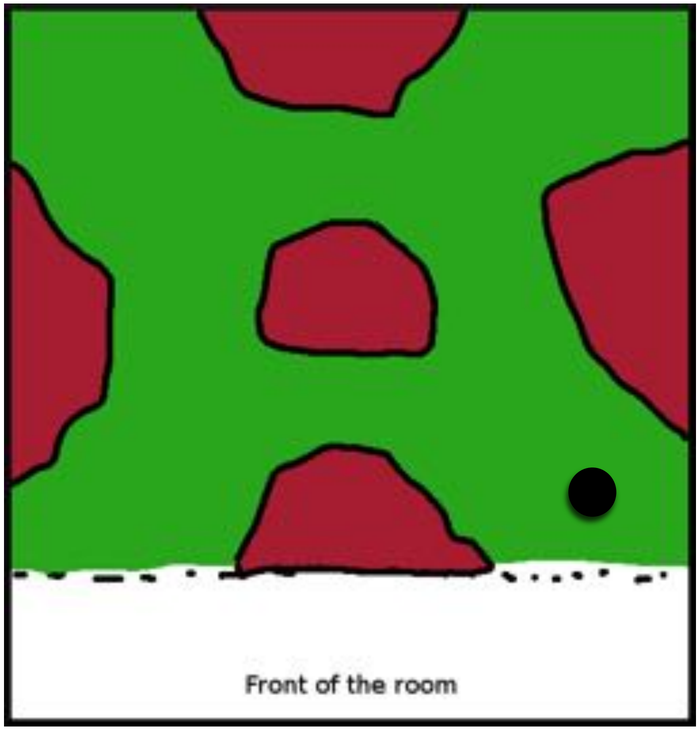

- Next Class: Competition “Split”
  - First round: 5 robots and teams
  - Second round: 5 robots and teams
  - *You only need to come for your round*

Layer 0: Don’t collide!

