4a Architectures for Autonomy

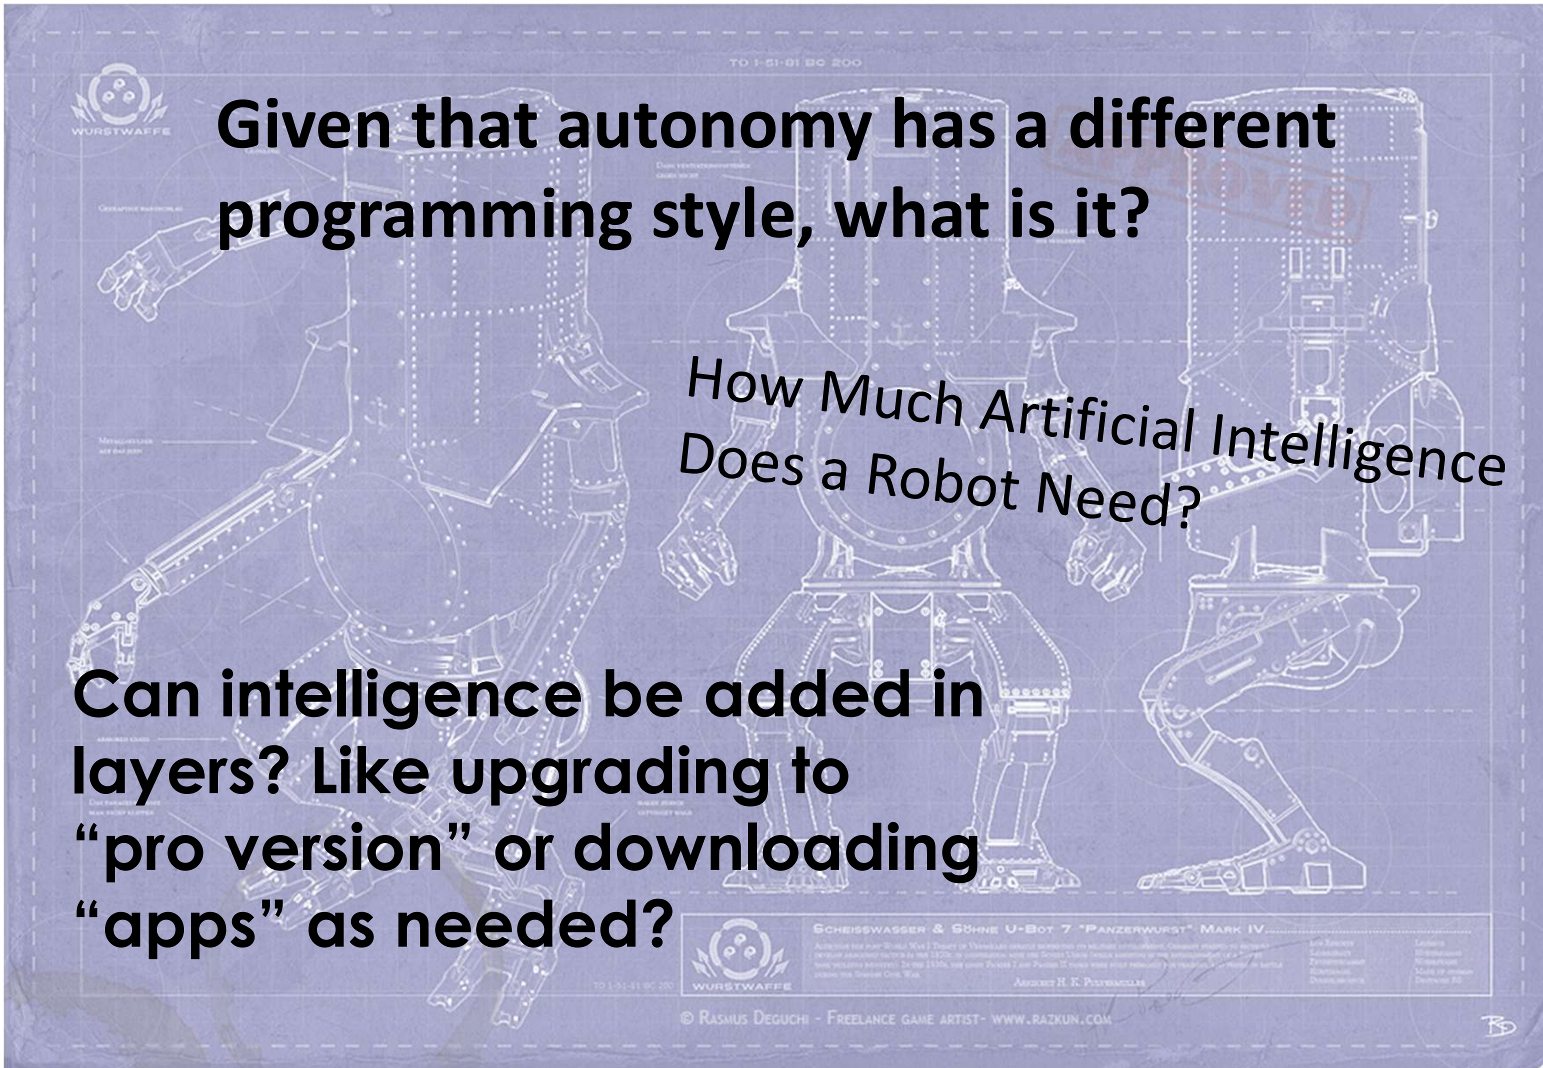

Given that autonomy has a different programming style, what is it?

How Much Artificial Intelligence Does a Robot Need?

Can intelligence be added in layers? Like upgrading to “pro version” or downloading “apps” as needed?

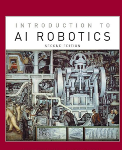

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

http://www.gameartisans.org

4a Specific Learning Objectives

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- Apply Levis' definition of architectures to organizing software in an intelligent robot and be familiar with how they can guide the design process

- Name and describe the layers (behavioral, deliberative, interface) within a canonical operational architecture of an intelligent robot in terms of the five attributes

- Name the four primitives of robotic intelligence (sense, plan, act, learn)

- List vulnerabilities (e.g., where poor choices or implicit assumptions can lead to failures) in the canonical operational architecture

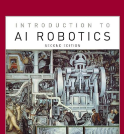

ROBIN R. MURPHY

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Recap: Where We've Been…

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- AI researchers view robotics as an issue of *increasing autonomy* (doing the "right thing" in an open world),
- whereas many engineering researchers view robotics as an issue of *extending automation* (focus on creating a closed world and guaranteeing actions).
- This gives rise to very different, though not necessarily mutually exclusive, approaches especially in terms of programming.

