



# **MIDCA, Version 1.5**

## User Manual and Tutorial for the Metacognitive Integrated Dual-Cycle Architecture

Technical Report Number COLAB<sup>2</sup>-TR-?

Dustin Dannenhauer<sup>1</sup>, Venkatsampath R. Gogineni, Sravya Kondrakunta,  
Anthony Mitchell and Michael T. Cox

**Wright State University**  
*College of Engineering & Computer Science*  
*Department of Computer Science and Engineering*  
*3640 Col. Glenn Hwy.*  
*Dayton, OH 45435*

**<sup>1</sup>Parallax Advanced Research**  
*4035 Col. Glenn Hwy.*  
*Beavercreek, OH 45431*

## Quick Start

This section describes how to install the core MIDCA system and shows how to run a very simple example that illustrates fundamental execution within the architecture. These instructions do not describe the full installation. See Installing MIDCA on page 25 instead. For more details on the following example, see Section 2.2 page 5.

### *Installing MIDCA*

1. Obtain a copy of MIDCA by cloning the repository or downloading the source directory (<https://github.com/COLAB2/midca>).
2. Make sure the name of the top-level folder is spelled exactly 'midca'. If midca has been downloaded as a zip file, you will have to rename it because it saves the folder as 'midca-master'.
3. Run the command `python setup.py install`.
  - If you plan to make changes to MIDCA, run `python setup.py develop` instead. Any changes you make will be immediately updated when you run MIDCA.
  - Note that NumPy (<http://www.numpy.org>) should be installed automatically when you run `python setup.py install` (or with the 'develop' option). If for any reason that fails, you will need to install the package yourself. You can check that NumPy is installed by running python and typing `import numpy`.

### *Example: The Running of the Chickens*

The `chicken_run.py` script runs a simple version of MIDCA in a domain where chickens can cross a road. The state of the environment (i.e., the world) consists of a left and right side of the road. The possible goal predicates are "onleft" and "onright" and the possible goal arguments are "clucky" and "lady cluck", which refer to two chickens. The goal predicates dictate which side of the road the selected argument (i.e., chicken) will try to reach. For example, if `onleft(clucky)` was selected as the goal, then the executed one-step plan for this goal would result in the chicken named clucky being on the left side of the road after the goal is completed.

Instructions for running the example.

1. Open up a command line in a terminal.
2. Go to `midca/midca/examples` directory.
  - `cd midca/midca/examples`
3. Run the `chicken_run.py` program.
  - `python chicken_run.py`

Press `?` and then `<enter>` for help. Or just pressing `<enter>` will advance one cognitive step. In MIDCA, these steps are called phases.

- You will see the following:

Next MIDCA command:

```
***** Starting Simulate Phase *****
```

Simulator: no actions selected yet by MIDCA.

Next MIDCA command:

4. Press `<enter>` to continue to the next phase.

- You will see the following

```
***** Starting Perceive Phase *****
```

World observed.

Next MIDCA command:

- Typing `show` displays Clucky on the right and lady clucky on the left.

Once at the Interpret phase

- You will see the following:

```
***** Starting Interpret Phase *****
```

Please input a goal if desired. Otherwise, press enter to continue

- Now enter a goal: `onleft(clucky)`.

5. You will see the plan and its remaining execution if you continue to press `<enter>` repeatedly.

6. To exit the script, type `q` and hit enter.

### *Runtime Commands*

The following interactive commands are useful when MIDCA pauses execution (see the section Runtime Commands on page 27 for a full list).

1. **help (or ‘?’)**. Displays the possible commands that can be given to MIDCA during runtime (i.e. the commands detailed here).
2. **memorydump**. Allows user to see memory variables. The user can either see all variables and their values or enter a single variable name and just see its value. If MIDCA has been running a long time, the output may take up more than the screen, therefore just looking for the variable can save space.
3. **printtrace**. Outputs a text representation of execution up to the last phase run.
4. **q**. Quits MIDCA.
5. **show**. Displays the current state of the world.
6. **skip (&optional x=1)**. Skips ahead x cycles or one full cycle if x is not given.

## Table of Contents

Quick Start .....	ii
Table of Contents .....	iv
Table of Figures .....	vi
List of Tables .....	vii
1. Introduction .....	1
1.1. MIDCA, Version 1.5 .....	1
1.2. Differences between MIDCA Versions 1.5 and 1.4 .....	1
1.3. How to Get Help .....	2
1.4. Outline .....	2
2. MIDCA Overview .....	2
2.1. The Ground Level, Object Level and Meta-level .....	2
2.2. The MIDCA Phase Structure .....	4
The cognitive cycle .....	4
The chicken_run example .....	5
The metacognitive cycle .....	5
3. Planning .....	7
3.1. The Blocksworld Domain .....	8
3.2. MIDCA's Planners .....	8
3.3. Examples of Planning Operators and Methods .....	9
Pyhop planner .....	9
JSHOP planner .....	10
3.4. How to Run an Example in MIDCA_1.5 .....	10
Blocksworld domain with Pyhop planner .....	10
Blocksworld domain with JSHOP Planner .....	11
4. Interpretation .....	12
4.1. D-Track Goal Generation .....	13
4.2. K-Track Goal Generation .....	14
5. The Implementation of MIDCA, Version 1.5 .....	17
5.1. Phases and Modules in MIDCA .....	17
5.2. Adding a New Module .....	19
5.3. The Goal Graph and Examples .....	19
Example with a plan that fails to achieve all goals .....	21
Examples of goal graph drawings .....	22
5.4. Goal Operations .....	23
5.5. How to Install and Run Version 1.5 of the MIDCA Architecture .....	25
Installing MIDCA .....	25

Using MIDCA with simulated worlds using a predicate representation .....	25
Running MIDCA .....	26
Understanding how MIDCA works from browsing the source code.....	27
6. Defining a New Domain .....	28
7. Logging and Debugging in MIDCA_1.5 .....	30
7.1. Logging .....	30
Initiating and disabling a log file .....	30
Location of log files .....	30
Files in the log folder: .....	30
7.2. Debugging .....	31
Debugging in MIDCA_1.5 .....	31
Debugging through log files .....	31
8. Multiagent Interaction.....	32
8.1. Goal delegation .....	32
9. Advanced Features .....	32
9.1. The MIDCA_1.5 Topic Interface to ROS.....	32
9.2. Using the MIDCA interface to ROS for the Baxter Humanoid Robot .....	34
Blocksworld domain for the Baxter .....	34
ROS topics .....	34
External steps (sensors and effectors).....	35
MIDCA setup: all steps from baxter_run.py.....	35
What MIDCA does while running.....	35
Camera calibration. ....	37
9.3. The Baxter and Gazebo .....	38
9.4. The MIDCA ROS API.....	43
The language. ....	43
Templates. ....	45
Instructions for running the MIDCA API.....	46
Acknowledgements.....	48
References.....	49
Appendix A: Frequently Asked Questions .....	51
MIDCA_1.5 General Questions .....	51
Questions regarding ROS and the Baxter robot .....	52
Appendix B: List of Modules .....	53

## Table of Figures

Figure 1. The Metacognitive Integrated Dual-Cycle Architecture (MIDCA) .....	3
Figure 2. An example of task decomposition. ....	7
Figure 3. A blocksworld problem to put block A on B. ....	8
Figure 4. A task decomposition tree for goal on(A, B) in blocksworld domain.....	8
Figure 5. MIDCA_1.4 output during the cogsci_demo example. ....	11
Figure 6. Depiction of the TF-Tree used in cycling through the 3 block configurations. ....	14
Figure 7. TF-Tree that generates goals to put out fires.....	14
Figure 8. Module execution sequences.....	18
Figure 9. Example showing interleaving of metacognitive modules.....	18
Figure 10. Complete sequence of all modules in the order they are executed.....	19
Figure 11. Goal graph for Example 1. ....	23
Figure 12. Goal graph for Example 2. ....	23
Figure 13. Example specification of the MIDCA_1.4 log file. ....	30
Figure 14. Code snippet of a log file.....	<b>Error! Bookmark not defined.</b>
Figure 15. Code snippet of memory access file.....	<b>Error! Bookmark not defined.</b>
Figure 16. Code snippet of the MIDCA_1.4 output file....	<b>Error! Bookmark not defined.</b>
Figure 17. Interfaces between MIDCA and the external world and between cognition and metacognition.....	32
Figure 18. Baxter Robot in gazebo, after executing the roslaunch command. ....	40
Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands. ....	41
Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.....	42

## List of Tables

Table 1. The arsonist explanation pattern .....	15
Table 2 Code snippet of module to phase assignment.....	17
Table 3. Phases of MIDCA interaction with goal graph.....	20
Table 4. Fundamental set of goal operations (adapted from Cox, Dannenhauer, & Kondrakunta, 2017) .....	24

## 1. Introduction

The *Metacognitive Integrated Dual-Cycle Architecture (MIDCA)* (Cox et al., 2016; Paisner, Cox, Maynard & Perlis, 2014), is a cognitive architecture that models both cognition and metacognition for intelligent agents. It consists of “action-perception” cycles at both the cognitive level and the metacognitive level. In general, a cycle performs problem-solving to achieve its goals and tries to comprehend the resulting actions and those of other agents. Problem solving consists of intention, planning, and action execution phases, whereas comprehension consists of perception, interpretation, and goal evaluation.

The MIDCA system is meant to serve as a basis for intelligent control of a software agent or physical platform in a complex environment. MIDCA can be applied to several different environments (i.e., domains) and can be used for various decision tasks. People interested in using MIDCA are likely to fall into two categories.

1. People who want to build an agent to carry out tasks in an environment.
2. People who want to build upon MIDCA as they conduct new research on cognitive architectures and artificial intelligence.

Regardless of category, if this is your first time using MIDCA, begin with the Quick Start section preceding the Table of Contents of this manual (pp. ii-iii).

### 1.1. MIDCA, Version 1.5

MIDCA is a computational theory of cognition and metacognition implemented in an open-source software system written mainly in python. MIDCA is highly extensible and can be easily tailored for custom use as will be explained in this manual. Version 1.5 of MIDCA (i.e., MIDCA\_1.5) is publicly available on the GitHub code repository and has several new features including the following.

1. Multiple automated planners;
2. Multiple goal operations;
3. Expectations and explanation module at the meta-level;
4. Multiple domains;
5. Multiagent support for multiple concurrent MIDCA executions.

Project website: <http://www.midca-arch.org>

GitHub repository: <https://github.com/COLAB2/midca>

GitHub wiki: <https://github.com/COLAB2/midca/wiki/MIDCA-v1.5-Home>

### 1.2. Differences between MIDCA Versions 1.5 and 1.4

For the user manual and tutorial of MIDCA\_1.4 see Dannenhauer, Schmitz, Eyorokon, Gogineni, Kondrakunta, Williams & Cox (2019). However, note that many implementation features have changed, and numerous new ones exist. For example, the implementation can now operate as a

multiagent system with multiple MIDCA agents acting independently in a shared environment and or coordinating together. To implement this feature, we have changed the mechanism by which goals are input. [Explain requested predicate]

[Elab]

### 1.3. How to Get Help

Questions and comments with respect to MIDCA\_1.5 are welcome. Please email your concerns to [wsri-midca-help@wright.edu](mailto:wsri-midca-help@wright.edu). Resources and videos are also available from the site of the MIDCA workshops. See [www.midca-arch.org/workshops](http://www.midca-arch.org/workshops).

### 1.4. Outline

The remainder of this document is organized as follows. First, we present an overview of the MIDCA architecture. In this section, we distinguish metacognition from cognition, and then we discuss the system's organization in terms of information processing phases. Section 5 describes the implementation of phases as python modules and details how to add custom modules. This is followed by a description of the goal graph data structure along with an example. The next section describes how to define a new domain, and we discuss the fundamental notion of goal operations in MIDCA that distinguishes it from similar cognitive systems. This is followed by a section that explains logging and debugging within the architecture. A subsequent section on advanced features includes details for using MIDCA\_1.5 with the *Robot Operating System (ROS)* framework. Finally, this document concludes with references and two appendices. Appendix A contains frequently asked questions and Appendix B enumerates the full list of the basic system modules.

## 2. MIDCA Overview

MIDCA's functional structure is depicted in Figure 1. This illustration shows the architecture's two reasoning cycles and their constituents. MIDCA's cycles are organized into functional phases at different levels, and the phases are implemented by python modules. This section describes the former, and the next section deals with the latter.

A reasoning cycle starts with input of the domain (either the world domain or cognition itself) via the Perceive (or Monitor) phase. The Interpret phase takes the extracted predicate relations and the expectations in memory in order to determine whether the agent is making sufficient progress. At this point, if the domain presents problems or opportunities for the agent, then a new goal is created. The Evaluate phase incorporates the concepts inferred from Interpret and notes whether existing goals are achieved. In a cycle, Intend commits to a current goal from those available. The Plan phase then generates a sequence of actions. The Act (or Control) phase executes the plan one step at a time to change the domain through the effects of the planned actions. In the next cycle, MIDCA will then use expectations from a model of these actions to evaluate the execution of the plan. The metacognitive cycle (in Figure 1, upper cycle in blue) is analogous to the cognitive cycle (lower cycle in orange).

### 2.1. The Ground Level, Object Level and Meta-level

One way of organizing MIDCA’s capabilities is to describe them in terms of three levels: (1) the ground level; (2) the object level; and (3) the meta-level (Cox & Raja, 2011). The *ground level* contains the Perceive phase, Act phase, and the world simulator (or when using MIDCA with ROS or another API, the world itself) and focuses on interactions with MIDCA’s environment. This level contains the information needed for the agent to acquire observations from the world (Perceive) and perform changes to the world (Act). Perceive and Act phases may need to be customized based on the environment or world the agent is operating within.