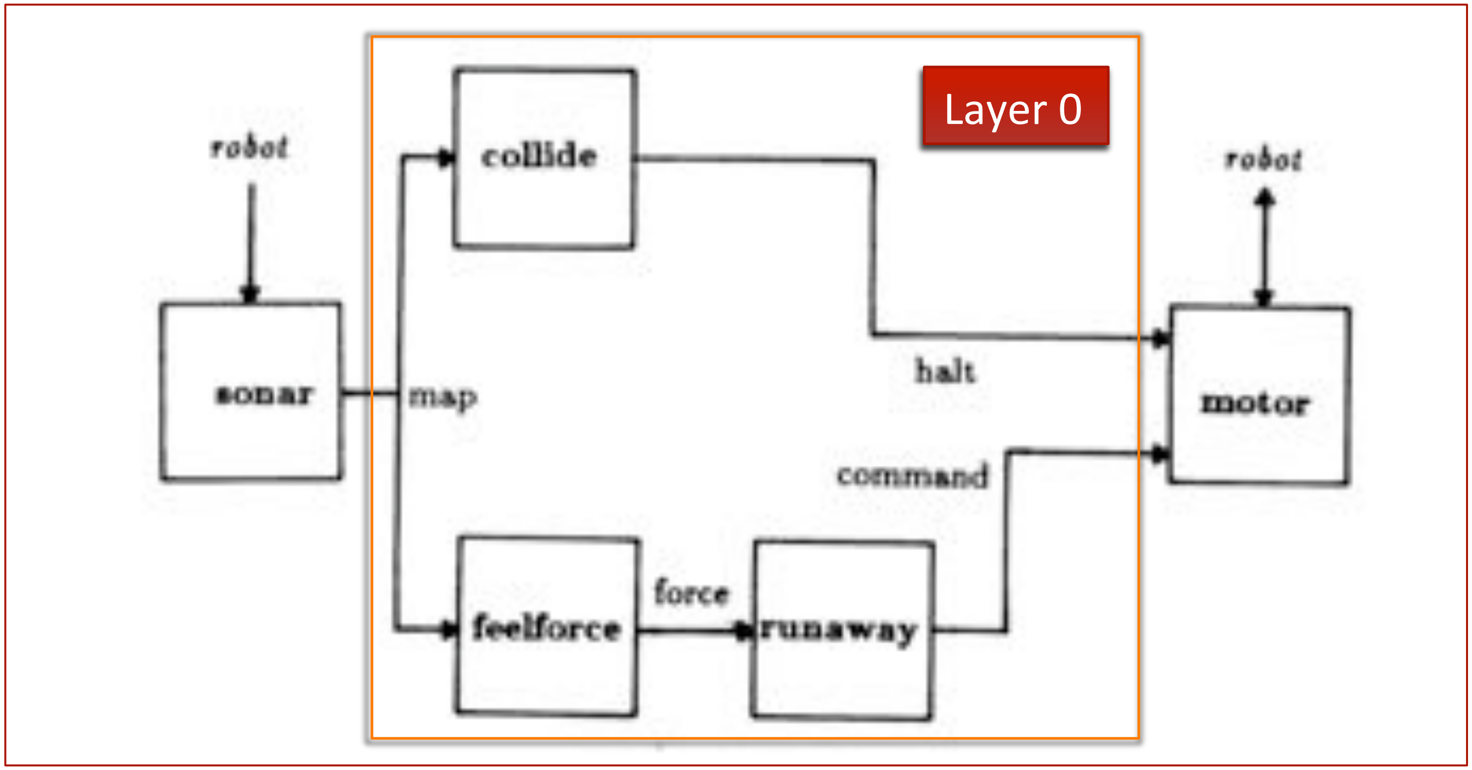

If obstacle appears, then “runaway”. Otherwise just stand still…
(Allows a person to “herd” the robot)