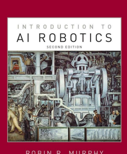

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

The Big Picture

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- AI has converged on a canonical operational architecture after years of exploration
  - This will be the backbone of additional lectures as we delve into the systems architecture and choice for algorithms, coordination functions, etc.
  - Historical development is of interest because it helps identify legacy code and to see barriers that led to the convergence (separate lecture)

- Operational architectures have advantages
  - S/W engineering principles of abstraction, modularity
  - Semiformal design specification and completeness
  - Can reveal (some) fundamental vulnerabilities

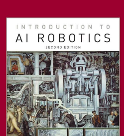

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Organizing Software

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- Overall style of design or organization is called an "**architecture**"
  - provides a principled way of organizing a control system. However, in addition to providing structure, it imposes constraints on the way the control problem can be solved *[Mataric]*
  - describes a set of architectural components and how they interact *[Dean & Wellman]*

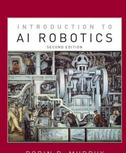

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Types of software architectures
[Levis, George Mason University]

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- **operational architecture:** describes ***what*** the system does, or its functionality, at a high level, but not how it does it

- **systems architecture:** describes how a system is decomposed into ***major subsystems***

- **technical architecture:** describes how a system works in terms of ***implementation details***, language

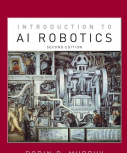

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

Focus of this Course

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- **operational architecture:** describes what the systems does, not how it does it
- **systems architecture:** describes how a system works in terms of major subsystems
- **technical architecture:** implementation details, language

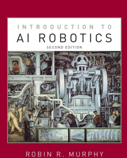

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Often Changes with New Languages, Applications

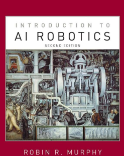

- **operational architecture:** describes what the systems does, not how it does it

- **systems architecture:** describes how a system works in terms of major subsystems

- **technical architecture:** implementation details, language

Subject of Standards Wars

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Why Worry About Software Organization?

- **Recall at least 7 distinct areas of Artificial Intelligence**, each with own algorithms and data structures; these have to be “knitted” together somehow

- Software engineering is necessary for a successful software enterprise

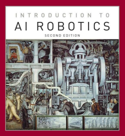

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Thinking About Architectures is Good Software Engineering

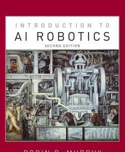

- **Abstraction**
  - Ignores details to permit focus for thinking about general organization of intelligence
  - Semiformal description
- **Modularity**
  - high cohesion (do one thing well, may be able to substitute different algorithms which produce same results “logical sensors”, “logical behaviors”, etc)
  - low coupling (which means may be able to add “apps”)
  - Supporting unit testing and debugging
- **Anticipation of change, Incrementality**
  - How to adapt, support evolution
- **Generality**
  - Not re-invent the wheel each time

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Operational Architecture of Biological Intelligence*

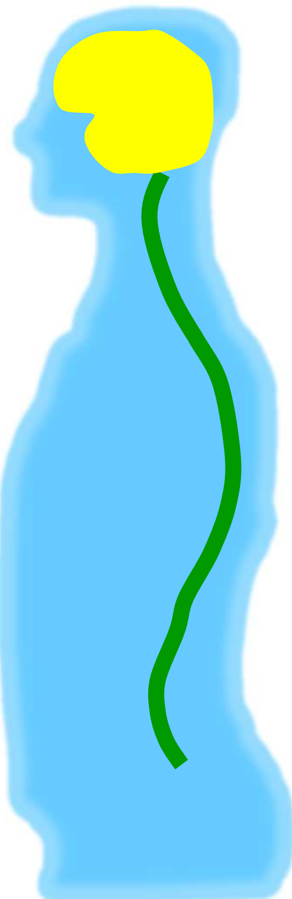

**An amazingly sweeping generalization for the purpose of metaphor*

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

Lower Central Nervous System

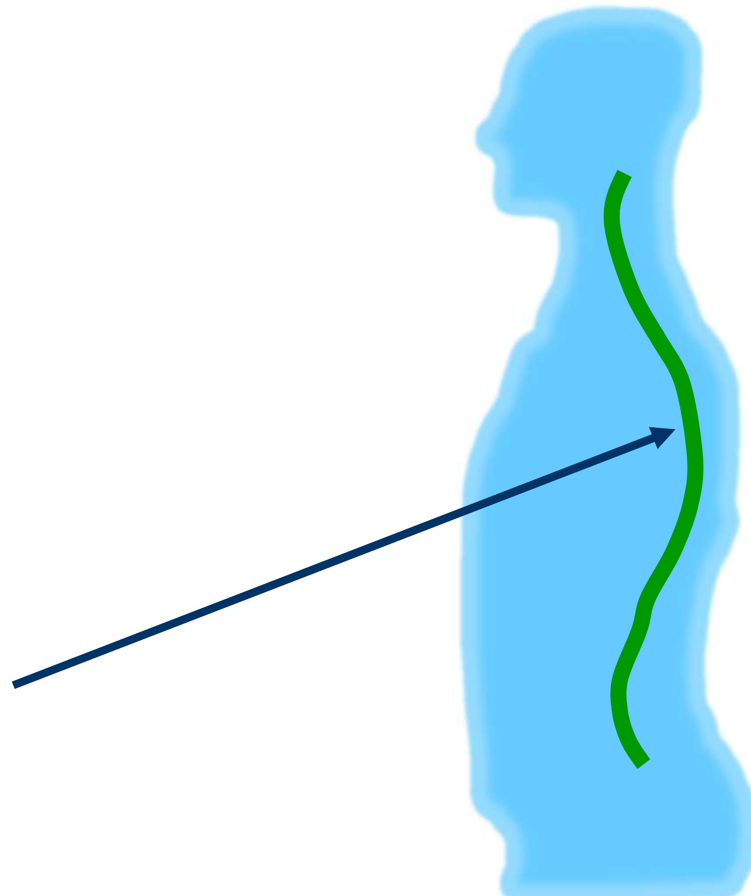

Spinal Cord and "lower brain"
Skills and responses

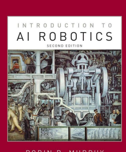

**An amazingly sweeping generalization for the purpose of metaphor*

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Loop 1: Reflexes, Reactive, Unconsciousness…