Figure 1. The Metacognitive Integrated Dual-Cycle Architecture (MIDCA)

The *object level*, which is sometimes referred to as the cognitive level, contains phases that use information directly related to the world. These are non-metacognitive phases. The Plan, Intend, Evaluate, and Interpret phases are all part of the object level, since they are each concerned with world-related information. The Plan phase determines what actions need to be executed in the world to achieve the goal. The Intend phase selects which goals to commit to; these goals are future world states the agent would like to achieve. The Evaluate phase assesses the completion of active goals. The Interpret phase is concerned with detecting anomalies, explaining failures (in the

world), and generating new goals. Within the phases of the object level, activity is generally related to the environment. Goals are concrete or abstract world states to be achieved.

The *meta-level* is concerned with monitoring and controlling the object level. In the same way the object level perceives the environment and acts to change it, so does the metacognitive level monitor the object level and act, via the meta-level control phase, to change the object level. The structure of the phases at the meta-level mirror the object level, yet the focus of the phases is different. The source of input to the meta-level (through the Monitor phase) is a trace of behavior of the object level. Note that the meta-level is still in early development and so the implementation is preliminary.

## 2.2. The MIDCA Phase Structure

The fundamental computational unit for reasoning cycles in MIDCA is an information processing phase.

### *The cognitive cycle*

**Perceive:** The objective of the Perceive phase is to obtain knowledge of the world. If using a simulator, this may be to simply query the simulator for a current state and then store that state in MIDCA's memory. If not using a simulator, other capabilities may be required such as instructing a physical camera to capture photos or videos of the agent's environment (see Section 9.2 for an example of this). In general, all observations should happen in the Perceive phase.

**Interpret:** The Interpret phase directly follows Perceive and is meant to provide a number of interpretation-related capabilities. These capabilities derive useful information from an initial world state which is obtained from the Perceive phase. These capabilities are likely to include but are not limited to: inferring new facts (e.g. by running a reasoner over the initial world state), noting any anomalies or unexpected conditions in the world state, explaining why the world state is the way that it is or why an anomaly has occurred, and generating new goals for MIDCA to pursue. Prior research on goal formulation has shown that certain states may prompt the agent to adopt new goals (Paisner, Cox, Maynard, & Perlis, 2014). Section 4 discusses this phase in detail.

**Evaluate:** Evaluate follows from the Interpret phase and the primary focus is to review and update the agent's current goals. Most importantly, if the agent has achieved a goal, then the goal should be dropped in the Evaluate phase. In general, any kind of behavior involving evaluation of the agent's knowledge should happen in this phase.

**Intend:** The Intend phase determines which goals to pursue and follows after the Evaluate phase. New goals are typically generated and inserted into MIDCA's goal graph (the goal graph is the system's internal structure for its goals and plans; see Section 5.3). During the Intend phase, MIDCA may commit to 1 or more goals. There are two selection strategies used to select one goal from all available goals. The first method is First In First Out (FIFO) and the second is a smart selection criterion to select a goal. For more information on the latter, see Kondrakunta (2017).

**Plan:** The primary purpose of the Plan phase is to generate the sequence of actions (i.e., a plan) to be executed by MIDCA to achieve the current goal set. As discussed in section 3.2, multiple planners exist in MIDCA that can be used to perform this function. The current goals and state of the world stored in MIDCA's memory are used as input to the planner linked to this phase. Because

planning often requires specific knowledge (e.g., action models that represent the steps composing a plan or domain-specific heuristics), there are special files that contains planner specific information (see Section 3.3).

**Act:** The Act phase carries out the next action by MIDCA, usually the next action of the current plan. If no plan exists, then no actions are performed. If an action is chosen, it is sent to the world simulator, which uses it to compute the next world state. An example of such an action might be `crossleft()` in the chicken domain below. Any behavior where MIDCA executes an action to change the world happens in this phase. For an example of the Act phase operating on a robot, see Section 9.2.

### *The chicken\_run example*

`Chicken_run.py` is a cognition-only example in MIDCA which demonstrate the basics of MIDCA 1.5. This section will go through the cognitive-level phases of MIDCA in the example and explain each phase. See also the Quick Start on pages ii-iii.

**Perceive:** In the beginning of this example, Perceive does little more than copy the state of the world (i.e., the chicken and the road) and put it back in memory.

**Interpret:** This phase interacts with the user. It waits for a goal predicate to be given. For example, the user could give the goal “`onleft(clucky)`”. This establishes the goal agenda for MIDCA as the single goal for the chicken to cross to the left side of the road.

**Evaluate:** This phase checks to see if the current goal is achieved in the state. Because the goal in the agenda has not yet been committed to, this phase does nothing.

**Intend:** This phase now selects the most recent goal from the agenda and makes it the current goal, thereby committing to achieve this goal

**Plan:** This phase creates a plan with actions that will allow a goal to be accomplished. For example, if the user wants the chicken to cross to the left side of the road, MIDCA might create a plan with a single action “`crossleft`”.

**Act:** This phase takes an action from the plan and attempts to execute it.

**Simulation:** The action is now simulated, and the results of the action changes the state of the world if the action was valid.

**Perceive:** This phase takes the new state from the simulator and updates the memory. For example, if the chicken has crossed to the left side of the road it will appear in the world on the left side of the road. This perceive marks the beginning of the second cycle through the cognitive level.

**Interpret:** Nothing occurs during Interpret at this point in the example. If something unexpected had happened (i.e., the “`crossleft`” action had failed), a new goal may have been generated.

**Evaluate:** This phase notices that the state of the world entails the goal state and removes the “`onleft(clucky)`” goal from the agenda. The agenda is now empty.

### *The metacognitive cycle*

Metacognitive phases are similar to their cognitive-level counterparts. The primary differences are between Perceive (cognitive) and Monitor (metacognitive), and between Act (cognitive) and Control (metacognitive). Perceive and Monitor are both concerned with obtaining the state of the “world.” However, the world for the Monitor phase is not the ground-level world, but instead it is the domain of cognition. Thus, Monitor obtains information pertaining to activity at the cognitive level. Likewise, the cognitive level Act phase is concerned with actions that change the ground-level environment while the meta-level Control phase is concerned with modifying some part of cognition at the object level. This could include changing goals stored in MIDCA’s memory or removing or adding a module from a particular cognitive phase. Below is a brief summary of each phase of the metacognitive layer.

**Monitor:** The Monitor phase obtains the most recent s of cognition from memory. The cognitive trace is constructed by recording what each module in each phase takes as an input and produces as an output. Specifically, the cognitive trace is composed of an ordered dictionary that can be indexed via the phase and cycle (cycle refers to loop iterations). The inputs and outputs are changes in the data stored in MIDCA’s memory. For example, an input to the module performing planning during the Plan phase would be the world state and current goal of MIDCA (again these are stored in MIDCA’s memory) and the output would be a plan (which is also stored in MIDCA’s memory).

**Interpret:** The Interpret phase is responsible for detecting cognitive-level behavior that may warrant a metacognitive response. Currently, Interpret has modules for detecting discrepancies, explaining the cause, and generating an appropriate goal. However, only discrepancy detection is implemented; explanation and goal generation are simple, hand-engineered approaches. Currently there are two ground-level domain-independent expectations that are used in the discrepancy detection approach found in the *MRSimpleDetect* module.

**Evaluate:** Since the current implementation of the metacognitive layer ensures that all metacognitive goals are achieved in one metacognitive cycle, Evaluate does not have an implementation. Evaluate will be responsible for monitoring progress on metacognitive goals that take more than a cycle to achieve. The metacognitive layer runs for only a single loop between each cognitive module; see Figure 8, Figure 13, and Figure 10).

**Intend:** The Intend phase is meant to decide which goal to pursue. The current implementation is straightforward, and any goals that are pending will be selected.

**Plan:** The Plan phase decides what actions to take to achieve a metacognitive goal (which is different than a cognitive goal, since a metacognitive goal is concerned with a future ‘mental state’ in cognition; goals at the metacognitive layer are another area of future work). Currently, the Control phase can only execute three actions: *RemoveModule*, *AddModule*, and *TransformGoal*. The *RemoveModule* and *AddModule* actions are responsible for removing and adding cognitive level modules, respectively. The *TransformGoal* action is responsible for performing a transformation on the goal. The code to carry-out these actions are found in the Control Phase.

**Control:** The Control phase contains the information for carrying out the actions described in Plan. Currently, the Control phase carries out an entire plan as opposed to a single action like the Act (cognitive) phase.

We now turn our attention to a more detailed discussion of the components that make up MIDCA's phases. Most important are the planning and interpretation components. We examine them in turn.

### 3. Planning

In the literature, planning is a means for calculating a sequence of steps that achieve a goal or perform a task, given the current configuration of the world. Steps are represented as actions that if executed will change the world; whereas the configuration of the world is represented as a state consisting of the relationships between objects that are true at some point in time. These relationships are often represented as a set of logical predicates along with their arguments. For example,  $S = \{p(a, b)\}$  is a state where the predicate  $p$  holds between the objects  $a$  and  $b$ . State-space planners seek to achieve a goal state given an initial state; whereas *hierarchical task network (HTN)* planners seek to refine an abstract task into a set of subtasks given an initial state. Both produce as output an ordered sequence of actions we call a plan. Given that MIDCA mainly uses HTN planners, we will assume this for the subsequent discussion.

We denote an action  $\alpha = (name(\alpha), precond(\alpha), effects(\alpha))$  that accomplishes a primitive task  $t$  in state  $s$  if  $name(\alpha) = t$  and  $precond(\alpha)$  is satisfied in  $s$ . The preconditions are satisfied if all positive conditions are elements of  $s$  and all negative conditions are not elements of  $s$ . A method is a tuple  $m = (name(m), task(m), precond(m), subtasks(m))$  in which  $name(m)$  is the name of the method;  $task(m)$  is a non-primitive task; and  $precond(m)$  is a set of literals called the method's preconditions.  $Precond(m)$  specifies what conditions the current state must satisfy in order for  $m$  to be applied, and  $subtasks(m)$  specifies one or more subtasks to perform in order to accomplish  $task(m)$ .




Figure 2. An example of task decomposition. Method  $m_2$  decomposes the non-primitive task  $t_{12}$  into  $t_{21}$  and  $t_{22}$ . Operators  $op_1$  and  $op_2$  accomplish each of these primitive subtasks respectively.

An HTN planning problem is a tuple  $P = (s, T, D)$ . It takes the initial state,  $s$ , which is a symbolic representation of the state of world, and a set of tasks  $T = \langle t_1, \dots, t_k \rangle$  to be accomplished. It also takes a knowledge base,  $D$ , including operators and methods. A plan  $\pi = \langle \alpha_1, \dots, \alpha_n \rangle$  is a solution (i.e., a task decomposition as in Figure 2) for a planning problem to accomplish  $T$ . This

means that there is a way to decompose  $T$  into  $\pi$  in such a way that  $\pi$  is executable in  $s$  and will transform the start state into the goal state once executed.

### 3.1. The Blocksworld Domain

Consider an example in the blocksworld domain with a goal  $on(A, B)$ . This goal is mapped to the root task *move-blocks* in Pyhop planner. Assume the initial state in panel (a) of Figure 3 with the three blocks A, B, and C on the table. Given that both blocks are clear, the planner generates a simple two-step plan  $\pi = pickup(A), stack(A, B)$ . Panels (b) and (c) show the execution of the plan steps. Figure 4 shows the task decomposition tree for this task.




Figure 3. A blocksworld problem to put block A on B. The first panel shows the initial state, and the remaining panels show the incremental execution of the plan steps that solve the problem.




Figure 4. A task decomposition tree for goal  $on(A, B)$  in blocksworld domain. The leaves are the operators.

### 3.2. MIDCA's Planners

Multiple planners exist that are available in MIDCA including:

- The java-based *Hierarchical Task Network (HTN)* planner JSHOP2 (Nau, et al., 2003; Nau, Muñoz-Avila, Cao, Lotem, & Mitchell, 2001) that is run as an external java program;
- A python-based version of JSHOP named Pyhop;
- An asynchronous version of Pyhop that is more commonly used on robotics platforms (we have used it on a Baxter robot);

- Fast downward (Helmert, 2006; 2009), a classical planning system that supports *Plan Domain Definition Language (PDDL)* (Edelkamp & Hoffmann, 2004; Fox & Long, 2003) action models;
- A state-space heuristic search planner that can be modified to be used as a BFS, DFS, or A\* planner.

SHOP (Nau, et al., 2003; Nau, Cao, Lotem & Muñoz-Avila, 1999) is an HTN planning algorithm that generates plans for tasks rather than goals. A procedure in MIDCA maps the goals to tasks for the SHOP Planner. SHOP creates plans by recursively decomposing tasks into smaller subtasks until only the primitive tasks are left which can be accomplished directly. SHOP uses methods and operators. A method specifies a way to decompose a non-primitive task into a set of subtasks while an operator specifies a way to perform a primitive task.

Pyhop is a SHOP-like planner ([www.cs.umd.edu/projects/shop](http://www.cs.umd.edu/projects/shop)) written in Python; whereas, JSHOP is the Java implementation of the SHOP planner. For both planners, we need to define a domain file and a problem file for the JSHOP planner. The domain file contains the definition of operators, methods and axioms. In the problem file, initial state and initial task list (goals) are specified.

### 3.3. Examples of Planning Operators and Methods

#### *Pyhop planner*

For each domain, there is an operator class and method class. For example, the operators and methods for blocksworld domain are in [midca/midca/domains/blocksworld/plan/](http://midca/midca/domains/blocksworld/plan/).