Layer 1: Wander

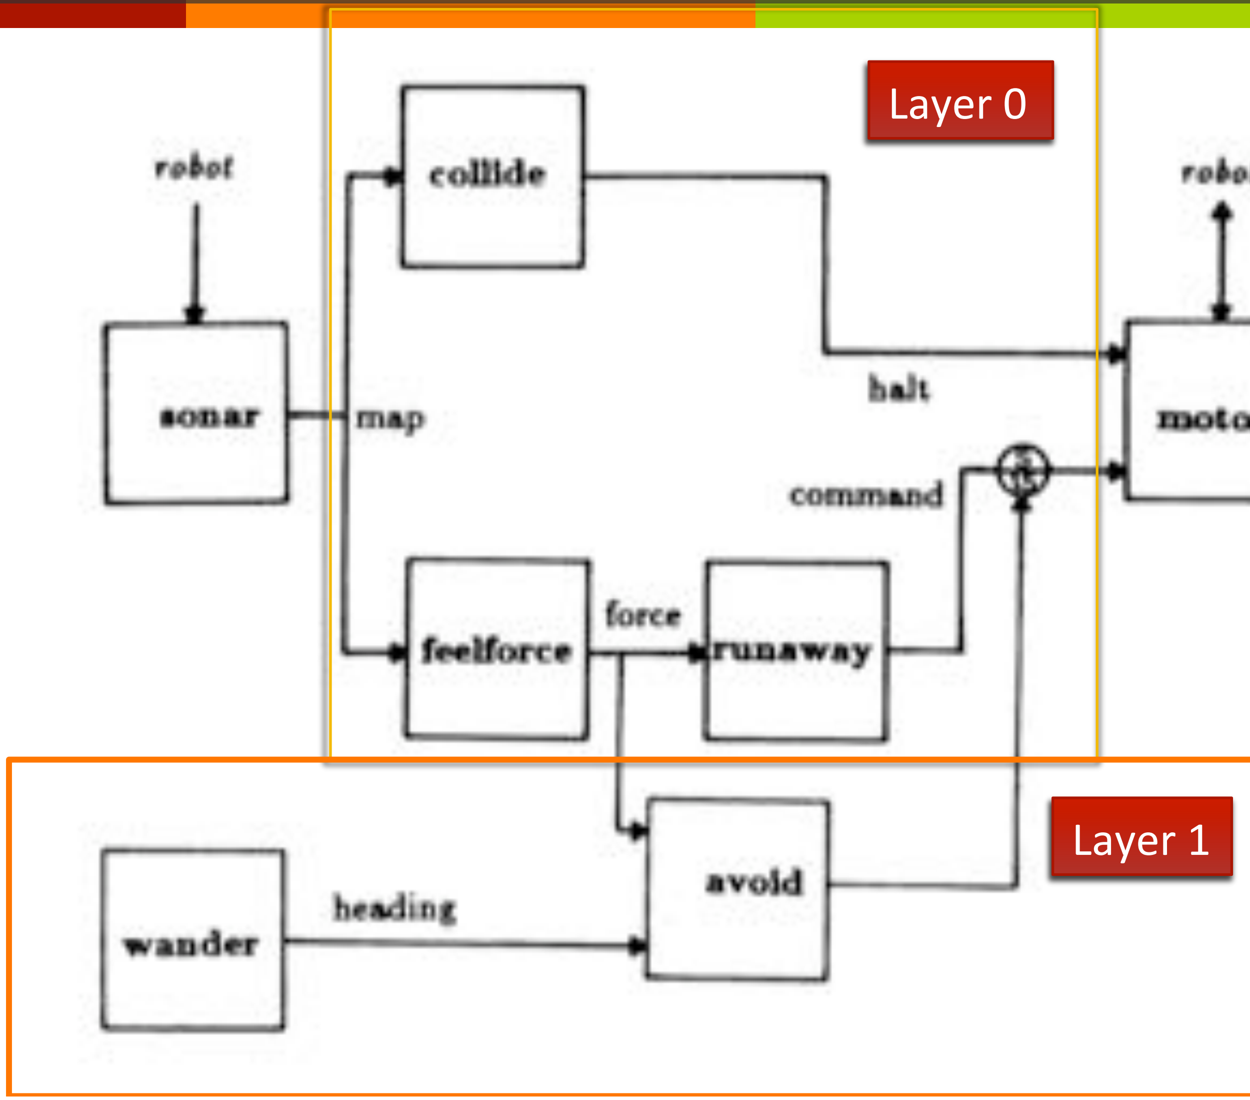

Layer 2: Explore

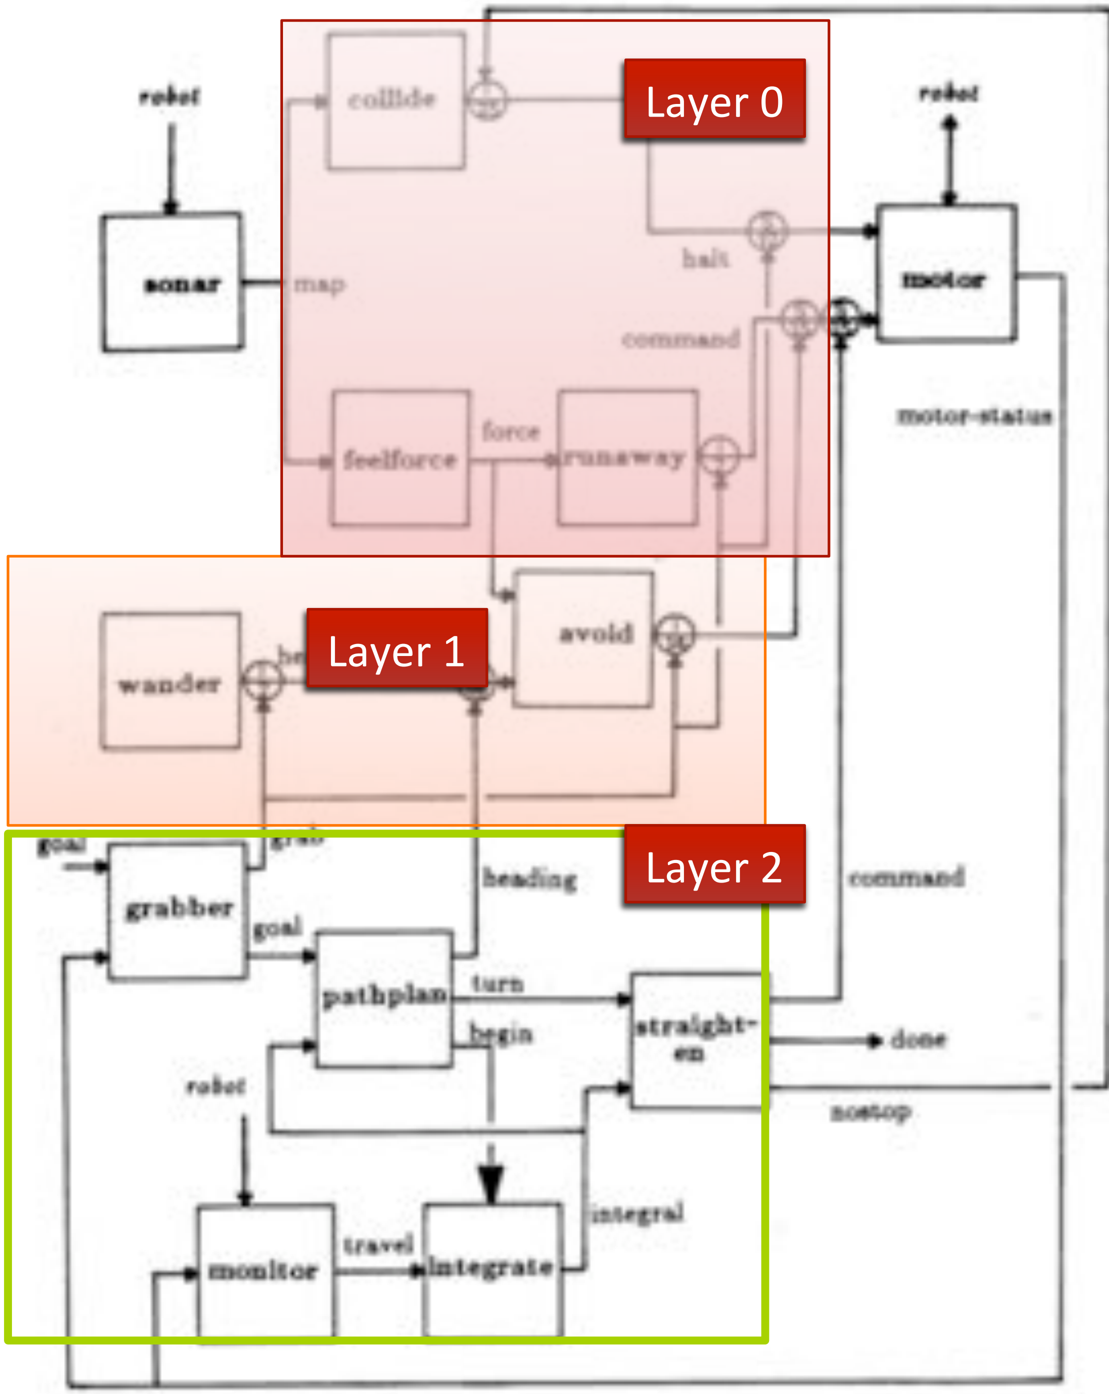

- **Reason about objects**
- **Plan changes to the world**
- **Identify objects**
- **Monitor changes**
- **Build maps**
- **Explore**
- **Wander**
- **Avoid objects**

Genghis and other "Veterans"

- Six-legged Robot
  - Many degrees of freedom, many sensors (hard!)
    - Standup layer
    - Simple Walk layer
    - Balance/Pitch Layer
    - Whiskers Behavior
    - Steer, Prowl (follow)
    - Every layer " subsumes" other!
    - "Decentralized intelligence"

    (seemingly purposeful behavior)

- Subsumption: Main Ideas
- • Modules should be grouped into *layers of competence*
- • Modules in a higher level can override or *subsume* behaviors in the next lower level
- – Suppression: substitute input going to a module – Inhibit: turn off output from a module
- • ***No internal state*** in the sense of a local, persistent representation similar to a world model.
- • Architecture should be *taskable:* accomplished by a higher level turning on/off lower layers