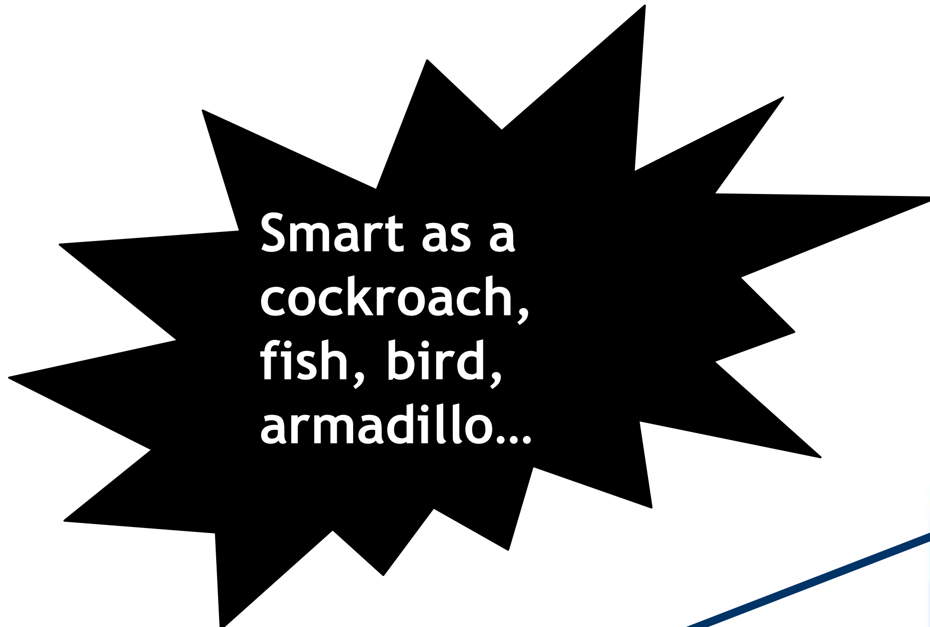

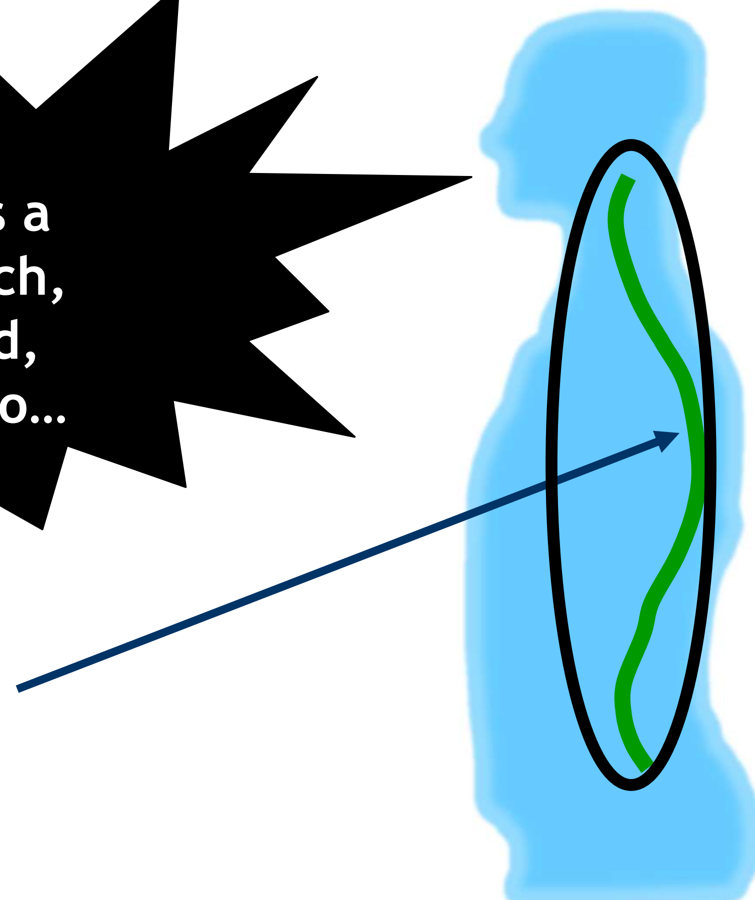

Spinal Cord and "lower brain"
Skills and responses

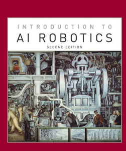

**An amazingly sweeping generalization for the purpose of metaphor*

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Brain

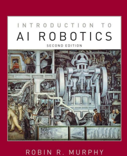

“Upper brain” or cortex
Reasoning over symbols (information) about goals

**An amazingly sweeping generalization for the purpose of metaphor*

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Loop 2: Deliberative, Thoughtful, Conscious…

“Upper brain” or cortex
Reasoning over symbols (information) about goals

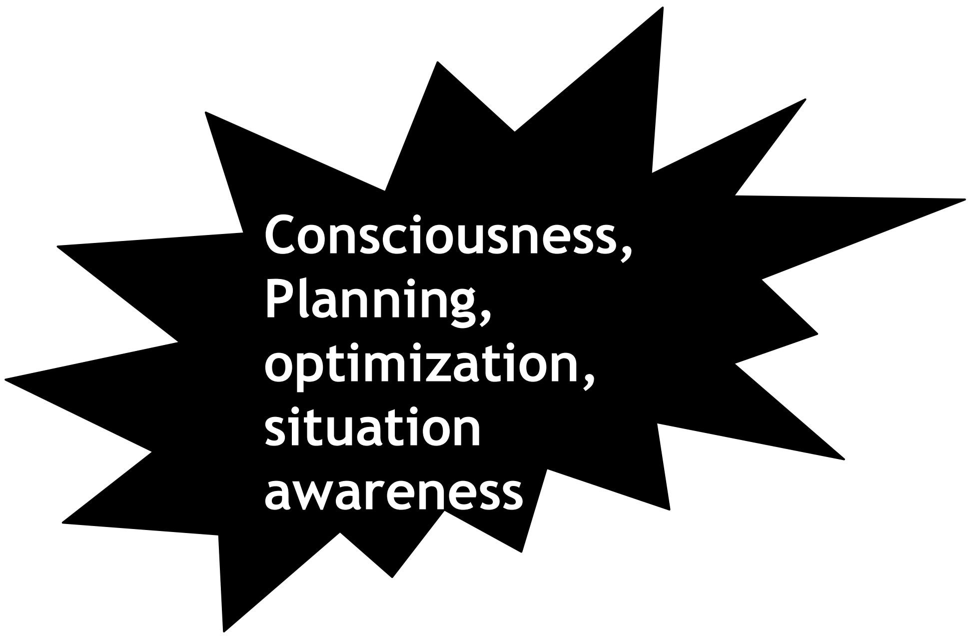

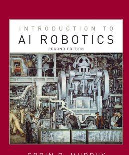

**An amazingly sweeping generalization for the purpose of metaphor*

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Direct Perception to Symbols

"Upper brain" or cortex
Reasoning over symbols (information) about goals

"Middle brain"
Converting sensor data into symbols (information)

Spinal Cord and "lower brain"
Skills and responses

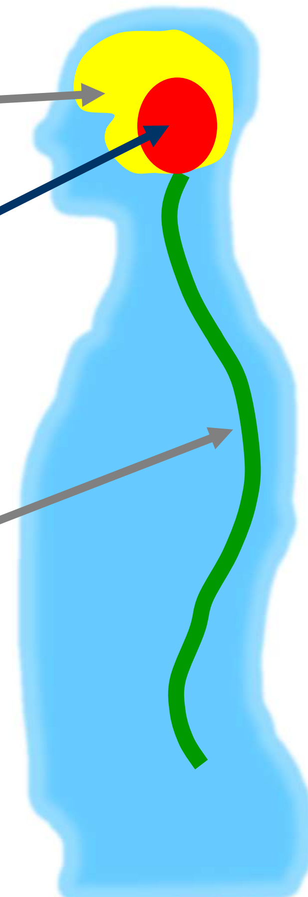

**An amazingly sweeping generalization for the purpose of metaphor*

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Two Loops Plus a Transform

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

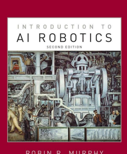

"Upper brain" or cortex
Reasoning over symbols (information) about goals

"Middle brain"
Converting sensor data into information

Spinal Cord and "lower brain"
Skills and responses

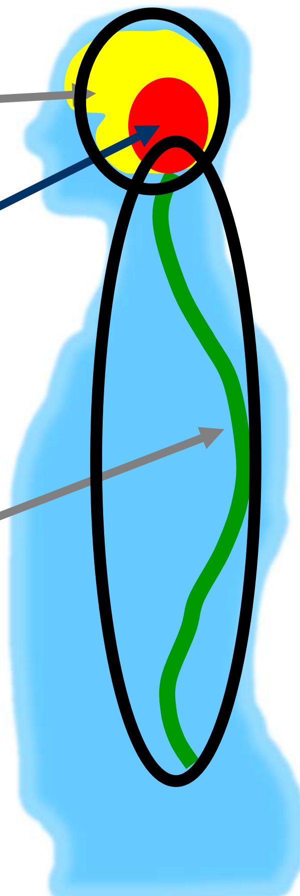

*An amazingly sweeping generalization for the purpose of metaphor

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a 2 Layers in Architecture

"Upper brain" or cortex
Reasoning over symbols (information) about goals

"Middle brain"
Converting sensor data into symbols (information)