The code below is an example of an operator in Pyhop. This is an operator to *pickup* a block from the table. The preconditions of this operator are (1) the position of the block is on the table, (2) the block is clear, and (3) the agent is not holding anything. The results of this operator are (1) the block is in the hand, (2) the block is not clear, and (3) the agent is holding the block.

```

1. def pickup(state,b):
2.     if state.pos[b] == 'table' and state.clear[b] == True and state.holding == False:
3.         state.pos[b] = 'hand'
4.         state.clear[b] = False
5.         state.holding = b
6.         return state
7.     else: return False

```

The code below is an example of a method in Pyhop. This method decomposes the high level task *move\_one* to a set of subtasks to get a block  $b_1$  and put it at destination. There are two different ways to decompose this task. If the position of the block  $b_1$  is in arm, then this task is decomposed to  $\langle (put, b_1, dest) \rangle$ , otherwise it is decomposed to the set of subtasks  $\langle (get, b_1), (put, b_1, dest) \rangle$ .

```

1. pyhop.declare_methods('move_one', move1)
2.
3. def move1(state,b1,dest):
4.     if state.pos[b1] == "in-arm":
5.         return [('put', b1,dest)]

```

```
6.     else:
7.         return [('get', b1), ('put', b1,dest)]
```

### *JSHOP planner*

The domain and state files for the blocksworld domain are located at the URL [midca/midca/domains/jshop\\_domains/blocks\\_world](http://midca/midca/domains/jshop_domains/blocks_world).

The code below is an example of an operator in JSHOP planner. The first line is the name of the operator. The second line specifies the preconditions that need to be met for that operator to be applied, and the last line is the results of performing that operator.

```
1. (:operator (!pickup ?a)
2.   ((clear ?a) (on-table ?a))
3.   ((holding ?a)))
```

The code below is an example of a method in JSHOP planner. The first line is the name of the method, the second line (which is empty here) is preconditions for this method, and the last line is a set of subtasks.

```
1. (:method (achieve-goals ?goals) ()
2.   ((assert-goals ?goals nil) (move-block nil)))
```

## 3.4. How to Run an Example in MIDCA\_1.5

### *Blocksworld domain with Pyhop planner*

You can run an example script named *cogsci\_demo.py* from [midca/midca/examples/](http://midca/midca/examples/) folder (see Figure 5), or you can write your own script.

(1) The domain file and state files need to be defined for the MIDCA\_1.5 simulator.

```
1. DOMAIN_FILE = DOMAIN_ROOT + "domains/arsonist.sim"
2. STATE_FILE = DOMAIN_ROOT + "states/defstate.sim"
```

For new domains:

[domains/blocksworld](http://midca/midca/blocksworld) contains the .sim files that contain the logic for states (types, predicates, and operators) that MIDCA's simulator will use. You can see an example of these files for different domains in [midca/midca/domains/](http://midca/midca/domains/). The operators' signature in this domain file needs to be identical to the operators' definition in the Pyhop planner. The state file specifies the initial state.

(2) Modify the path for the methods and operators for Pyhop planner. These two files are the definition of operators and methods for the planner.

```
1. DECLARE_METHODS_FUNC = methods.declare_methods
2. DECLARE_OPERATORS_FUNC = operators.declare_ops
```

(3) Modify the utility file for domain specific utility functions. *blocksworld/util.py* is a file that contains any domain specific utility functions for the *blocksworld* domain. The Pyhop planner uses its own state representation which requires translation to and from MIDCA states (the state


specified in `blocksworld/domains/*.sim` file). For *blocksworld* this translation happens in `util.py`, specifically the functions `pyhop_state_from_world()` and `pyhop_tasks_from_goals()`. Many other useful utility functions are located here, including how to draw an ascii representation of the world state in a terminal (optional, but useful).

(4) Modify the file that creates *the MIDCA object*. This object is an instance of the `PhaseManager` class created to insert, append and run modules of `MIDCA_1.5`. Setting the planner to use the `Pyhop` planner while passing the `util` file along with the methods and operators as parameters, is as follows.

```

1. myMidca.append_module("Plan",
2.                       planning.PyHopPlanner(util.pyhop_state_from_world,
3.                                             util.pyhop_tasks_from_goals,
4.                                             DECLARE_METHODS_FUNC,
5.                                             DECLARE_OPERATORS_FUNC))

```



```

examples — Python cogsci_demo.py — 80x24
Simulator: no actions selected yet by MIDCA.
  / \
 /   \
/     \
|   D_   \
|-----|
|   B_   |
|-----|
|-----|
|   A_   |   |   C_   |
|-----|   |-----|
-----

Next MIDCA command:

***** Starting Perceive Phase *****

World observed.
Next MIDCA command:

***** Starting Interpret Phase *****

TF-Tree goal generated: Goal(C_, B_, predicate: on)
Next MIDCA command: █

```

Figure 5. MIDCA\_1.5 output during the `cogsci_demo` example.

### *Blocksworld domain with JSHOP Planner*

You can run an example script named `simple_run_jshop.py` from `midca/midca/examples/` folder, or you can write your own script.

(1) The domain file and state files need to be defined for the MIDCA simulator.

```

1. DOMAIN_FILE = DOMAIN_ROOT + "domains/arsonist.sim"

```

```
2. STATE_FILE = DOMAIN_ROOT + "states/defstate_jshop.sim"
```

(2) The domain file and state file path need to be defined for the JSHOP planner.

```
1. JSHOP_DOMAIN_FILE = MIDCA_ROOT + "domains/jshop_domains/blocks_world/blocksworld.shp"  
2. JSHOP_STATE_FILE = MIDCA_ROOT + "domains/jshop_domains/blocks_world/bw_ran_problems_5.shp"
```

(3) Set the util file. *blocksworld/util.py* is a file that contains any domain-specific utility functions. The JSHOP Planner uses its own state representation which requires translation to and from MIDCA states. For *blocksworld*, this translation happens in *util.py*, specifically the functions *jshop\_state\_from\_world()* and *jshop\_tasks\_from\_goals()*. Many other useful utility functions are located here, including how to draw an ASCII representation of the world state in a terminal (optional, but useful).

(4) Set the planner to use JSHOP planner and pass the util functions and JSHOP domain file and state file as parameters.

```
1. myMidca.append_module("Plan",  
2.                       planning.JSHOPPlanner(util.jshop_state_from_world,  
3.                                             util.jshop_tasks_from_goals,  
4.                                             JSHOP_DOMAIN_FILE,  
5.                                             JSHOP_STATE_FILE))
```

## 4. Interpretation

Perception takes percepts from the environment as input and produces state predicates as output. Interpretation makes sense out of the state by matching it with expectations from the intended plan and knowledge structures in memory and creating a model of the current sequence of events and state of the world. In particular, if MIDCA's expectations fail to match the observation from perceive, it needs to explain the difference and formulate a new goal if necessary. The process it uses to perform this function is called *Goal-Driven Autonomy (GDA)* (Cox, 2007; Klenk, Molineaux, & Aha, 2013; Munoz-Avila, Jaidee, Aha, Carter, 2010).

The Interpret phase has been at the core of our research efforts. It is implemented as a GDA procedure that uses both a bottom-up, data-driven track and a top-down, knowledge rich track (Cox, Maynard, Paisner, Perlis, & Oates, 2013). MIDCA\_1.5 uses both of these processes to analyze the current world state and determine which, if any, new goals it should attempt to pursue. The details of this process are described below. In the *cogsci\_demo.py* example, this is the phase in which MIDCA notices an anomaly in the blocksworld (e.g., a block on fire) and decides what to do about it.

The Interpret phase of MIDCA is implemented by two GDA processes that combine to generate new goals based on the features of the world the agent observes. We call these processes the *D-track*, which is a data driven, bottom-up approach, and the *K-track*, which is knowledge rich and top-down (Paisner, Cox, Maynard, & Perlis, 2014). A statistical anomaly detector constitutes the first step of the D-track, a neural network identifies low-level causal attributes of detected anomalies, and a goal classifier, trained using methods from machine learning, formulates goals. The K-track is implemented as a case-based explanation process.

The representations for expectations significantly differ between the two tracks. K-track expectations come from explicit knowledge structures such as action models used for planning and ontological conceptual categories used for interpretation. Predicted effects in the former and attribute constraints in the latter constitute expectations. By contrast, D-track expectations are implicit. Here the implied expectation is that the probabilistic distribution from which observations are sampled will remain the same. When the difference between expected and perceived distribution is statistically significant, an expectation violation is raised.

#### 4.1. D-Track Goal Generation

The D-track interpretation procedure uses a novel approach for noting anomalies. We apply the statistical distance metric called the A-distance to streams of predicate counts in the perceptual input (Cox, Oates, Paisner, & Perlis, 2012), yielding a measurement of how the distributions of predicates differ from a base state. This enables MIDCA to detect regions in which statistical distributions of predicates differ from previously observed input. MIDCA's implicit assumption is that where change occurs problems may exist.

When a change is detected, its severity and type can be determined by reference to a neural network in which nodes represent categories of normal and anomalous states. This network is generated dynamically with the growing neural gas algorithm (Paisner, Perlis, & Cox, 2013) as the D-track processes perceptual input. This process leverages the results of analysis with A-distance to generate anomaly prototypes, each of which represents the typical member of a set of similar anomalies the system has encountered. When a new state is tagged as anomalous by A-distance, the GNG net associates it with one of these groups and outputs the magnitude, predicate type, and valence of the anomaly.

Goal generation is achieved in MIDCA\_1.5 using TF-Trees (Maynard, Cox, Paisner, & Perlis, 2013), machine-learning classification structures that combine two algorithms which work over the predicate representation of the blocksworld domain. The first of these algorithms is Tilde (Blockeel, & De Raedt, 1997), which is itself a generalization of the standard C4.5 decision tree algorithm. The second algorithm is FOIL (Quinlan, 1990), an algorithm which, given a set of examples in predicate representation reflecting some concept, induces a rule consisting of conjunctions of predicates that identify the concept. Given a world state, a TF-Tree first uses Tilde to classify the state into one of a set of scenarios. Each scenario is then associated with a rule generated by FOIL. Once that rule is obtained, groundings of the arguments of the predicates in that rule are permuted until either a grounding that satisfies the rule is found (in which case a goal is generated) or until all permutations have been eliminated as possibilities (in which case no goal is generated). The structure of a TF-Tree is a tree where in internal nodes are produced by Tilde and leaf nodes are rules produced by FOIL. Figure 6 depicts the structure of the TF-Tree MIDCA uses in cycling through the 3 block arrangements.




Figure 6. Depiction of the TF-Tree used in cycling through the 3 block configurations.

For example given the middle state of Figure 5, triangle D is clear, it is on the table, and the table is a table. Thus, we take the right branch labeled “yes.” Now triangle D is also a triangle, so again we take the “yes” branch to arrive at the right-most leaf of the tree. The leaf rule then binds the variable Y to the clear square C, and the resulting goal is to have triangle D on square C.

The construction of a TF-Tree requires a training corpus consisting of world states and associated correct and incorrect goals. In simple worlds TF-Trees can be constructed which have perfect or near perfect accuracy using small training corpora. Corpora have to be constructed by humans, as labels need to be attached to potential goals in various world states. For simple worlds corpus construction does not carry an excessive burden, but that burden increases with the complexity of the world. Because a TF-Tree is a static structure trained on the specifics of the world, when the world changes, even in minor ways, a new training corpus has to be constructed and a new TF-Tree trained. However, the corpus to create a simple tree for reacting to fires (see Figure 7) consisted of only four examples.




Figure 7. TF-Tree that generates goals to put out fires.

## 4.2. K-Track Goal Generation

The K-track GDA procedure uses the XPLAIN system (Cox & Burstein, 2008). XPLAIN is built on top of the Meta-AQUA introspective story understanding system (Cox and Ram 1999) and is used in MIDCA to detect and explain problems in the input perceptual representations. The system’s interpretation task is to “understand” input by building causal explanatory graphs that link subgraph representations in a way that minimizes the number of connected components. XPLAIN uses a multistrategy approach to this problem. Thus, the top-level goal is to choose a comprehension method (e.g., script processing, case-based reasoning, or explanation generation) by which it can understand an input. When an anomalous or otherwise interesting input is detected, the system builds an explanation of the event, incorporating it into the preexisting model of the story. XPLAIN uses case-based knowledge representations implemented as frames tied together by explanation-patterns (Cox & Ram, 1999) that represent general causal structures.

XPLAIN relies on general domain knowledge, a case library of prior plan schemas and a set of general explanation patterns that are used to characterize useful explanations involving that background knowledge. These knowledge structures are stored in a (currently) separate memory sub-system and communicated through standard socket connections to the rest of MIDCA\_1.5. XPLAIN uses an interest-driven, variable depth, interpretation process that controls the amount of computational resources applied to the comprehension task. For example, an assertion that triangle-D is picked up generates no interest, because it represents normal actions that an agent does on a regular basis. But XPLAIN classifies block-A burning to be a violent action and, thus according to its interest criterion, interesting. It explains the action by hypothesizing that the burning was caused by an arsonist. An abstract explanation pattern (see Table 1), or XP, retrieved from memory instantiates this explanation, and the system incorporates it into the current model of the actions in the input “story” and passes it as output to MIDCA.

The ARSONIST-XP asserts that the lighting of the block caused heat that together with oxygen and fuel (the block itself) caused the block to burn. The arsonist lit the block because he wanted the block’s burning state that resulted from the burning. The objective is to counter a vulnerable antecedent of the XP. In this case the deepest antecedent is the variable binding =l-o or the light-object action. This can be blocked by either removing the actor or removing the ignition-device. The choice is the actor, and a goal to apprehend the arsonist is thereby generated.