Vertical Behaviors: Many Active at Once

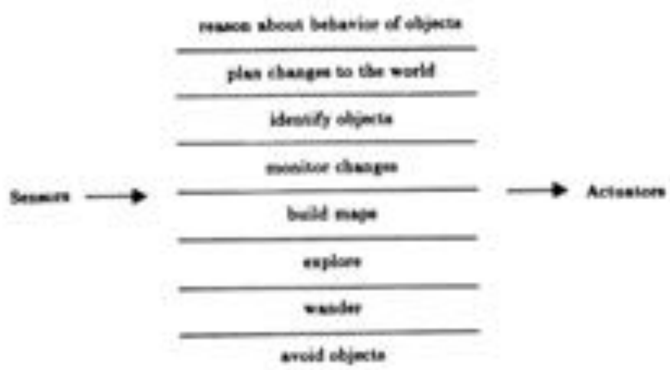

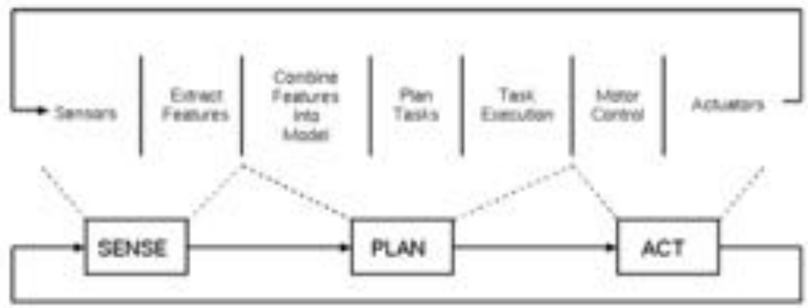

From Murphy 2000

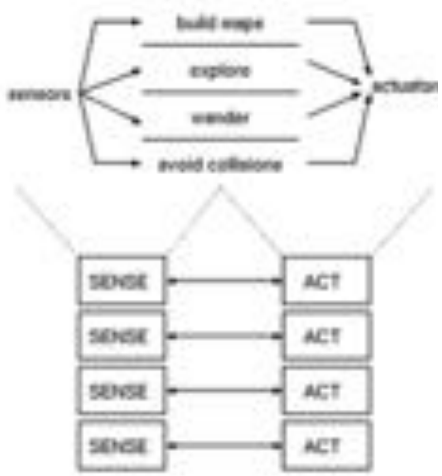

Example: Docking Behavior

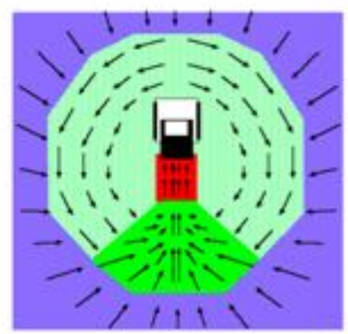

Reactive Architectures

Related Ideas
Braitenberg Vehicles
and
Society of Mind (Minsky)