Spinal Cord and "lower brain"
Skills and responses

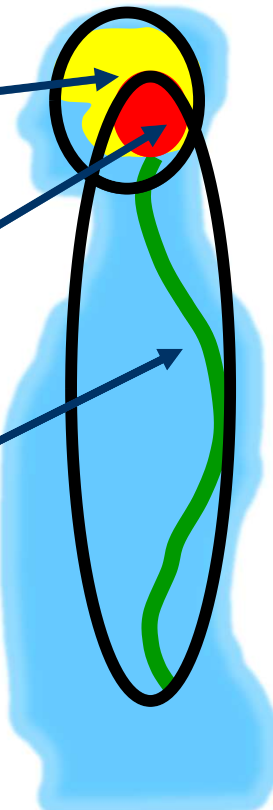

Deliberative Layer

Reactive (or Behavioral) Layer

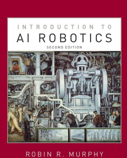

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

But There Is "Emotional Intelligence" Too

- User Interfaces
  - Displays, transparency of what robot is doing/thinking
  - Natural Language, gestures
- Working in teams
  - Explicit multi-agent coordination often relies on "social rules" (though implicit swarm intelligence may not)
- "Persona" we present to others
  - Security: what can you at your security grade see about me
  - Human-robot interaction: affective responses, natural language

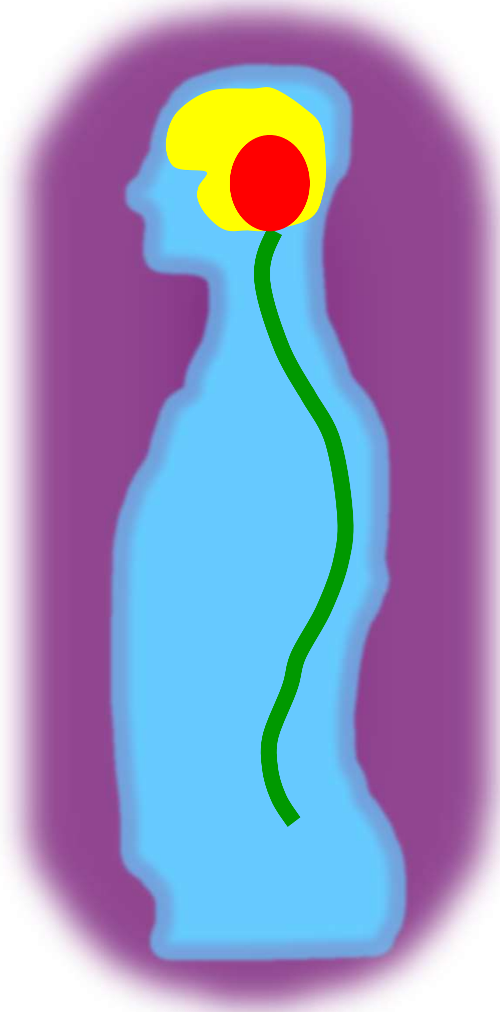

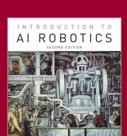

*An amazingly sweeping generalization for the purpose of metaphor

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Canonical Operational Architecture

“Upper brain” or cortex
Reasoning over symbols (information) about goals

“Middle brain”
Converting sensor data into symbols (information)

Spinal Cord and “lower brain”
Skills and responses

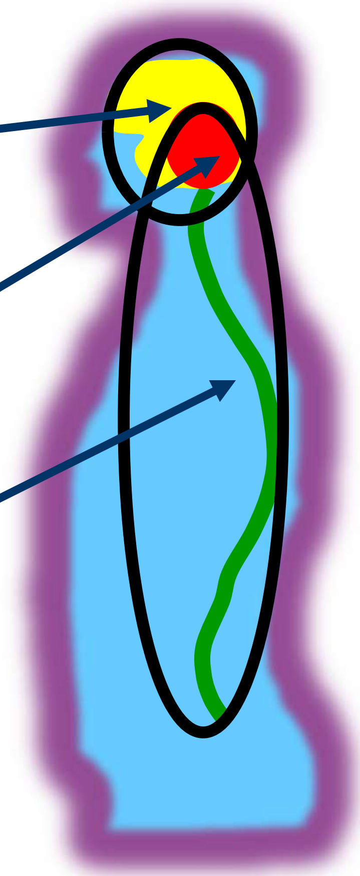

Interaction Layer

Deliberative Layer

Reactive (or Behavioral) Layer

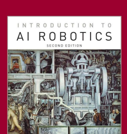

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Canonical Architecture

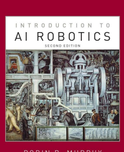

Interaction:
Persona and teaming

Deliberative loop:
Reasoning over symbols

Key function:
Converting sensor data into symbols

Reactive loop: ***Skills and responses***

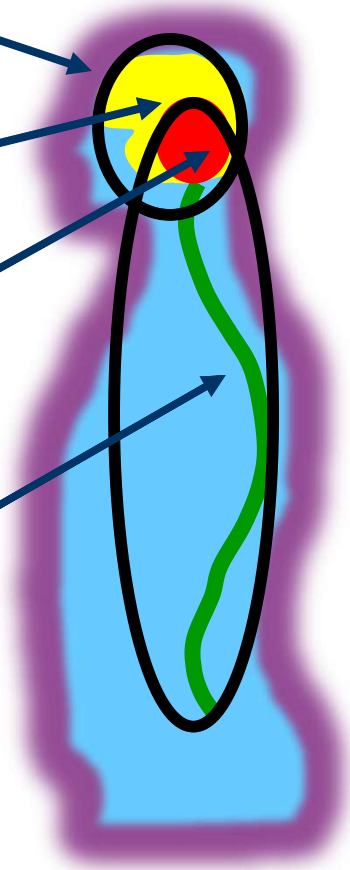

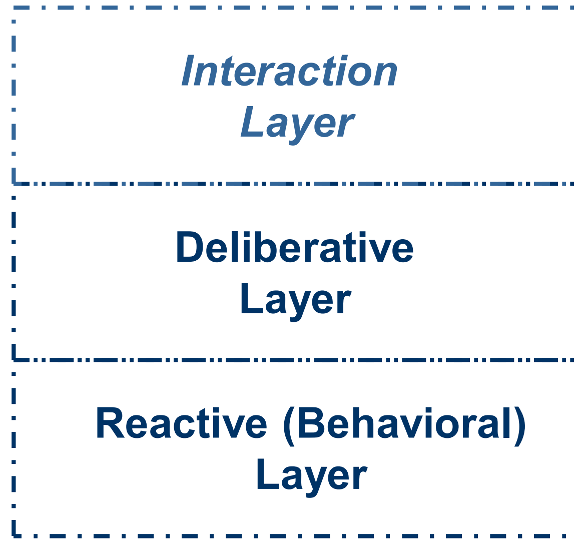

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

What Does This Mean for Programming?