Table 1. The arsonist explanation pattern

```
(define-frame ARSONIST-XP
  (actor (criminal-volitional-agent))
  (object (physical-object))
  (antecedent (ignition-xp
    (actor =actor)
    (object =object)
    (ante (light-object =l-o
      (actor =actor)
      (instrumental-object
        (ignition-device))))
    (conseq =heat)))
  (consequent (forced-by-states
    (object =object)
    (heat =heat)
    (conseq (burns =b
      (object =object)))))
  (heat (temperature (domain =object)
    (co-domain very-hot.0)))
  (role (actor (domain =ante)
    (co-domain =actor)))
  (explains =role)
  (pre-xp-nodes(=actor =consequent =object =role))
  (internal-nodes nil.0)
  (xp-asserted-nodes (=antecedent))
  (link1 (results
    (domain =antecedent))
    (co-domain =consequent)))
  (link2 (xp-instrumental-scene->actor
    (actor =actor)
    (action =l-o)
    (main-action =b)
    (role =role))))
```



## 5. The Implementation of MIDCA, Version 1.5

A series of python modules organized into phases and centered about a core memory structure implements Version 1.5 of MIDCA. We review MIDCA's phases and include implementation details. Then we describe how to add a new module to the MIDCA system. Following this description, we explain MIDCA's goal graph and provide examples of its operation and then briefly describe the concept of goal operations. Finally, we specify how to install and run the system code.

### 5.1. Phases and Modules in MIDCA.

Modules make up the phases discussed in the previous section. More than one module can make up a phase, in which case the order in which modules will run needs to be specified. Table 2 is an example code snippet of initializing the modules and phases of MIDCA.

Table 2 Code snippet of module to phase assignment.

```
1. myMidca.append_module("Perceive", perceive.PerfectObserver())
2. myMidca.append_module("Interpret", note.SimpleDetect())
3. myMidca.append_module("Interpret", assess.SimpleExplain())
4. myMidca.append_module("Interpret", guide.SimpleGoalGen())
5. myMidca.append_module("Eval", evaluate.SimpleEval())
6. myMidca.append_module("Intend", intend.SimpleIntend())
7. myMidca.append_module("Plan", planning.PyHOPPlanner())
8. myMidca.append_module("Act", act.SimpleAct())
```

Examining Table 2, we see that some phases only have a single module. For example, Perceive uses a single module called *PerfectObserver* when interacting with the MIDCA simulator. Other phases such as Interpret have multiple modules (as you can see with *SimpleDetect*, *SimpleExplain*, and *SimpleGoalGen*). The most common module for Evaluate is *SimpleEval* which checks to see if any goals have been completed, and if so, drops them. Another useful module (not shown here) for Eval is *Scorer* which is used to calculate the score MIDCA has received from a recently achieved goal. *SimpleIntend* chooses one or more goals which are stored in a list structure in the memory variable `mem.CURRENT_GOALS`. Then in this example, the *PyHOPPlanner* generates a sequence of actions that will achieve the goal(s). The input to the planner is the current state of the world (usually stored in `mem.STATE` or as the last item in `mem.STATES`) and the current goal(s). The output is stored on the goal graph (described in Section 5.3). Finally, the most common module for the Act phase is *SimpleAct* which chooses the next action from the current plan, if one exists. For an example of an Act module operating on a robot, see the *AsynchronousAct* module.

The order in which the modules are appended is the order the modules will run, though it is possible to pass an additional argument to *append\_module* which specifies the order (see code documentation). While there is no constraint that a module must only be used in a single phase, most often modules are designed for a specific phase and used for that phase only.

In an agent with only cognitive-level behavior, each module is executed sequentially within each phase, and when the last module of a phase is executed, the first module of the next phase is executed. When the last module of the last phase is executed, the first module of the first phase is executed next, completing the circle. Figure 8 shows an example of the order of execution, starting with the Perceive phase.




Figure 8. Module execution sequences (a) example execution sequence of a cognitive-only agent (orange, left); (b) example module execution at the metacognitive level (blue, right)

Metacognitive phases are implemented in a similar fashion as cognitive phases, and in the current version of MIDCA, a single metacognitive loop is run in-between each cognitive level module (see Figure 9). A single metacognitive loop refers to starting with the first module of the first metacognitive phase and executing each metacognitive module until the last metacognitive module of the last metacognitive phase (Control). Figure 10 shows the order in which all cognitive and metacognitive modules are executed. The cognitive-level phases starting with Perceive are described in the next subsection.




Figure 9. Example showing interleaving of metacognitive modules continuing with examples from Figure 8




Figure 10. Complete sequence of all modules in the order they are executed in a metacognitive-enabled agent (starting from the top left and continuing down, ending at the bottom right)

## 5.2. Adding a New Module

To add a cognitive or metacognitive module to MIDCA\_1.5, start by locating the modules/ and metamodules/ folders. The new module will be a .py file in subdirectories of either of these folders. It may make sense to simply add the new module into an existing .py file (i.e., adding a new class to interpret.py or planning.py). A module in MIDCA must meet the following criteria.

1. Inherit the BaseModule class found in midca/midca/base.py
2. Implement the init() method. Note that this is different than the default \_\_init\_\_() method. This init() method will be called by MIDCA.
3. Implement the run() method. This is the code that will run each time the module is executed.
4. Finally, after the module is ready, it must be assigned to a phase during the initialization of the agent, which happens in a startup script found in the midca/midca/examples/ directory. The startup script will contain code like the code shown in Table 3. Make sure to add the new module here.

Modules have access to MIDCA's memory from the init() method and can access this memory later during subsequent calls to run(). For more ideas on how to implement a module, browse the current modules found in the various .py files located in midca/midca/modules/ and midca/midca/metamodules/.

## 5.3. The Goal Graph and Examples

MIDCA's goal graph is an important data structure that maintains all goals and associated plans of MIDCA\_1.5. An example of a goal file can be found at `midca/midca/goals.py`. `Goals.py` contains three main classes: `Goal`, `GoalNode`, and `GoalGraph`. The current goals the agent is committed to achieving are stored outside the goal graph in the `CURRENT_GOALS` memory variable (found in `midca/midca/mem.py`). The graph structure maintains a partial ordering of unique goals. The root nodes of the graph contain goals that have precedence over the goals in subsequent child nodes. If a goal is inserted with less precedence than a root node, it will become either a child of that root node or a child of that node and so on. Each goal node has the current plan associated with achieving that goal. The following table shows how each module in the current version of MIDCA interacts with the function `GoalGraph`.

Table 3. Phases of MIDCA interaction with goal graph

Module	Interaction with Goal Graph
Percieve ( <code>percieve.py</code> : <code>PerfectObserver</code> )	No Interaction
Interpret #1 ( <code>note.py</code> : <code>ADistanceAnomalyNoter</code> )	No Interaction
Interpret #2 ( <code>guide.py</code> : <code>UserGoalInput</code> )	Inserts new goal into the goal graph
Interpret #2 ( <code>guide.py</code> : <code>TFStack</code> )	Iterates over the goals from <code>goalgraph.getAllGoals()</code> to check if a block stacking goal already exists. If not, inserts the goal from the TF-Tree into graph
Interpret #2 ( <code>guide.py</code> : <code>TFFire</code> )	Inserts the goal from the TF-Tree into graph Checks the result of the <code>goalgraph.insert</code> function.
Eval ( <code>evaluate.py</code> : <code>SimpleEval</code> )	Checks to see if all goals are achieved and if so calls <code>goalgraph.remove(Goal g)</code> for each goal. Subsequently calls <code>goalgraph.removeOldPlans()</code>
Intend ( <code>intend.py</code> : <code>SimpleIntend</code> )	Checks to see if <code>goalgraph</code> has been initialized. Calls <code>goalgraph.getUnrestrictedGoals()</code> and then sets those goals to the memory variable <code>CURRENT_GOALS</code>
Plan ( <code>planning.py</code> : <code>PyhopPlanner</code> )	Calls <code>goalgraph.getMatchingPlan(CURRENT_GOALS)</code> , if exists, it checks validity. If no matching plans or plans are not valid, calls <code>Pyhop</code> and calls <code>goalgraph.addPlan(midcaPlan)</code>
Act ( <code>act.py</code> : <code>SimpleAct</code> )	Iterates over each plan in <code>goalgraph.getAllMatchingPlans(CURRENT_GOALS)</code> to return the plan that achieves the most goals.

**Plans:** Plans are stored in a set in the goal graph. Plans are added to the goal graph and are stored in a set with all other plans. In the future, plans will be stored in the node of each goal. The following functions provide interaction with the current plans in the goal graph:

**addPlan(plan):** Adds the given plan into the current set of plans by calling the built-in set `add()` function.

**removePlan(plan):** Removes the given plan by calling the built-in set `remove()` function

**removePlanGoals(plan):** Removes all goals associated with the given plan. It checks the plan object for its goals, and removes each of those.

**removeOldPlans():** Removes every plan whose goals are no longer in the goal graph.

**allMatchingPlans(goals):** Given 1 or more goals, this will return all plans where the goal of the plan is the same as the goal passed in as the argument.

**getMatchingPlan(goals):** Returns a plan whose goalset contains all given goals. If multiple plans exist, it chooses the one with minimum extraneous goals, and if the plans tie, the tie is broken arbitrarily. If no plan succeeds all the goals, returns None.

**getBestPlan(goals):** Returns the plan that achieves the most goals in the set of goals passed in as an argument. Tries to achieve fewest extraneous goals, and if no plan achieves any of the goals, returns None.

*Example with a plan that fails to achieve all goals*

```
***** Starting Perceive Phase
*****
World observed.          (No goal graph interaction)

***** Starting Interpret Phase
*****
No anomaly detected.
TF-Tree goal generated: Goal(C_, B_, predicate:
on)
Check that no goal exists that uses 'on' predicate (calls getAllGoals())
Insert new goal into goal graph (insert(Goal(C_,B_, predicate: on)))

***** Starting Eval Phase *****
No current goals. Skipping eval (no goalgraph interaction b/c CURRENT_GOALS is different than
goalgraph)

***** Starting Intend Phase *****
(calls getUnrestrictedGoals() and sets these to be the current goals, of which there is only 1
right now)
Selecting goal(s): Goal(C_, B_, predicate: on)

***** Starting Plan Phase *****
(first a call to getMatchingPlan((Goal(C_,B_, predicate:on)) was made but returned empty so
proceeded to planning).
Planning...
Planning complete.
Plan: unstack(D_, B_) putdown(D_) pickup(C_) stack(C_, B_)
(after planning completed the plan (unstack(D_, B_) putdown(D_) pickup(C_) stack(C_, B_)) was
passed to goalgraph.addPlan())

***** Starting Act Phase *****
(calls goalgraph.getAllMatchingPlans() which only returns the following plan)
Selected action unstack(D_, B_) from plan:
unstack(D_, B_) putdown(D_) pickup(C_) stack(C_, B_)

***** Starting Simulate Phase *****
simulating MIDCA action: unstack(D_, B_)

***** Starting Perceive Phase *****
World observed.
Next MIDCA command:

***** Starting Interpret Phase *****
```

```

No anomaly detected.
MIDCA already has a block stacking goal. Skipping TF-Tree stacking goal generation
Please input a goal if desired. Otherwise, press enter to continue
on(A_,C_)
Goal added.
Insert new goal into goal graph (insert(Goal(A_,C_, predicate: on)))
Please input a goal if desired. Otherwise, press enter to continue

***** Starting Eval Phase *****
(no plans are finished achieving goals so there are no goals to remove)
Not all goals achieved; Goal(C_, B_, predicate: on) is not true.

***** Starting Intend Phase *****
(calls getUnrestrictedGoals() and sets these to be the current goals)
Selecting goal(s): Goal(C_, B_, predicate: on) Goal(A_, C_, predicate: on)
Next MIDCA command:

***** Starting Plan Phase *****
(calls getMatchingPlan(Goal(C_, B_, predicate: on) Goal(A_, C_, predicate: on)) which fails
and so the current plan doesn't change)
Planning...
Planning failed for Goal(C_, B_, predicate: on) Goal(A_, C_, predicate: on)
Next MIDCA command:

***** Starting Act Phase *****
(calls goalgraph.getAllMatchingPlans() which only returns the following plan)
Retrieved plan does not achieve all goals. Trying again.
Plan: unstack(D_, B_) putdown(D_) pickup(C_) stack(C_, B_)
Goals achieved: ['Goal(C_, B_, predicate: on)']
Best plan does not achieve all goals.
Plan: unstack(D_, B_) putdown(D_) pickup(C_) stack(C_, B_)
Goals achieved: ['Goal(C_, B_, predicate: on)']
Selected action putdown(D_) from plan:
unstack(D_, B_) putdown(D_) pickup(C_) stack(C_, B_)

```

### *Examples of goal graph drawings*

To produce a pdf drawing of the goal graph at any time, just run the following MIDCA command.  
**drawgoalgraph**

#### *Example 1*

Here is a goal graph when multiple fire and block stacking goals are given with the following compare function:

```

1. def preferFire(goal1, goal2):
2.     if 'predicate' not in goal1 or 'predicate' not in goal2:
3.         return 0
4.     elif goal1['predicate'] == 'onfire' and goal2['predicate'] != 'onfire':
5.         return -1
6.     elif goal1['predicate'] != 'onfire' and goal2['predicate'] == 'onfire':
7.         return 1
8.     return 0

```




Figure 11. Goal graph for Example 1.