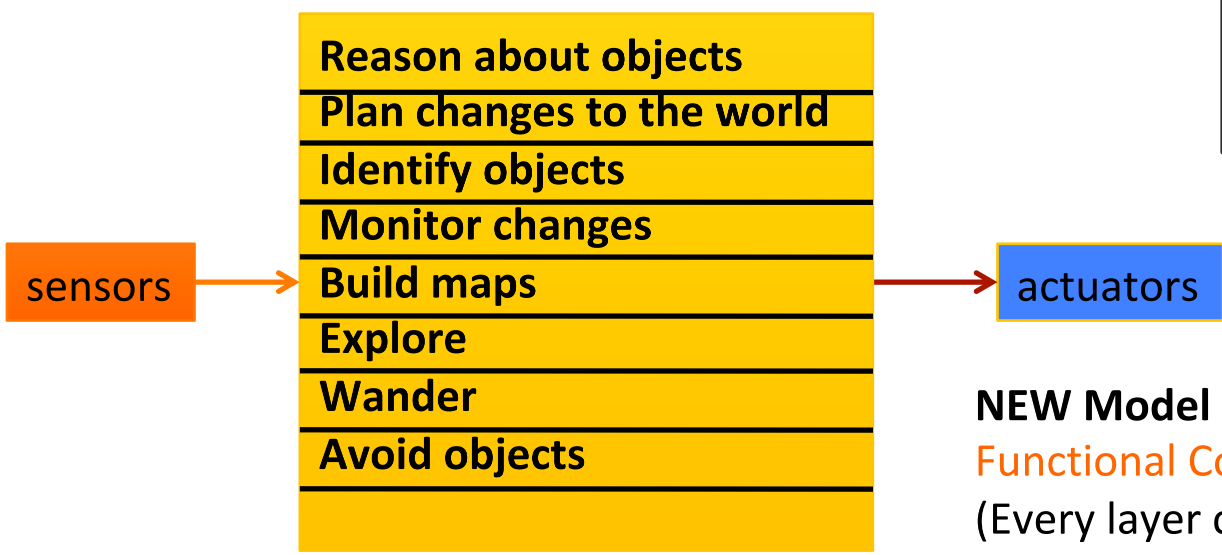

NEW Model
Functional Concurrent Decomposition!
(Every layer can access raw inputs
And affect raw outputs…)

"Veterans" Robots

- Survey of Robots
  - A Robust Layered Control System for a Mobile Robot, AI Lab Memo 1985

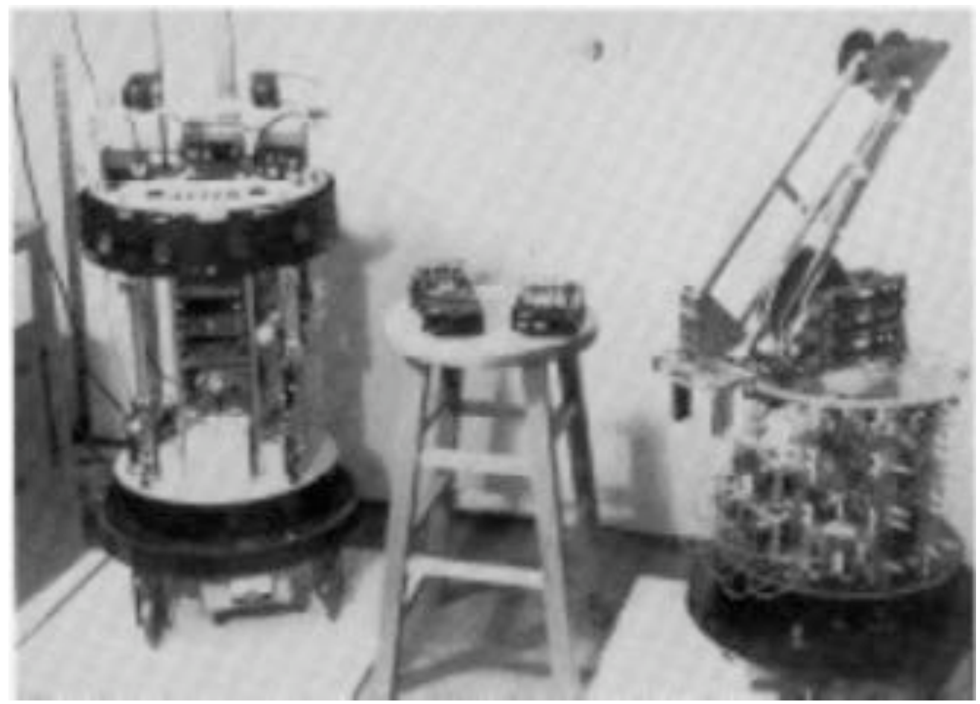

- Requirements
  - Multiple Goals
  - Multiple Sensors
  - Robustness
  - Extensibility
  - *Autonomy*

- Principles!
  - Simplicity in control
  - Complexity in world

"The world is its own best model"

Which way is better?

- Hybrid Architectures [Ron Arkin, GaTech]
  - Compose lower layers in behavior-based style
    - Sensor-act behaviors acting *quickly* and *concurrently*
  - Impose a higher-level reasoning system
    - Coarse grained knowledge of the world and goal
    - Slower time-loop within which to plan
  - Many Examples: Robot Soccer, Self-Driving Cars

- The Debate is not over!
  - E.g. Mapping vs Bio-inspired Optical-flow
  - Designing an autonomous Robot vacuum