Objectives
Review
Architecture
-types
Operational
-**Biological**
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- **From Reactive to Deliberative**
  - Two types of perception: DIRECT, RECOGNITION (symbols)
    - Impacts computer vision
  - Different time horizons
    - From Present to Present, Past, Future
    - Impacts sensing, storage, as well as algorithms reasoning, projecting
  - Different time scales of a function
    - Very fast (a reflex), fast (selection of functions), slow (reasoning about a problem)
  - Need a central structure (WORLD MODEL) to hold the symbols, history, knowledge but is tractable
- **From Reactive/Deliberative to Interaction**
  - Additional knowledge “theory of mind”– beliefs, desires, intentions (BDI) of the other agent, common ground

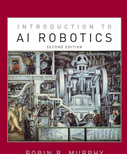

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a The Hard Parts

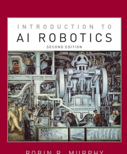

- **From Reactive to Deliberative**
  - Two types of perception: DIRECT, RECOGNITION (symbols)
    - Impacts computer vision
  - Different time horizons
    - From Present to Present, Past, Future
    - Impacts sensing, storage, as well as algorithms reasoning, projecting
  - Different time scales of a function
    - Very fast (a reflex), fast (selection of functions), slow (reasoning about a problem)
  - Need a central structure (WORLD MODEL) to hold the symbols, history, knowledge but is tractable
- **From Reactive/Deliberative to Interaction**
  - Additional knowledge “theory of mind”– beliefs, desires, intentions (BDI) of the other agent, common ground

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Primitives for Robot Intelligence

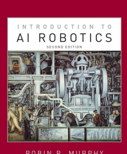

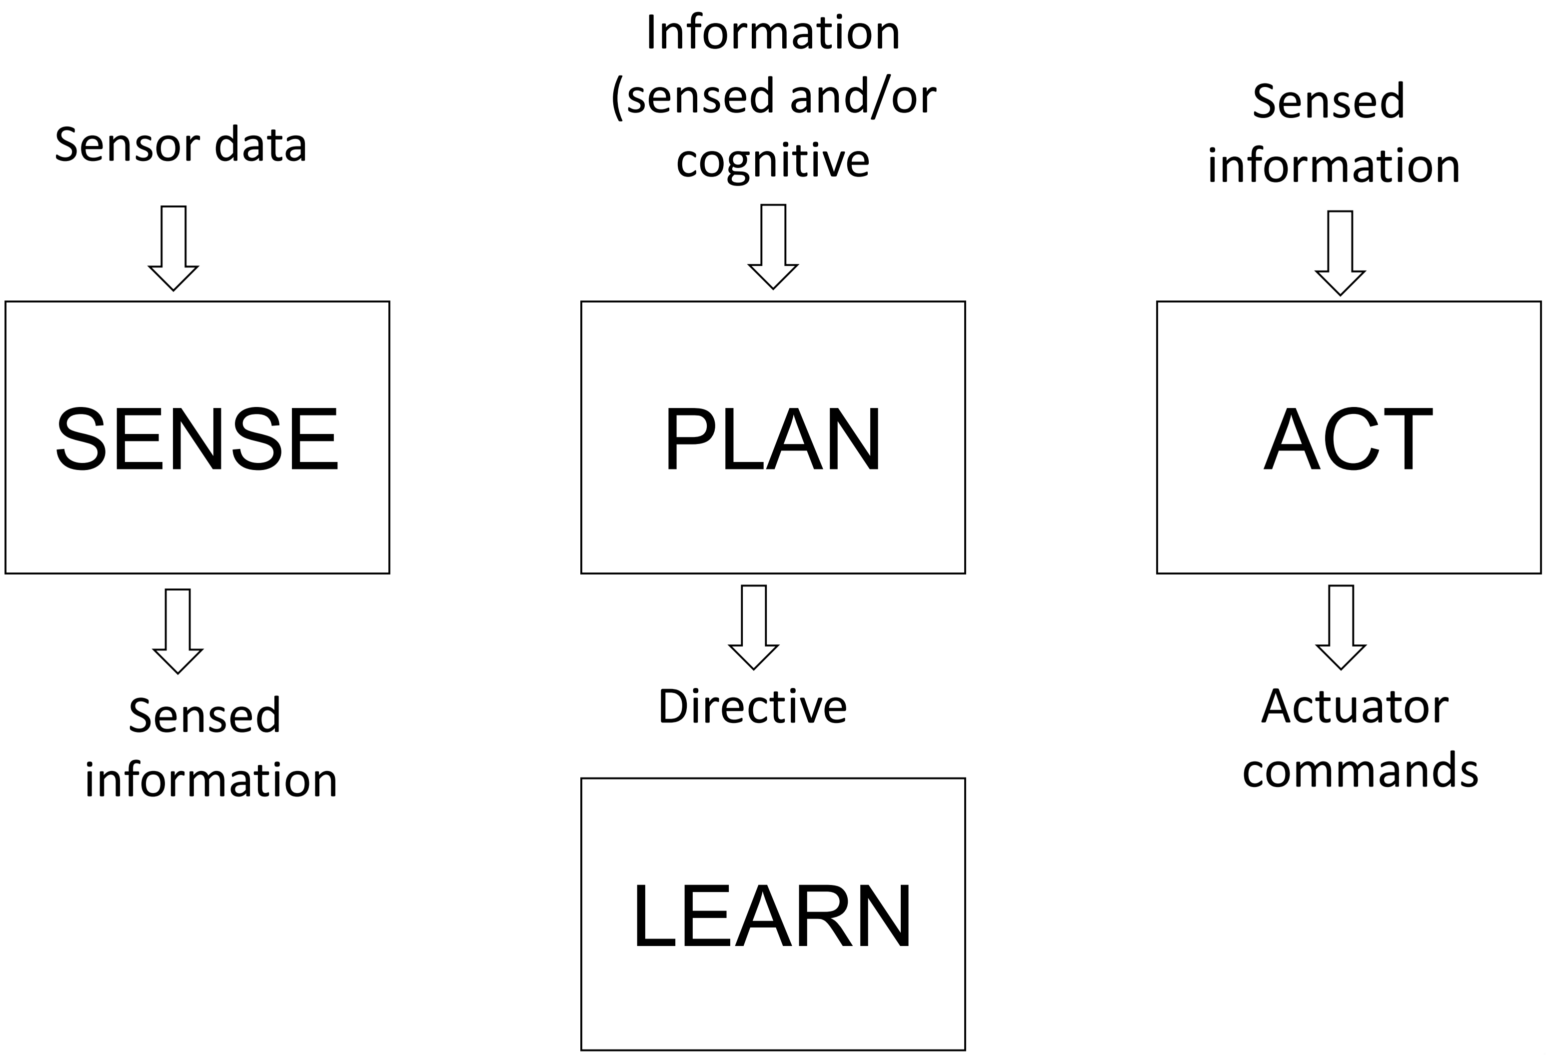

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

“Behavioral Robotics”