### Example 2

Figure 12 illustrates what a goal graph may look like when multiple fire and block stacking goals are given with the following compare function:

```

1. def preferAnythingButFire(goal1, goal2):
2.     if 'predicate' not in goal1 or 'predicate' not in goal2:
3.         return 0
4.     elif goal1['predicate'] == 'onfire' and goal2['predicate'] != 'onfire':
5.         return 1
6.     elif goal1['predicate'] != 'onfire' and goal2['predicate'] == 'onfire':
7.         return -1
8.     return 0

```




Figure 12. Goal graph for Example 2.

## 5.4. Goal Operations

We recognize a number of operations on goals and distinguish them from operations on plans (Cox, Dannenhauer, & Kondrakunta, 2017). Although the purpose of a plan is to establish a state of the world that satisfies a goal or a set of goals, we argue that the separation of goal and planning operations provides at least an organizational benefit within a cognitive architecture. However, like Roberts and colleagues, who combine both types of operations into a goal life-cycle framework (Roberts et al., 2015), we acknowledge the close relationship between the two. Table 3 classifies the ten primary operations on goal expressions.

MIDCA Version 1.5 implements many of the goal operations shown in Table 4 which are as follows:

1. MIDCA formulates goals (goal formulation) when an anomaly is detected in the blocks world domain. The user can find an example in the “cogsci\_demo” script.
2. MIDCA selects a goal (goal selection) from all the goals based on two selection methods: FIFO and a smart selection method. User can find the implementation in the example files “cogsci\_demo\_mortar\_construction” and “restaurant”.
3. MIDCA suspends the goal (goal suspension) which it is currently working on when an anomaly is detected and performs the formulated goal.
4. MIDCA resumes a previously suspended goal (goal resumption) and continues to work on that goal.
5. MIDCA performs the goal change operation when the resources are not sufficient for the agent to continue on with the goal. User can observe this in the example file “cogsci\_demo\_mortar”. The agent changes the goal of “stable-on” to “on”.
6. MIDCA also performs the goal monitoring operation (Dannenhauer & Cox, 2018)
7. MIDCA checks if the goal is achieved (goal achievement) this operation is performed in the evaluate phase of MIDCA.

Table 4. Fundamental set of goal operations (adapted from Cox, Dannenhauer, & Kondrakunta, 2017)

No.	Operation	Description	Phase
1	Goal formulation	Create a new pending goal	Interpret
2	Goal selection	Commit to an active goal from the set of pending goals	Intend
3	Goal suspension	Pause in pursuit of a currently committed goal	Not Implemented
4	Goal resumption	Resume pursuit of a suspended goal	Not Implemented
5	Goal change	Change a goal into a similar one that is close to the original goal	Control
6	Goal monitoring	Track that a goal maintains its usefulness	Interpret
7	Goal delegation	Find another agent willing to pursue a goal for you	Not Implemented
8	Goal interpretation	Infer the meaning of a stated intent by another agent	Not Implemented
9	Goal abandonment	Remove a pending or committed goal from consideration	Interpret
10	Goal achievement	Verify that a goal state is satisfied in some environment	Evaluate

## 5.5. How to Install and Run Version 1.5 of the MIDCA Architecture

### *Installing MIDCA*

1. Obtain a copy of MIDCA by cloning the repository or downloading the source directory. (<https://github.com/COLAB2/midca>)
2. Make sure the name of the top-level folder is spelled exactly 'midca' (if midca has been downloaded as a zip file, you will have to rename it as it saves the folder as 'midca-master')
3. Run the command `python setup.py install`
  - If you plan to make changes to MIDCA, do `python setup.py develop` instead. Any changes you make will be immediately updated when you run MIDCA.
  - Note that NumPy (<http://www.numpy.org/>) should be installed automatically when you run `python setup.py install` (or with the 'develop' option). If for any reason that fails, you will need to install the package yourself. You can check that NumPy is installed by running python and typing `import numpy`. If nothing happens, you have successfully installed it. Otherwise you will get an error message.
4. (Optional) Graphviz is needed for drawing graphs and saving them to a pdf. It must be installed manually.
  - See [www.graphviz.org](http://www.graphviz.org)

### *Using MIDCA with simulated worlds using a predicate representation*

1. Create a simple MIDCA version which allows text-based goals to be input at runtime.

```
1. #set locations of files defining domain and world state
2. domainFilename = "myDomainFile"
3. stateFileName = "myStateFile"
4.
5. from MIDCA.examples import predicateworld
6.
7. myMIDCA = predicateworld.UserGoalsMidca(domainFilename, stateFileName)
```

2. See all phases in a MIDCA instance.

```
1. print myMIDCA.get_phases()
```

3. Add/remove phases

```
1. myMidca.insert_phase(phaseName, i)
2. myMidca.append_phase(phaseName)
3. myMidca.remove_phase(phaseName) #throws ValueError if phase named phaseName not present
```

4. See the classes that are implementing a phase.

1. `print myMidca.get_modules(phaseName)`
5. a) Create a custom phase implementation (module) without inheriting from the `BaseModule` class

Create a python class with at least these two methods:

1. `init(self, world, mem)`
2. `run(self, cycle, verbose)`

- world is the initial world state
- mem is MIDCA's central memory
- cycle is the cycle # (starting with 1)
- verbose is the level of output requested
- the init method should do any setup the module requires. It will be called once for each module during MIDCA's initialization. Init methods will be called in phase order.
- the run method will be called once per cycle in phase order. This method will define what the module actually does. Within a phase, modules will be called in the order listed, which can be modified as shown in 6.

### Running MIDCA

For examples, see `midca/midca/modules/*`

- b) Create a custom phase implementation (module) - new style:

- create a subclass of the `base.BaseModule` class. To do this, you must implement the `run(self, cycle, verbose)` method, which defines the module's behavior. In this method, you can access MIDCA's memory through the `self.mem` field.

6. Add/remove custom or predefined modules to/from MIDCA

```

1. myModule = MyModule() #
2. assert hasattr(myModule, 'run') and hasattr(myModule, 'init')
3.
4. myMidca.append_module(phaseName, myModule)
5. myMidca.insert_module(phaseName, myModule, i)
6. #i is the index of where this module should be placed during the phase. This is for ordering when more than one module is used in a single phase.
7.
8. myMidca.clear_phase(phaseName) #removes all modules implementing the phase

```

7. Initialize

```
1. myMidca.init()
```

8. By default, MIDCA runs in interactive mode. There are two modes: interactive and non-interactive.

1. `myMidca.run()` #by Default; runs in interactive mode.
2. `myMidca.run(usingInterface=False)` # non-interactive mode

## 9. logging

From 'outside' MIDCA use the following.

1. `myMidca.logger.log(msg)`

From inside a module that inherits `BaseModule` from a MIDCA module's run method, use the following.

1. `self.log(msg)`

**Note:** by default, MIDCA will automatically log everything sent to standard output. To turn this off, set the MIDCA constructor argument 'logOutput' to False.

**Note:** by default, MIDCA also logs each memory access. To turn this off, set the MIDCA constructor argument `logMemory` to False, or set `myMidca.mem.logEachAccess` to False.

### *Understanding how MIDCA works from browsing the source code*

1. Start with the `base.py` file. The method `PhaseManager.run()` defines the behavior of MIDCA in interactive mode, and follows the relationship between user inputs and associated function calls which illustrates what MIDCA is doing.
2. Each module is defined independently and they interact only through memory. In the `mem.py` file, the `Memory` class has a list of constants that define keys for the default MIDCA memory structures (e.g. the goal graph, observed world states). The built-in MIDCA modules generally interact only through reading/writing to the values referred to by these keys.
3. To understand MIDCA behavior at a more fine-grained level, it is necessary to look through module by module to see what each one is doing. Check the MIDCA object to see what modules it runs in each phase (see docs above - printing a module should show the file and class name of its implementation), then go to the file in the modules folder to see what it does. Especially note the calls to the memory structure (`self.mem`), since these are the I/O.

### ***Runtime Commands***

The following commands can be used when MIDCA is run in interactive mode.

1. **change.** Allows the user to change the state of the world. The user can either give a file name to be loaded or enter atoms one at a time.
2. **drawgoalgraph.** Generates a visual representation of the current goal graph.  
**Requirement:** Using this command requires Graphviz to be installed ([www.graphviz.org](http://www.graphviz.org)).

3. **help.** Displays the possible commands that can be given to MIDCA during runtime (the commands detailed here).
4. **log.** Prompts the user to enter a message which will be written to the log file. If the user doesn't enter a message, no log message is written.
5. **memorydump.** Allows the user to see variables in MIDCA's memory. This is useful when the user wants to see the value of a variable. The user can either see all the variables and their values, or enter a single variable name and just see that value. If MIDCA has been running a long time, the output may take up more than the screen, therefore just looking for the variable can save space.
6. **printrace.** Outputs a text representation of the execution up to the last phase run.
7. **q.** Quits MIDCA.
8. **show.** Displays the world.
9. **skip (&optional x=1).** Skips ahead x cycles or one full cycle if x is not given.
10. **toggle meta verbose.** Turns off/on meta output. It is useful to turn this off to reduce the amount of text seen while running MIDCA. The first time running this command it will turn the meta output OFF and then it can be turned on later by running the command again.

All the above commands can be typed in the terminal while MIDCA is running in interactive mode.

## 6. Defining a New Domain

To add a new domain for MIDCA, all the domain-specific material should be in the domains/ folder, named for your new domain. Domain folders are structured as follows:

- midca/midca/domains/
  - New\_Domain\_Script.py
  - your-new-domain/
    - your-new-domain.sim
    - init.py
    - util.py
    - plan/
      - domain.pddl (Only if using PDDL)
    - states/
      - problem1.sim
      - problem2.sim
    - domain.cfg (optional)
    - ros/ (optional)

The file *your-new-domain/your-new-domain.sim* contains the logic for states (types, predicates, and operators) that MIDCA's simulator will use (see *domains/blocksworld/arsonist.sim* for an example). If the planner for the domain uses pddl for its action representations, the domain.sim file can be automatically generated by running *New\_Domain\_Script.py* with the domain.pddl file.

The folder *your-new-domain/plan/* contains any material the planner will make use of (e.g. HTN methods and operators, domain-specific heuristics).

The file *your-new-domain/plan/domain.pddl* is the domain file provided by you and will be passed into `New_Domain_Script.py` to create the rest of the directory. The name for this file is arbitrary but must be provided as the command line argument. Your domain directory and this `.pddl` file are the only requirements to run the new domain script. (Only if using PDDL)

The folder *your-new-domain/states/* contains the starting states for problems MIDCA solves. See the file `domains/blocksworld/states/defstate_fire.sim` for an example. Any various starting states for this domain will be stored in this subfolder.

The folder *your-new-domain/ros/* is optional but it contains any ROS functions needed for performing robotic actions in a real or simulated world.

The file *your-new-domain/domain.cfg* is a configuration file used to specify entities and interactions within your domain. The file is parsed to create the interfaces between MIDCA and ROS. See Section 9.4, The MIDCA ROS API.

The file *your-new-domain/init.py* contains nothing, but an init file is mandatory for all folders in MIDCA. It is automatically created when running the new domain script.

The file *your-new-domain/util.py* contains any domain specific utility functions. For example, the blocksworld domain uses an HTN planner Pyhop (the planner can be found in `midca/modules/plan/PyHopPlanner.py`). The Pyhop Planner uses its own state representation which requires translation to and from MIDCA states (the state specified in `_your-new-domain/domain.sim` file). For blocksworld, this translation happens in `util.py`, specifically via the functions `pyhop_state_from_world()` and `pyhop_tasks_from_goals()`. Many other useful utility functions are located here, including how to draw an ascii representation of the world state in a terminal (optional, but useful).

Adding a new domain into MIDCA requires various domain knowledge, and this is all organized into the structure shown above. Mostly, this is to make it convenient to locate anything specific to a domain. When writing a startup script for MIDCA, the domain location needs to be given (see `examples/cogsci_demo.py` for an example).

This is summarized in the following steps:

1. Add a new folder under `midca/midca/domains/` (i.e., `midca/midca/domains/my_new_domain/`)
2. Add a new folder called `plan` to the `my_new_domain` directory
3. Add your `domain.pddl` file to the `/plan` directory (Only if using PDDL)
4. Run `python New_Domain_Script.py my_new_domain <path_to_domain.pddl>` (Only if using PDDL)
5. Add state files to the directory `midca/midca/domains/my_new_domain/states`

6. Implements necessary domain functions in util.py file
7. (IMPORTANT) Modify midca/midca/setup.py to include the new domain folder.

## 7. Logging and Debugging in MIDCA\_1.5

### 7.1. Logging

#### *Initiating and disabling a log file*

Logging is initialized in MIDCA\_1.5 by default. If the user does not want to use log files, then they can change the variable “logenabled” in phase manager class of base.py file to False.

#### *Location of log files*

Log files are stored in the directory: midca/midca/examples/log. Each time the user runs MIDCA, there will be a new folder created. The folder name is the current date and time and contains the log information. The path to the folder can also be viewed on the screen when we run MIDCA. An example is shown in the Figure 13.

```
Sravyas-MBP:examples sravyakondrakunta$ python cogsci_demo.py
Logger: logging this run in /Users/sravyakondrakunta/Documents/git/midca/midca/examples/log/2017-10-24 13_48_27
```

Figure 13. Example specification of the MIDCA\_1.5 log file.

#### *Files in the log folder:*

The folder contains three files named “log”, “Memory Access”, and “MIDCA output”. These files can be opened in any text editor.

- “log” – log file contains the information about the variables accessed in the memory for each step and the verbose statements printed on the screen for each phase and during initialization. **Error! Reference source not found.** shows a part of log file.

Table 5. Code snippet of a log file.

```
.000331 - Memory access at key __goals
.000460 - Goal Graph initialized.
.000600 - [cognitive] Initializing Simulate module 1
FireReset...done.
.000697 - [cognitive] Initializing Simulate module 2
MidcaActionSimulator...done.
.000788 - [cognitive] Initializing Simulate module 3
ArsonSimulator...done.
.000942 - [cognitive] Initializing Simulate module 4
ASCIIWorldViewer...done.
.001038 - [cognitive] Initializing Perceive module 1
PerfectObserver...done.
.002258 - Memory access at key A-Distance memory
```

- “Memory Access” – This is the subset of information within the log, which contains only the information pertaining to accessing memory variables in MIDCA. **Error! Reference source not found.** represents a part of memory access file. This gives the description of variables present in memory when they are accessed.

Table 6. Code snippet of memory access file.

```
.000331 - Memory access at key __goals
.002258 - Memory access at key A-Distance memory
.002707 - Memory access at key Last Scored Goal
.002764 - Memory access at key Score
.002973 - Memory access at key __PlanningCount
.003167 - Memory access at key __goals
.223371 - Starting cycle 0
.223945 - ***** Starting Simulate Phase *****
```

- “MIDCA output” – This file represents the information related to only the verbose statements for each phase. All of the user experiments can be saved in this and a user can look at this to trace all the previous experiments. **Error! Reference source not found.** shows a snippet of a MIDCA output file.

Table 7. Code snippet of the MIDCA\_1.5 output file.

```
.000460 - Goal Graph initialized.
.000600 - [cognitive] Initializing Simulate module 1 FireReset...done.
.000697 - [cognitive] Initializing Simulate module 2
MidcaActionSimulator...done.
.000788 - [cognitive] Initializing Simulate module 3 ArsonSimulator...done.
.000942 - [cognitive] Initializing Simulate module 4
ASCIIWorldViewer...done.
.001038 - [cognitive] Initializing Perceive module 1
PerfectObserver...done.
.002337 - [cognitive] Initializing Interpret module 1
ADistanceAnomalyNoter...done.
.002414 - [cognitive] Initializing Interpret module 2 TFStack...done.
```

## 7.2. Debugging

The debugging action can be performed differently for various files, i.e., the examples can be debugged by tracking different phases.

### *Debugging in MIDCA\_1.5*

MIDCA always runs in a cycle of phases. Each phase is displayed on the screen, which helps the user to identify the verbose statements belonging to the specific phase. To use full MIDCA output, use the “toggle meta verbose” command if in interactive mode.

### *Debugging through log files*

Log files contain the information of the variable accessed in memory for each phase. With this information we can easily debug to find what memory variable is accessed and if an error exists,

we can find the last access to the memory variable and take it from there. With the help of log files, we can also store our experiments run in a MIDCA session.

## 8. Multiagent Interaction

### 8.1. Goal delegation

## 9. Advanced Features

This section describes a topic interface between MIDCA and ROS, shows how it can be applied to controlling a humanoid robot, and then explains an associated interface with the Gazebo simulator. It also describes an API between MIDCA and ROS, then shows how to run it.

### 9.1. The MIDCA\_1.5 Topic Interface to ROS

We added an interface to MIDCA to communicate with ROS and a Baxter humanoid robot (see Figure 14). It is responsible for sending messages to ROS as requested by MIDCA, and for placing messages received in appropriate queues for MIDCA to process. We created other ROS nodes which are responsible for doing specific actions, such as moving the Baxter's arms, and for getting object representations. These communicate with MIDCA through ROS Topics (<http://wiki.ros.org/Topics>).




Figure 14. Interfaces between MIDCA and the external world and between cognition and metacognition.

The types of ingoing and outgoing messages on each ROS topic and their meaning is specified. As these messages are asynchronously received, a set of MIDCA handlers put them in appropriate buffers within a partition of MIDCA memory. During the Perceive phase, these messages will be accessed and stored in MIDCA's main memory.

The interface treats MIDCA as a ROS node which can send and receive messages to other ROS nodes. A ROS node is an executable that uses ROS to communicate with other ROS nodes. At the beginning, a RosMidca object is created which is responsible for sending messages to ROS as

requested by MIDCA and for placing messages received in appropriate queues for MIDCA to process. Different unique topics will be used for different ingoing and outgoing messages.

An external voice recognition node is constantly running which publishes utterances as string messages on UTTERENCE\_TOPIC. Once MIDCA receives any message on this topic it puts the message on an appropriate queue which will be processed in the perceive phase later.

An example of a RosMidca object is:

```
class DoRaise(AsynchAction):

    def __init__(self, mem, midcaAction, objectOrID, maxAllowedLag, maxDuration, topic,
msgID):
        self.objectOrID = objectOrID
        self.maxAllowedLag = maxAllowedLag
        self.maxDuration = maxDuration
        self.lastCheck = 0.0
        self.topic = topic
        self.complete = False
        self.msgID = msgID
        executeAction = lambda mem, midcaAction, status: self.send_point()
        completionCheck = lambda mem, midcaAction, status: self.check_confirmation()
        AsynchAction.__init__(self, mem, midcaAction, executeAction,
completionCheck, True)

    def send_point(self):
        raising_point = self.mem.get(self.mem.RAISING_POINT)

        self.msgDict = {'x': raising_point.x, 'y': raising_point.y,
'z': raising_point.z, 'time': self.startTime, 'cmd_id': self.msgID}

        sent = rosrun.send_msg(self.topic, rosrun.dict_as_msg(self.msgDict))
        if not sent:
            if verbose >= 1:
                print "Fail"
            self.status = FAILED
```

In the Perceive phase, MIDCA reads messages from all the queues, processes them and stores the processed data in MIDCA's main memory. In the Interpret phase, MIDCA checks to see if it has received any instruction. If it detects the message 'get the red block', it will create the goal 'holding the red block'. Once the goal is created, it will be stored in the goal graph.

In the Plan phase, after it creates a high level plan for the selected goal, it operationalizes each action using a mapping between high-level actions and directly executable methods for the robot. For example, the high level action reach(object) is instantiated in a method which sends out a ROS message to the node which operates the arm, then repeatedly checks for feedback indicating success or failure. Once all actions in a plan are complete, the plan itself is considered complete. If any action fails, the plan is considered failed.

In the Act phase, MIDCA loads a plan for the current goal from memory. In each cycle, one action starts and if it is completed in a certain amount of time, the next action will start in the next cycle. If any action fails, the rest of the actions won't start and the plan fails.

We created other ROS nodes for our purpose which are responsible for doing specific actions, such as moving the Baxter's arms, and for getting object representations. These communicate with MIDCA through the topics interface. These processes listen to different ROS topics for MIDCA command and act on them appropriately. As previously mentioned, the Act phase contains different modules. When an action is chosen, these modules publish the correct command, which the low-level processes are listening to.

## 9.2. Using the MIDCA interface to ROS for the Baxter Humanoid Robot

We are working toward cooperative interaction between humans and machines by starting with a small instructional problem for the robot to accomplish. A Human asks the robot to pick up a colored block. The Robot needs to understand what the human wants and create a plan to achieve it. The MIDCA architecture provides the reasoning to processes the command and execute the right actions.

1. **Speech to text:** Audio signal is translated to text string percept.
2. **Infer goal:** The text is mapped to user intent.
3. **Create plan:** The SHOP2 planner creates a plan to achieve goal.
4. **Execute plan:** Actions sent to topics.

In this section, we describe a demo with a Baxter robot performing in Blocksworld domain. The robot is asked to stack blocks, unstack them, put them on the table, or give a block to the user.

### *Blocksworld domain for the Baxter*

Consider an example when Baxter is given the goal  $on(G, R)$ . This goal is mapped to the root task *move-blocks* in Pyhop planner. The planner decomposes *move-blocks* to the non-primitive  $\langle pickupT(G), stackT(G, R) \rangle$  tasks in that order. The task  $pickupT(G)$  decomposes to the primitive tasks  $\langle moveto(loc(G)), grasp(G) \rangle$ , and  $stackT(G, R)$  decomposes to the primitive tasks  $\langle moveto(loc(R)), stack(G, R) \rangle$ .

### *ROS topics*

The Topic Names used in this demo are listed below.

1. **OBJ\_LOC\_TOPIC:** the topic on which object detection messages are published
  - must be changed in MIDCA run script(`baxter_run.py`) and object detection node (e.g. `OD.py`)
2. **UTTERANCE\_TOPIC:** the topic on which utterances (e.g. commands) baxter hears are published
  - must be changed in MIDCA run script(`baxter_run.py`) and utterance listener node
3. **POINT\_TOPIC:** the topic on which point commands from MIDCA are published.
  - must be changed in `asynch.py` and in point effector node (`grabbing.py` reads `asynch's` value)

### *External steps (sensors and effectors)*

These can be started in any order, but must all be started for the demo to work.

1. Start an external object detection ROS node which publishes a PointStamped ROS msg on OBJ\_LOC\_TOPIC. Simulated Implementation: drone\_location\_simulator.py, OD.py.
2. Start an external voice recognition node which publishes utterances as string messages on UTTERANCE\_TOPIC. Implementation is currently in the baxter\_cog repository, to be added to MIDCA.
3. Start an external pointing effector node which listens for point commands on POINT\_TOPIC. A point command will be in the form of a String message encoding a Dictionary containing x,y,z coordinates. rosrn.py contains methods for transforming between String and dict; see examples/\_baxter/pointing.py for an implementation. This node should also publish feedback when it encounters an error or completes its task. This too is implemented in pointing.py.

### *MIDCA setup: all steps from baxter\_run.py.*

1. Create a new MIDCA object and add robot domain-specific modules to it.
2. Create a RosMidca object. This object is responsible for sending messages to ROS as requested by MIDCA, and for placing messages received in appropriate queues for MIDCA to process. At present, all topics which will be used for incoming or outgoing messages must be specified at creation.
3. Pass in handlers as arguments to the RosMidca constructor for incoming and outgoing messages. In this demo, MIDCA uses 3 incomingMsgHandlers and 1 outgoingMsgHandler:
  - A FixedObjectLocationHandler receives information about the location of a single, prespecified object.
  - An UtternanceHandler receives utterances as Strings.
  - A FeedbackHandler receives feedback regarding the success or failure of requested actions
  - An outgoingMsgHandler sends out String messages representing point commands.
4. Call ros\_connect() on the RosMidca object. Note that the ROS master node must already be started or this method will fail.
5. Call run\_midca() on the RosMidca object. This will run MIDCA asynchronously. If certain rate (phases/second) is desired, it can be input as the cycleRate argument of this method (default 10).

### *What MIDCA does while running*

1. Asynchronously to cyclical behavior, RosMidca's handlers listen for incoming messages. As they are received, handlers place them into appropriate queues in a partition of MIDCA's memory that could be thought of as the subconscious,

or perhaps preconscious.

- **Note:** if external perception changes its output style or capabilities, the handlers - defined in `rosrun.py` - are responsible for adjusting to process the new input. Specifically, for each new input type or format, a new handler should be created.
2. In the perceive phase, MIDCA reads messages from all queues, processes them as necessary, adds a time stamp to indicate when each message was received, and stores the processed data in MIDCA's main memory. Note that only the perceive phase accesses the incoming message queues.
  3. In the interpret phase, MIDCA checks to see if it has received any verbal instructions. If it gets the message 'point to the quad', it will create the goal: `Goal(objective = "show-loc", subject = "self", directObject = "quad", indirectObject = "observer")`. Currently it also interprets the phrase "goodbye baxter" in the same way, simply because in testing it was sometimes difficult for the voice recognition software to understand "point to the quad". Once a goal is created it will be stored in the goal graph. In this demo, since all goals are identical and identical goals are only stored once, there will never be multiple goals in the goal graph, though the same goal may be added again after it is achieved and removed.
  4. In the Evaluate phase, MIDCA checks to see if its current plan is complete. If it is, it declares the goal of that plan completed and removes the goal and plan from memory.
  5. In the Intend phase, MIDCA selects all goals of maximal priority from the goal graph. In this demo, there is never more than one goal, so MIDCA will select that goal if it exists in the graph.
  6. In the planning phase, MIDCA checks to see if an old plan exists for the current goal. If not, it creates a high level plan by using the pyhop planner, then transforms it into an actionable plan using a mapping between high-level actions and methods to carry them out. For example, the high level action `point_to(object)` is instantiated in a method which sends out ROS messages to the pointing effector node, then repeatedly checks for feedback indicating success or failure. Once all actions in a plan are complete, the plan itself is considered complete. If any action fails, the plan is considered failed.
  7. In the Act phase, MIDCA attempts to load a plan for the current goal from memory. If a plan exists, it follows this pattern:

```
currentAction = plan.firstAction
while currentAction != None:
    if currentAction.complete:
        currentAction =
plan.nextAction()
        continue
    else if currentAction.not_started:
        currentAction.start()
        if currentAction.failed or
currentAction.isBlocking:
            break
```

In other words, actions begin successively until either one fails or a blocking action is reached. Actions are assumed to be running asynchronously from when they start to when they are declared completed.

## 8. Lower-level details of planning and action

- Planning methods and operators for this demo are in the `_planning/asynch` folder.
- Low-level methods - see the `point_to` example in 6) - are defined in `_planning/asynch/asynch.py`.
- Each `Asynch[ronous]Action`, a low-level method, defines an incomplete python function and an `executeFunc` function, which are passed into the constructor as arguments. These methods fully define the behavior of the action. So, for example, the `do_point()` `AsynchAction`'s `isComplete` function checks MIDCA's memory for feedback indicating the action's completion or failure, then updates its status appropriately. The `execute` function searches memory for the last known location of the object given as an argument (from the high-level plan), then creates a ROS message containing that location and a command id, for later feedback, and requests that `RosMidca` broadcast the message.
- This setup means that if external effectors change their input requirements, MIDCA's high-level planning can stay the same, but the interface between the two defined in `asynch.py` must change. Specifically, a new `AsynchAction` must be created for each new behavior type, though this process could be automated to some degree.
- As an aside, the mirror of the last point with respect to perception rather than action is also true. See the note after 1).