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

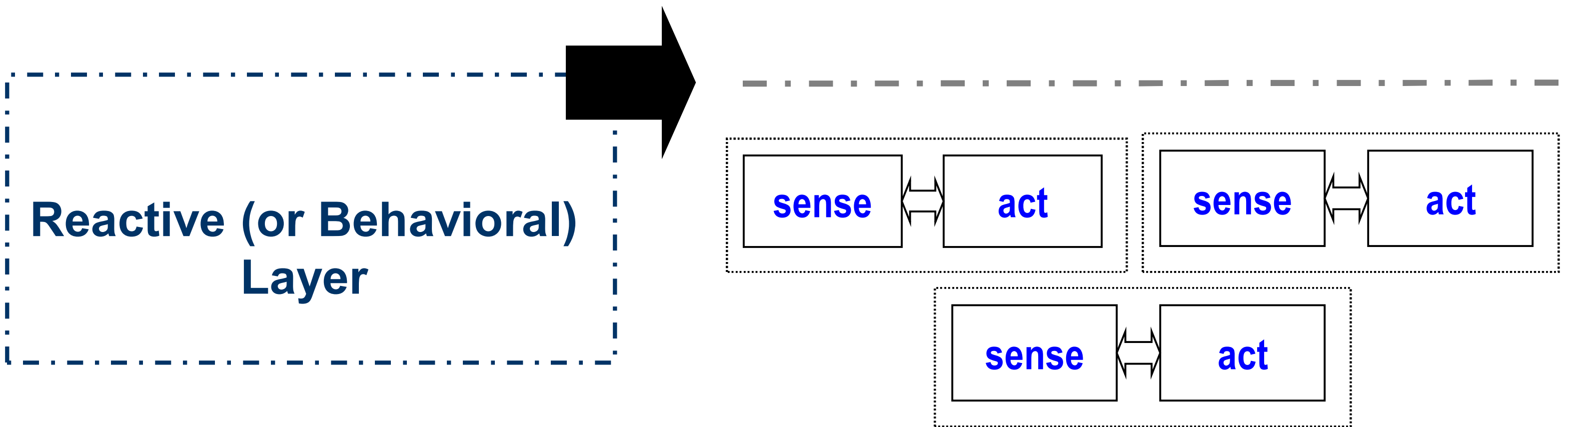

- “born” with sets of SENSE-ACT couplings called behaviors that get turned on/off based on stimulus
- No PLAN – So, the planning horizon is the present

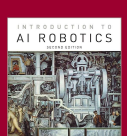

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Control Theory is "Lower Level" But Doesn't Necessary Capture it All

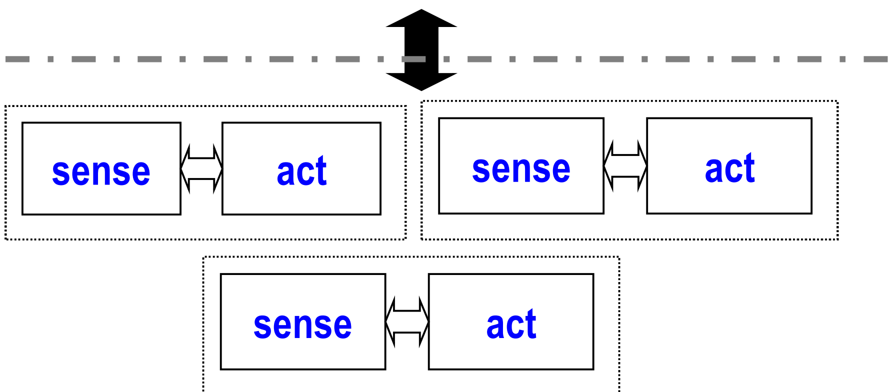

Reactive (inner loop control, behaviors):

- **Tightly coupled with sensing, so very fast**
- **Many concurrent stimulus-response behaviors, strung together with simple scripting with FSA**
- **Action is generated by sensed or internal stimulus**
- **No awareness, no mission monitoring**
- **May have local world model, serving as short term memory**

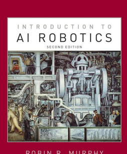

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a Behavioral Robotics

- Focus in AI from 1986-1996

- Things you can do with behavioral robotics
  - Roomba, Aibo
  - Guarded motions (don't hit anything), panic behaviors (stop!), "Macros" such as self-righting in teleoperation

- Advantages
  - Direction perception is usually simple
  - High modularity, add new behaviors without reprogramming old working behaviors, degrades gracefully (depending on the technical architecture)

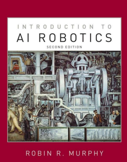

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Behavioral Robotics Issues

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- **Hidden Costs**
  - Behaviors plus some sort of ***coordination function*** to fuse outputs to effectors
- **Where it breaks**
  - "Fly at window" effects due to local scope (but lots of ways around this)
  - Poor choice of coordination functions (go left, go right = go in middle and hit obstacle)
- **Where it doesn't break but scares/annoys people**
  - Often can't predict whether will go left or right to avoid obstacle
  - Not optimal

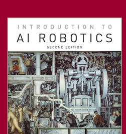

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

“Hybrids”