### *Camera calibration.*

We use the Baxter' right hand camera to observe the table and use the color of the objects to find where it is in the image. Then using camera calibration we find the location of the object on the table plane. The ROS service called "Right\_hand\_camera" was created to retrieve the image from Baxter' right hand camera that subscribed to the topic `/cameras/right_hand_camera/image`. This is where the images coming from camera are published. Upon calling this service, it stores and returns the last image published.

To get the coordinates of a pixel on the table plane from coordinates of the object in the image, we performed a calibration task before running MIDCA. In this calibration, we marked four points on the table and sample them moving the Baxter' left arm to each point to get the position of the end-effector of the left arm. Then we can visualize them in the image and the corresponding pixels by clicking on those points on the image. With those points, we can calculate a matrix  $H$  that represents a linear transform between points in the floor and points in the image, called a

homography. Using the inverse matrix  $H$ , given a pixel in the image we can calculate the coordinate of an object in the table plane.

We use OpenCV library to recognize an object by its color in the way filters color from a certain range in HSV color space that corresponds to each color, transforming it in black and white image. Let's say red is the current color, it uses the filter to change every red pixel to white. Then it finds the largest contour on the image which represents the object. Then the algorithm finds the center of this contour as the selected pixel and using the inverse matrix  $H$  it finds the object position on the table plane. For each defined HSV range, this algorithm finds the location of that colored object on the table plane and adds it to a list. MIDCA will receive a list of objects that are visible in the current scene, with the information on their color and location.

To make this work, we used the Baxter' right hand camera to monitor the table and get the objects position based off the color of the object ([www.nildo.me/organizer-baxter](http://www.nildo.me/organizer-baxter)). We manually place the end effector of Baxter' left arm in four different points over the table, this way we can easily map the table. After this, we click on the exact points on the image given by the camera. So, we can make a cross reference between those points and compute the homography matrix from the image to the table.

### 9.3. The Baxter and Gazebo

Gazebo is an open-source robot simulator whose code is publicly available at the URL [gazebo.org](http://gazebo.org). The requirements to run the demo are as follows.

1. UBUNTU 14.04 LTS
2. Python 2.7

The package includes a graphical user-interface, a programmable software interface, a physics engine, and high-quality graphics rendering.

#### Installation Instructions.

1. Workstation Setup: Follow steps from 1 to 4 from the link below:
2. SIMULATOR SETUP  
Follow steps to install GAZEBO from the reference link: [Simulator Installation](#) (ROS Indigo version)
3. MIDCA  
Follow instructions in Section 5.5 to install MIDCA.
4. DEPENDENCIES
  - Pocket Sphinx  

```
$ pip install --upgrade pip setuptools wheel
```

```
$ pip install --upgrade pocketsphinx
```
  - Robot State Publisher  

```
$ sudo apt-get install ros-indigo-robot-state-publisher
```

#### Running Example Scripts

This is the demo on the gazebo simulator working with the simulated Baxter robot on the blocksworld domain. We use three blocks in this domain and perform operations like stack, unstack, get and put.

1. Launch the Baxter Robot in the Gazebo Simulator (see Figure 18. Baxter Robot in gazebo, after executing the roslaunch command.

Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands. Figure 18. Baxter Robot in gazebo, after executing the roslaunch command.

Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands.

Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands. Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands. Figure 18. Baxter Robot in gazebo, after executing the roslaunch command.

2. Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands. Figure 18. Baxter Robot in gazebo, after executing the roslaunch command.).

- Open a new terminal window (ctrl + T)

```
$ cd ~/ros_ws (Go to the workspace directory)
$ ./baxter.sh sim
$ roslaunch baxter_gazebo baxter_world.launch
```




Figure 15. Baxter Robot in gazebo, after executing the roslaunch command.

Figure 16. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands. Figure 17. Baxter Robot in gazebo, after executing the roslaunch command.

Figure 18. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands.

Figure 19. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands. Figure 20. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands. Figure 21. Baxter Robot in gazebo, after executing the roslaunch command.

Figure 22. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands. Figure 23. Baxter Robot in gazebo, after executing the roslaunch command.

3. Get the models (table and blocks) into the simulator (see Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands).

Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands. Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands.

Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands. Figure 19. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands.

4. Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands. Figure 19).

- Open a new terminal (ctrl + T)
- Change current directory to the MIDCA Folder

```
$ ./baxter.sh sim  
$ cd examples/_gazebo_baxter  
$ python model.py
```




Figure 24. Baxter Robot along with table and blocks in gazebo, after executing the second set of commands

5. Position Right Hand Camera to view Blocks (Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

6. Figure 20. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.)
- Open a new terminal (ctrl + T)
  - Change to the MIDCA Folder

```
$ ./baxter.sh sim
$ cd examples/_gazebo_baxter
$ python tuck_arms.py -u
$ python position_right_camera.py
$ python publish_right_camera_to_xdisplay.py
```




Figure 32. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

Figure 33. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

Figure 34. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

Figure 35. Baxter Robot with left arm in tucked position and right arm in a position to view blocks, after executing the third set of commands.

7. Script to Detect Objects and send coordinates to MIDCA: Detects the blocks and puts the coordinates into MIDCA's Incoming Message Handler.

- Open a new terminal (ctrl + T)
- Direct to the MIDCA Folder

```
$ ./baxter.sh sim
$ cd examples/_gazebo_baxter
$ python OD.py
```

8. Script for Receive command from MIDCA and execute on MIDCA: During the Act phase MIDCA initiates the command into outgoing message handler that is received through the script `grabbing.py` and implements the actions to be performed physically.

- Open a new terminal (ctrl + T)
- Direct to the MIDCA Folder

```
$ ./baxter.sh sim
$ cd examples/_gazebo_baxter
$ python grabbing.py
```

9. Run MIDCA: Receives the instructions through voice commands and executes the phases. MIDCA runs in a continuous loop until terminated.

- Open a new terminal (ctrl + T)
- Direct to the MIDCA Folder

```
$ ./baxter.sh sim
$ cd examples
$ python baxter_run_OD_sim.py
```

10. Run Voice Commands: Use headphones to speak out voice commands, after running the scripts below format:

- stack the [block name] on the [block name]
- put the [block name] on table
- unstack the [block name]
- get the [block name]
- Example: Stack the “blue block” on the “red block”
  - o Open a new terminal (ctrl + T)
  - o Direct to the MIDCA Folder

```
$ ./baxter.sh sim
$ cd examples/_gazebo_baxter
$ python voice_cmds_sr.py
```

A video of MIDCA controlling a simulated Baxter robot through the Gazebo application is shown in the link below:

<https://drive.google.com/file/d/0B20nF1x2Cg5MX3ZGMC13dU1Na1U/view?usp=sharing>

#### 9.4. The MIDCA ROS API

To unify the naming process between ROS vision nodes and the MIDCA node, we have defined a language which relates all entities, attributes, possible values for the attributes, relations of entities,

and actions that can be taken in the world. This is the MIDCA API. A sample of this language for the popular domain, Blocksworld, is found below. We call this a configuration (.cfg) file, which is standard to store in the top level of your domain.

### *The language.*

```
middleware, ROS
type, block, entity
attribute, block, color
attribute, block, x
attribute, block, y
attribute, block, z
attributeValue, color, red, green, blue
attributeValue, x, INT
attributeValue, y, INT
attributeValue, z, INT
relation, on, block, block
relation, clear, block
action, pickup, loc, grab, raise
action, putdown, loc, release
action, stack, loc, release
action, unstack, loc, grab, raise
```

**Middleware:** Middleware may be different depending on the simulator or robot being used. The syntax for this is “middleware,<name\_of\_middleware>”. ROS and MOOSE are two examples of such middleware and have different code associated with them so the generated code stubs will be different. This keyword allows the API to generate only the wanted code. Currently there is only support for the ROS middleware.

**Type:** As above, a type denotes any sort of entity within the domain. The syntax for this is “type,<name\_of\_object>,<name\_of\_object’s\_parent>”. Entity is enforced as the top level of types, the parent of all within the domain. This enables subcategories to be made for both physical and non-physical objects. Each leaf node in this tree is some sort of definable entity within the domain, thus ROS should have object detector code for it. This is where the language creates object detector code stubs.

**Attribute:** Attributes provide contextual information about an entity within a domain. For the Blocksworld domain, there is the color of the block, along with its x, y, and z coordinates. The syntax for this is “attribute,<name\_of\_object>,<name\_of\_object’s\_attribute>”. Attributes can be assigned to any level of type within the language, all of which pass along to their children’s attributes.

**AttributeValue:** Attribute values are possible values for each attribute. The syntax for this is “attributeValue,<name\_of\_attribute>,<list\_of\_attribute’s\_values>”. This is more concrete with something like colors, where the names may be finite. There are plans to have reserved words for integers (“INT”) but implementing this will be down the line.

**Relation:** Relations describe what are commonly called predicates. The syntax for this is “relation,<name of relation>,<list of entities in the relation>”. In the Blocksworld domain they are *on* and *clear*. *On* specifies two blocks, the first being the block that is stacked on the second block. *On(A,B)* would denote block A is stacked on block B. *Clear* references one block, which does not have anything stacked on it. The language also creates code stubs for determining these predicates within the domain during MIDCA’s Perceive cycle.

**Action:** Actions are used to define interactions with a domain. The syntax for actions is “action,<name of action>,<list of topics associated with action>”. This is part of the influence on the change to attribute values. The name of the action is the high-level operator used in planning and defines a discrete action in the domain. The list of topics is a reference to the ROS architecture of individual nodes exchanging information over named buses called topics. A current example of this in MIDCA is found below. MIDCA is considered a node here, along with `grabbing.py` and Object Detection. The pass information between each other on the ROS level on the `obj_pos`, `loc_cmd`, `grabbing_cmd`, `raise_cmd`, and `release_cmd` topics. These are the basis for the topics listed for the actions in the Blocksworld configuration file above.

### *Templates.*

The MIDCA API generates generic code based on the domain configuration file provided at runtime. The base for these templates can be found in `midca/api/templates`, and they are all initially placed in the domain’s ROS subdirectory after being run. These are code templates as described below, where we will start with object detection on the ROS side.

**Detectors:** (*detectorsTemplate.txt*) The generated file is named based on the domain name, ending up as `<domain-name>Detectors.cpp`. It has code stubs for detecting each leaf node entity defined in the domain’s configuration file. The code below shows an example of this generated function.

```
std::string detect_stone ()
{
    //TODO: create stone_attr string with the following attributes:
    //['craftable', 'breakable']

    return stone_attr;
}
```

From here, the researcher will implement whatever appropriate vision system they decide to detect each specific entity in their domain. In the sample above, this is stone. Because the configuration file for this domain listed craftable and breakable as stone’s attributes (more accurately inherited attributes, more on this later), they are expected to be part of the ROS topic as their status is important. These will be determined from whatever vision system the researcher implements and wrapped in the string `stone_attr`, which is what is passed on the associated topic.

```
ros::Publisher stone_pub = n.advertise<std_msgs::String>("stone_attr", 1);
```

As this is on the ROS side, a publisher is created for this entity, and is later called with the `stone_attr` string from the detector function above.

```
stone_pub.publish(stone_attr);
```

The code this template is based on can be found at:

[https://github.com/COLAB2/baxter\\_pcl/blob/master/src/detect\\_objects.cpp](https://github.com/COLAB2/baxter_pcl/blob/master/src/detect_objects.cpp)

**Entity Handlers:** (*entitiesHandlerTemplate.txt*) The generated file is named based on the domain name, ending up as `<domain-name>EntitiesHandler.py`. This is the code that subscribes to the topics created in the detector code above. Conceptually, this puts it in the “perceive” MIDCA node.

Most of the generated code is boilerplate, with a list of relations (commonly called predicates in planning) to detect in the code, as seen below.

```
#TODO: Code should handle detecting the following relations (predicates):  
#block: ['on', 'clear']
```

There is one line for every leaf node entity in the configuration file, followed by the relations that have been defined for said entities. For a simple example, a Blocksworld domain defines a block with several possible relations. The relation “on” defines block A as on block B if A is stacked directly on top of B. This is written as `on(A, B)`, and should be determined if it is true in the code written by the researcher. A block may be “clear” if there are no blocks placed on top of it, written as `clear(A)` if A is clear and `¬clear(A)` if a block is stacked on A.

The code this template is based on can be found at:

<https://github.com/COLAB2/midca/blob/c3e4cc9fab5683f44435816d90a7a2eb591c4ce4/midca/osrun.py>

**Async:** (*asyncTemplate.txt*) The generated file is named based on the domain name, ending up as `<domain-name>_async.py`. Conceptually, this is located in the “act” node of MIDCA, as it is the publisher to the topics defined in the domain’s configuration file. All unique topics from the Actions within a domain are used, with each high-level action being added to and looked for in the `midcaPlan`’s list of actions. Each action then has a class generated for it based on the `doOperatorTemplate`. These are all found in the resulting `<domain-name>_async.py` file so that template is listed below as a subsection.

- **Action:** (*doOperatorTemplate.txt*) A helper template for generating the action classes within the `entitiesHandlerTemplate`. The initialization function lists all associated topics and provides a place to specify their formats. The `send_topics` function will be where the topic strings are assembled and published.

The code these templates are based on can be found at:

[https://github.com/COLAB2/midca/blob/master/midca/modules/\\_plan/async/async.py](https://github.com/COLAB2/midca/blob/master/midca/modules/_plan/async/async.py)

One difference to note between the `async` template and above source is the addition of a list of topics for each action. This has been added this to extend the range of real-world actions taken by one high level conceptual action. For example, one robot may break the “pickup” action into move, extend, grasp, and retract while another already provides support for the “pickup” action.

TODO comments have been left in the templated code sparingly, to provide a place to start from. For more information on specific implementations of the generated code stubs, check the manual’s provided links to see their sources and how they have been implemented. As implementations may change with time, no template is considered complete and should be completed by the researcher creating the domain. Generated code may not run without interaction and modification. These code stubs are intended as a starting point for new researchers and a way to standardize code between implementations.

*Instructions for running the MIDCA API.*

- 1) Create your `.cfg` file for your domain
- 2) Place your `.cfg` file in `midca/midca/domains/your-new-domain`

- 3) Open up a command line in a terminal.
- 4) Go to midca/midca/api directory (important).
  - a) `cd midca/midca/api`
- 5) Run Midca\_Domain\_API.py.
  - a) `python Midca_Domain_API.py [configuration file path]`  
 The file path may take several forms
    - i) The whole path: `../domains/coloredBlocksworld/coloredBlocksworld.cfg`
    - ii) Partial path: `domains/coloredBlocksworld/coloredBlocksworld.cfg`
    - iii) Partial path: `coloredBlocksworld/coloredBlocksworld.cfg`
- 6) Manage your created templates

Verbose logging was implemented for clearer user interaction and to provide more insights to your domain configuration file. There are 5 distinct levels for this logging, as explained below. Higher levels output their specifics and those of all levels below them.

1. No verbosity: No output, no clutter.
2. `-v`: Prints all edges within the type, attribute, attribute value, and relation trees.

```
Edges of attribute tree: [('block', 'color'), ('block', 'z'), ('block', 'y'), ('block', 'x')]
```

3. `-vv`: Prints the structure of the type, attribute, attribute value, and relation trees.

```
Attribute Tree: defaultdict(<type 'list'>, {'block': ['color', 'z', 'y', 'x']})
```

4. `-vvv`: Displays tree structure with hierarchy of types within the domain. (This assumes entity is the top level parent)

```
+-- entity
  +-- block
```

5. `-vvvv` (or higher): Displays tree structure with hierarchy of types AND their attributes within the domain.

```
+-- entity
  []
  +-- block
    ['color', 'z', 'y', 'x']
```

These attributes are not displayed as being passed down from parents, even though they are inherited in such a way for the detector template. For example, in the Minecraft domain, this is a portion of the max verbosity output.

```
+-- item
  ['craftable']
  +-- block
    ['breakable']
    +-- stone
      []
    +-- grass
      []
```

We see here items have the craftable attribute, blocks have the breakable attribute, but the two specific types of blocks have neither listed under their attributes. These were left out for readability, as large or very precise domains with lots of specific attributes become very bloated and difficult to read. A look at the generated detector code for this domain reveals the accumulated attributes, 'craftable' and 'breakable'. This shows the language follows the principles of inheritance.

```
std::string detect_stone()  
{  
    //TODO: create stone_attr string with the following attributes:  
    //['craftable', 'breakable']  
  
    return stone_attr;  
}
```

## Acknowledgements

This work was funded in part by AFOSR contract #FA2386-17-1-4063, by ONR grant #N00014-18-1-2009, and by NSF grant #1849131.

## References

- Blockeel, H., & De Raedt, L. (1997). [Lookahead and discretisation in ILP](#). *Proc. of the 7th intl. workshop on inductive logic programming* (pp. 77–84) Berlin: Springer
- Cox, M. T. (2007). [Perpetual self-aware cognitive agents](#). *AI Magazine*, 28(1), 32-45.
- Cox, M. T., Alavi, Z., Dannenhauer, D., Eyorokon, V., Munoz-Avila, H., & Perlis, D. (2016). [MIDCA: A metacognitive, integrated dual cycle architecture for self regulated autonomy](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, Vol. 5* (pp. 3712-3718). Palo Alto, CA: AAAI Press.
- Cox, M. T., & Burstein, M. H. (2008). [Case-based explanations and the integrated learning of demonstrations](#). *Künstliche Intelligenz (Artificial Intelligence)* 22(2), 35-38.
- Cox, M. T., Dannenhauer, D., & Kondrakunta, S. (2017). [Goal operations for cognitive systems](#). In *Proceedings of the Thirty-first AAAI Conference on Artificial Intelligence* (pp. 4385-4391). Palo Alto, CA: AAAI Press.
- Cox, M. T., Maynard, M., Paisner, M., Perlis, D., & Oates, T. (2013). [The integration of cognitive and metacognitive processes with data-driven and knowledge-rich structures](#). In *Proceedings of the Annual Meeting of the International Association for Computing and Philosophy*. IACAP-2013.
- Cox, M. T., Oates, T., Paisner, M., & Perlis, D. (2012). [Noting anomalies in streams of symbolic predicates using A-distance](#). *Advances in Cognitive Systems* 2, 167-184.
- Cox, M. T., & Raja, A. (2011a). [Metareasoning: An introduction](#). In M. T. Cox & A. Raja (Eds.) *Metareasoning: Thinking about thinking* (pp. 3-14). Cambridge, MA: MIT Press.
- Cox, M. T., & Ram, A. (1999). [Introspective multistrategy learning: On the construction of learning strategies](#). *Artificial Intelligence*. 112 1-55.
- Dannenhauer, Z. A., & Cox, M. T. (2018). [Rationale-based perceptual monitors](#). *AI Communications*, 31 (2), 197-212.
- Dannenhauer, D., Schmitz, S., Eyorokon, V., Gogineni, V. R., Kondrakunta, S., Williams, T., & Cox, M. T. (2019). MIDCA Version 1.4: User manual and tutorial for the metacognitive integrated dual-cycle architecture (Tech. Rep. No. COLAB<sup>2</sup>-TR-3). Dayton, OH: Wright State University, Collaboration and Cognition Laboratory.
- Edelkamp, S., & Hoffmann, J. (2004). PDDL2.2: The language for the classical part of the 4th International Planning Competition, Tech. Rep. 195, Albert-Ludwigs University at Freiburg, Institut für Informatik (2004).
- Fox, M., & Long, D. (2003). PDDL2.1: An extension to PDDL for expressing temporal planning domains, *Journal of Artificial Intelligence Research* 20, 61–124.

- Helmert, M. (2006). The Fast Downward Planning System. *Journal of Artificial Intelligence Research* 26, 191-246.
- Helmert, M. (2009). Concise finite-domain representations for PDDL planning tasks. *Artificial Intelligence* 173, 503-535. <http://www.informatik.uni-freiburg.de/~ki/papers/helmert-aij2009.pdf>
- Klenk, M., Molineaux, M., & Aha, D. (2013). Goal-driven autonomy for responding to unexpected events in strategy simulations. *Computational Intelligence*, 29(2), 187–206.
- Kondrakunta, S. (2017). [\*Implementation and evaluation of goal selection in a cognitive architecture\*](#). Masters thesis. Computer Science and Engineering Department, Wright State University, Dayton, OH.
- Munoz-Avila, H., Jaidee, U., Aha, D. W., Carter, E. (2010). Goal-driven autonomy with case-based reasoning. I. Bichindaritz & S. Montani (Eds.), *Case-Based Reasoning. Research and Development, 18th International Conference on Case-Based Reasoning, ICCBR 2010* (pp. 228-241). Berlin: Springer.
- Nau, D., Au, T.-C., Ilghami, O., Kuter, U., Murdock, W., Wu, D., & Yaman, F. (2003). [\*SHOP2: An HTN Planning System\*](#). *JAIR*, 20, 379-404.
- Nau, D., Cao, Y., Lotem, A., & Muñoz-Avila, H. (1999). [\*SHOP: Simple Hierarchical Ordered Planner\*](#). In *Proceedings IJCAI-99* (pp. 968-973).
- Nau, D., Muñoz-Avila, H., Cao, Y., Lotem, A., & Mitchell, S. (2001). [\*Total-Order Planning with Partially Ordered Subtasks\*](#). In *IJCAI-2001*. Seattle.
- Paisner, M., Cox, M. T., Maynard, M., & Perlis, D. (2014). [\*Goal-driven autonomy for cognitive systems\*](#). In *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 2085–2090). Austin, TX: Cognitive Science Society.
- Quinlan, J. R. (1990). [\*Learning logical definitions from relations\*](#). *Machine Learning* 5, 239-266.
- Roberts, M., Vattam, S., Alford, R., Auslander, B., Apker, T., Johnson, B., & Aha, D. W. (2015). [\*Goal reasoning to coordinate robotic teams for disaster relief\*](#). In A. Finzi, F. Ingrand, & A. Orlandini (Eds.), *Planning and Robotics: Papers from the ICAPS Workshop*. Palo Alto, CA: AAAI Press.

## Appendix A: Frequently Asked Questions

This appendix lists some commonly asked questions about MIDCA\_1.5 together with answers to those questions.

### MIDCA\_1.5 General Questions

#### How do I get started with MIDCA? Where do I begin?

See the Quick Start section on pages ii-iii.

#### Where can I find more information about MIDCA?

A number of student-initiated workshops have been held on MIDCA. The following URL contains pointers to the web sites for these events and includes slides, recorded oral presentations and workshop proceedings.

<http://www.midca-arch.org/workshops>

See the References section of this manual for further reading. Additionally, the link below contains pointers to many publications pertaining to MIDCA.

[http://www.wright.edu/~michael.cox/colab2/publications\\_copy\\_chronological.html](http://www.wright.edu/~michael.cox/colab2/publications_copy_chronological.html)

#### I modified the source code, but when I run MIDCA it's not detecting the changes.

If you installed MIDCA by doing `python setup.py install` you will need to re-run the command again every time a change is made because MIDCA is using the version stored in python's Lib/site-packages every time it is executed. This is (most likely) different than the location of the source code you're editing. When you run the command `python setup.py install`, one of the things it does is copy the source code from where you are editing it to the Lib/site-packages directory of your python installation. Instead, consider doing `'python setup.py develop'` which creates a symbolic link to your current code directory. Therefore, when you run MIDCA you are using the code you are editing.

## Questions regarding ROS and the Baxter robot

### How do I install ROS?

Follow the instructions from

[sdk.rethinkrobotics.com/wiki/Baxter\\_Setup](http://sdk.rethinkrobotics.com/wiki/Baxter_Setup)

### Which ROS distribution should I install?

Find your linux distribution and choose a compatible ROS distribution. For example, the compatible ROS for Trusty is indigo. If you chose Hydro for Trusty, you will have to install ROS from source and which will lead to dependency issues.

### How do I find the linux distribution?

```
lsb_release -a
```

### How do I create a ROS service?

ROS services are defined by srv files which specifies contains a request message and a response message. Follow this link to create a ROS service:  
[wiki.ros.org/ROS/Tutorials/CreatingPackage](http://wiki.ros.org/ROS/Tutorials/CreatingPackage)

example: `roscpp create_pkg baxter_cog std_msgs roscpp`

You need to adjust your CMakeLists.txt file; you have to uncomment `roscpp_genmsg()` and `roscpp_gensrv()`

### How do I install Opencv?

[docs.opencv.org/doc/tutorials/introduction/linux\\_install/linux\\_install.html#linux-installation](http://docs.opencv.org/doc/tutorials/introduction/linux_install/linux_install.html#linux-installation)

[opencv.org](http://opencv.org)

### How do I install Moveit?

Follow instructions from <https://github.com/RethinkRobotics/sdk-docs/wiki/MoveIt-Tutorial>. Baxter using moveit was detecting a collision between its hand camera and gripper base, as well as `[right/left]_hand_range` and `gripper_base`, `gripper` and `gripper base`. Added exceptions to ignore those to `baxter_ws/src/moveit_robots/baxter/baxter_moveit_config/config/baxter.srdf`

### How do I install ORK?

Follow instructions from

[wgperception.github.io/object\\_recognition\\_core/install.html#install](http://wgperception.github.io/object_recognition_core/install.html#install)

## Appendix B: List of Modules

The following is a list of all modules in MIDCA Version 1.5.

```
ADistanceAnomalyNoter (file:modules/note.py)
AsynchronousAct (file:modules/act.py)
DeliverGoal (file:modules/guide.py)
EvalPointingFromFeedback (file:modules/evaluate.py)
GenericPyhopPlanner (file:modules/planning.py)
JSHOP2Planner (file:modules/planning.py)
JSHOPPlanner (file:modules/planning.py)
MAQuery (file:modules/assess.py)
MAREporter (file:modules/perceive.py)
NBeaconsGoalGenerator (file:modules/guide.py)
NBeaconsSimpleAct (file:modules/act.py)
PerfectObserver (file:modules/perceive.py)
PerfectObserverWithThief (file:modules/perceive.py)
PyHopPlanner (file:modules/planning.py)
ReactiveApprehend (file:modules/guide.py)
SimpleAct (file:modules/act.py)
SimpleEval (file:modules/evaluate.py)
SimpleEval2 (file:modules/evaluate.py)
SimpleIntend (file:modules/intend.py)
SimpleMortarGoalGen (file:modules/guide.py)
SimpleNBeaconsExplain (file:modules/assess.py)
SimpleNBeaconsGoalManager (file:modules/guide.py)
StateDiscrepancyDetector (file:modules/note.py)
TFFire (file:modules/guide.py)
TFStack (file:modules/guide.py)
UserGoalInput (file:modules/guide.py)
WarehouseIntend (file:modules/intend.py)
```

[Perhaps not complete list. For instance metamodules are missing.]