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

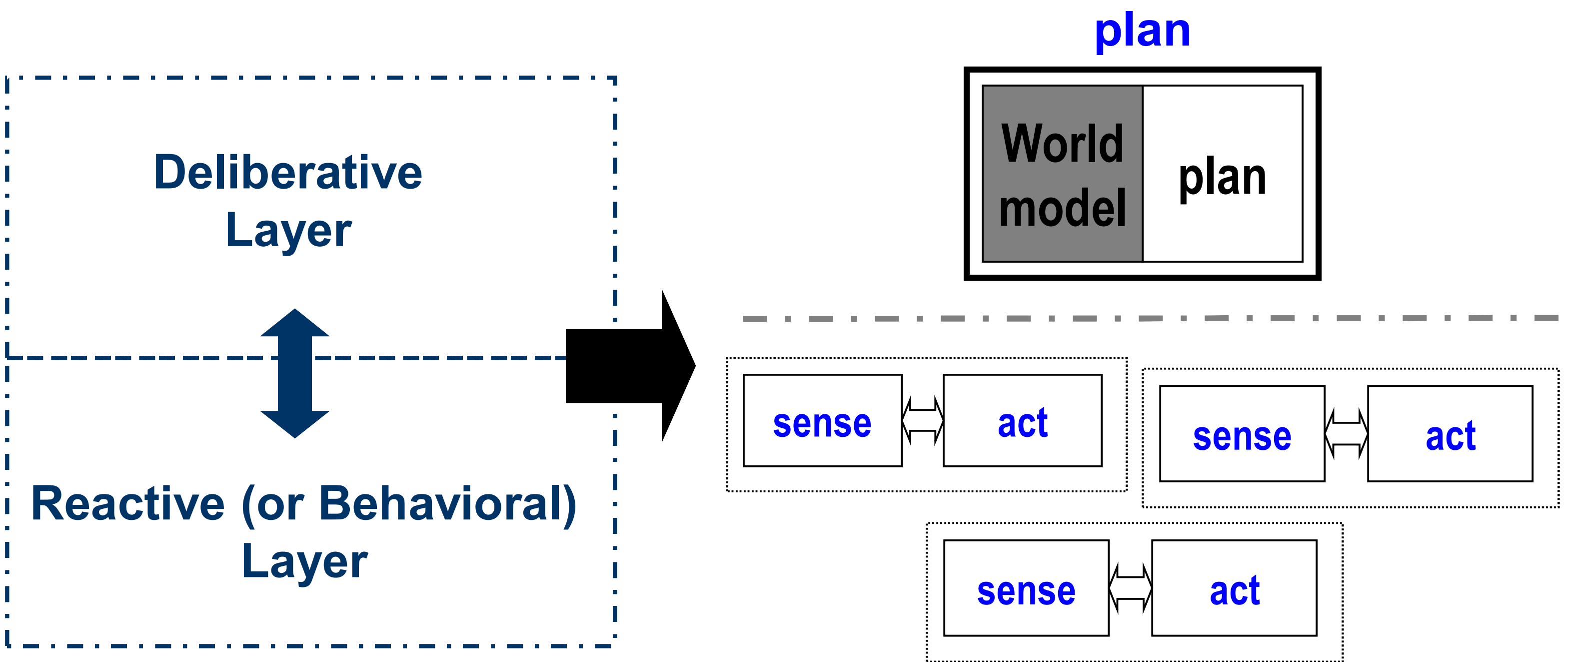

- PLAN, then instantiate appropriate SENSE-ACT behaviors, until next step in plan, …
- PLAN requires a World Model (though it is bounded) plus the actual planning algorithms

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

3 Tier Variant

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

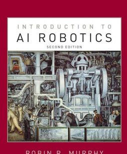

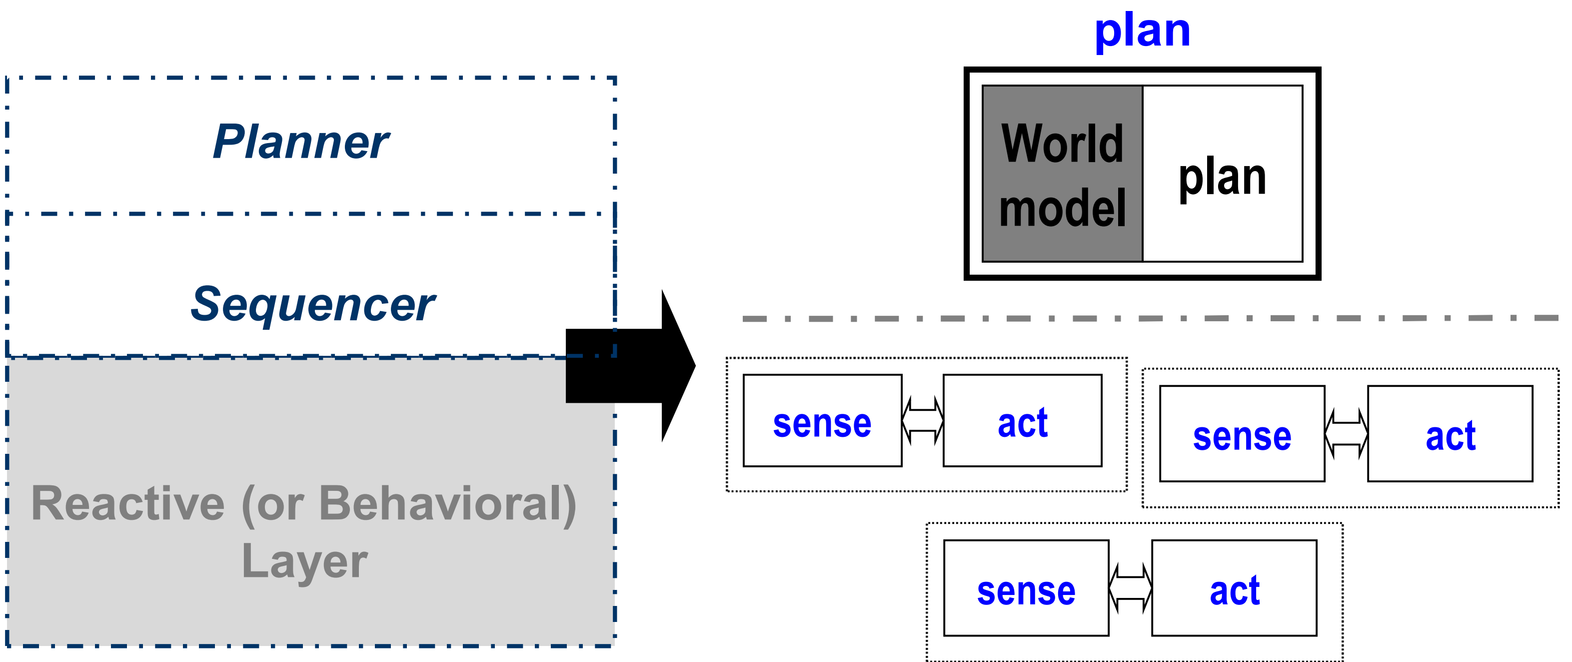

- PLAN, then **instantiate appropriate** SENSE-ACT behaviors, **until next step** in plan, …
- Technical PLAN execution but a significant programming aspect

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

LEARN

- Different types of learning, different things to learn
  - Is very diffuse
  - So, more permeates the architecture

- AI robots consist of four primitives

SENSE **ACT** **PLAN** **LEARN**

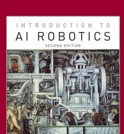

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

So Far…

- Biological organization suggests three layers of intelligence with distinctly different perception, knowledge, planning horizons, and time scales
- The AI Robotics field has converged on PLAN, then SENSE-ACT with LEARN as needed at different points
  – Technically this is SENSE-PLAN, SENSE-ACT but historically the sensing for planning just like the execution monitoring is lumped in “PLAN”
- *But still not as tangible as desired…*

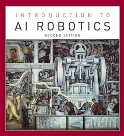

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

Consider as Functions

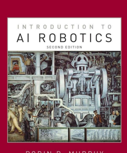

- Large scale factory automation (not factory robots) has 4 functions
  - generating
  - selecting
  - implementing
  - Monitoring
- May involve different contributions/combinations of primitives
  - Ex. monitoring requires sensing, understanding what was planned, and perhaps learning of what is normal and not normal
- Factory automation functions don't capture implementation but does seem to capture deliberation
  - what was implemented was usually a single, monolithic routine or control loop; natural intelligence usually have many low-level "routines" running concurrently
- Problem: still missing data/knowledge structures; not prescriptive: if I want to do X, then pick these components or have this set of layers…

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a PLAN got Bigger and Matches 3T

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

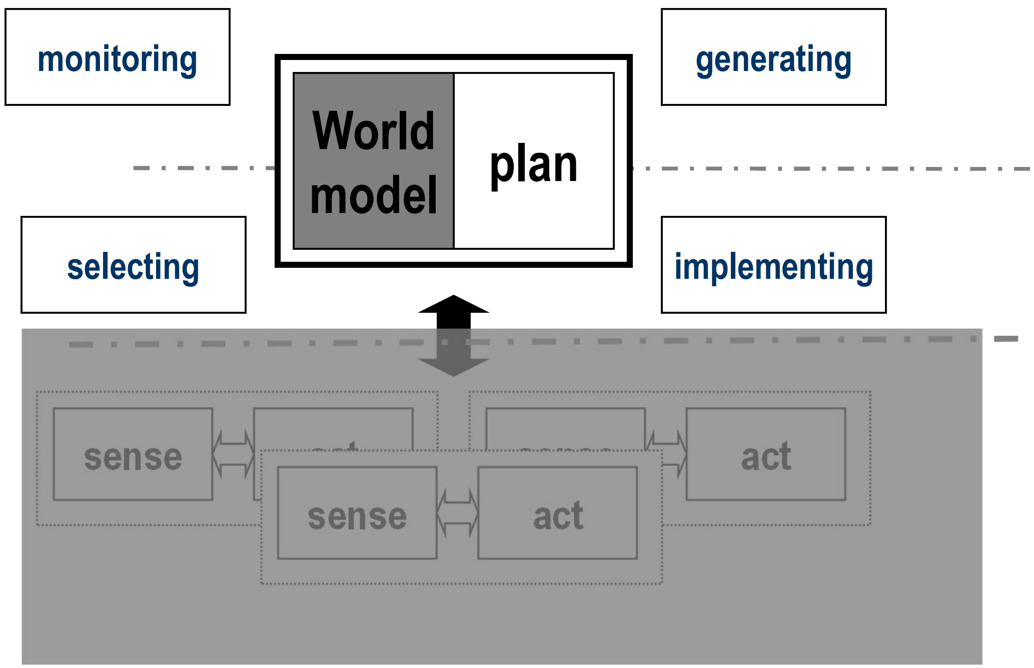

- Upper level (PLANNER)
  is *mission generation* & *monitoring*
  - Need Past, Present, Future

- Lower level (SEQUENCING)
  is *selection* of behaviors to accomplish task (*instantiation*) & local monitoring
  - Need Past, Present

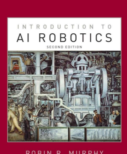

4a

Other Ramifications: Planning Horizon, Time Scales

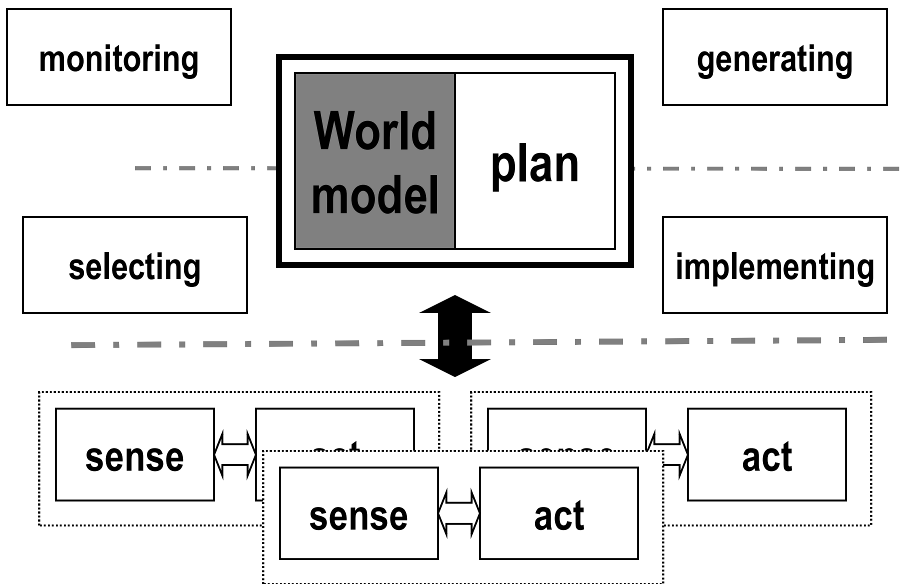

PRESENT+PAST+FUTURE, SLOW

PRESENT+PAST, FAST

PRESENT, VERY FAST

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Canonical Operational Architecture

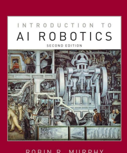

Interaction Layer

Deliberative Layer

Behavioral Layer

- Each layer has a different style of program organization
- Not concerned with details of interaction layer

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Programming Languages

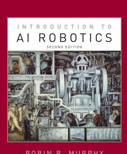

Layer	Programming languages
Interaction Layer	Procedural, Functional, Ontological languages such as OWL
Deliberative Layer	Functional languages such as Lisp
Behavioral Layer	Procedural languages such as C, C++, Java

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Advantages of programming in layers with different styles

- Decomposition of a complex system
  - Can use a separate processor(s) for each layer or behaviors
  - Can split between on-board, off-board
- Matching right tools and mindset for the task
  - Ex. C++ for behaviors, Lisp for planning
- Add to working, verified code

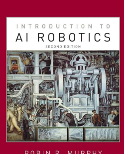

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

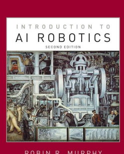

Can intelligence be added in layers? Like upgrading to “pro version” or downloading “apps” as needed?

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Can intelligence be added in layers? Like upgrading to "pro version" or downloading "apps" as needed?

- Yes, intelligence is organized in layers
- Yes, Can create libraries of equivalent ("logical") algorithms
- But as will be seen in later lectures…
  - Adding new behaviors or algorithms is non-trivial
    - Coordination functions at reactive layer impose certain assumptions, restrictions, and side effects
  - Adding another layer is non-trivial
    - Different attributes such as perception, models require significant design investment
    - "Hidden" coordination between layers

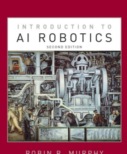

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

HOW MUCH ARTIFICIAL INTELLIGENCE DOES A ROBOT NEED?

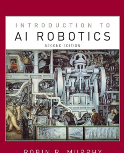

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a How Much Artificial Intelligence Does a Robot Need?

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- It depends…
- What functions does the robot need to do?
  - Generate? Monitor? Select? Implement? Execute behaviors? Learn?
- What planning horizon does the functions require?
  - Present, Present+Past, Present+Past+Future
- How fast do the algorithms have to update?
  - May have to use a closed world and guaranteed execution rates (control theory)
- What type of model does the robot need? Local? Global? Both?
  - Note: go with the minimum

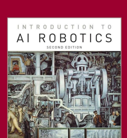

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

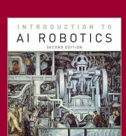

SUMMARY AND ADDITIONAL THOUGHTS

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Summary: Architectures

- An architecture is the Big Picture of how to program an intelligent robot
- "Architecture" refers to either
  - The *operational architecture* which described what the system does on an abstract level
  - The *system architecture* developed by a manufacturer (or research group) and may look like a data flow diagram
  - The *technical architecture*, specifies the actual techniques and code organization
- The canonical operational architecture for an AI robot consists of three layers which generally represent different programming styles or even implementation languages:
  - Behavioral
  - Deliberative
  - Interaction

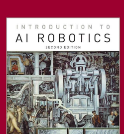

ROBIN R. MURPHY

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)

4a

Other Observations

Objectives
Review
Architecture
-types
Operational
-Biological
-Process
-Functions
-Ramifications
Layered Int.
How much?
Summary

- Current practice is good with deliberative functions operating on symbols, good with behaviors using direct perception
- Major barrier is going from sensory data to symbols: recognition and labeling as unique instances
  - May be able to see a coffee cup but not that it is my coffee cup and yours from the same dish collection is over there
- Major barrier in understanding human intention (which is often implied but never spoken); AI robots currently require explicit directions

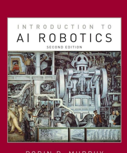

© 2019 Robin Murphy Introduction to AI Robotics 2nd Edition (MIT Press 2019)