Humans Teaching Robots: Challenges to Decoding the Intention Behind Natural Instruction

IJCAI 2011 Workshop on Agents Learning Interactively from Human Teachers (ALIHT)
Barcelona, Spain

Presenter: Tasneem Kaochar

Work done in collaboration with Raquel Torres Peralta, Ian R. Fasel, Clayton T. Morrison, Thomas J. Walsh and Paul R. Cohen

THE UNIVERSITY OF ARIZONA®

Human-Instructable Computing

<u>Research Focus:</u> build an intelligent agent that is capable of learning a task from a naïve human teacher

- Complex, multi-tasking intelligent devices will soon become ubiquitous in the home and workplace
  - Examples: household robot, networked home entertainment system,
- Such devices will be interacting daily with untrained and naïve human users
- Users may wish to extend or customize a device's capabilities beyond its factory-manufactured settings

Human-Instructable Computing

- To build a capable electronic student we need to first understand how humans teach
  - “Natural” human teaching is dynamic, interactive and much less structured than formal programming
  - We want to bridge the gap between human natural instruction methods and machine learning algorithms
- We performed an exploratory study using Wizard of Oz protocol to better understand human teaching patterns

How Do Machines Learn?

Imagine you wanted to teach a robot to help you clean the dishes. How might you teach the robot? **How might the robot learn?**

- through ***concept*** definitions
  - Robot can learn the distinction between objects (such as a *cup* and a *plate)* based on observed characteristics of each object

- by observing ***demonstration*** of how to perform a task
  - Robot can watch how you (teacher) place the dishes into the dishwasher and attempt to imitate

- by using teacher ***feedback***
  - to reinforce learning with a numerical value (or simply a thumbs up or down)
  - to explain what went wrong, i.e., critique

Can the modes of interaction of machine learning (examples, demonstrations, feedback) be a basis for natural instruction?

Next Step: Build a teaching interface that allows a human teacher to provide natural instruction to an electronic student using the modes of interaction from machine learning?

BLUI: Bootstrapped Learning User Interface

- Domain: X-Plane simulated flying environment
- Student is the control system of a simulated unmanned aerial vehicle (UAV) that will be taught to carry out missions
- UAV is equipped with 3 sensors: wide-range camera, high-resolution camera and radiation sensor

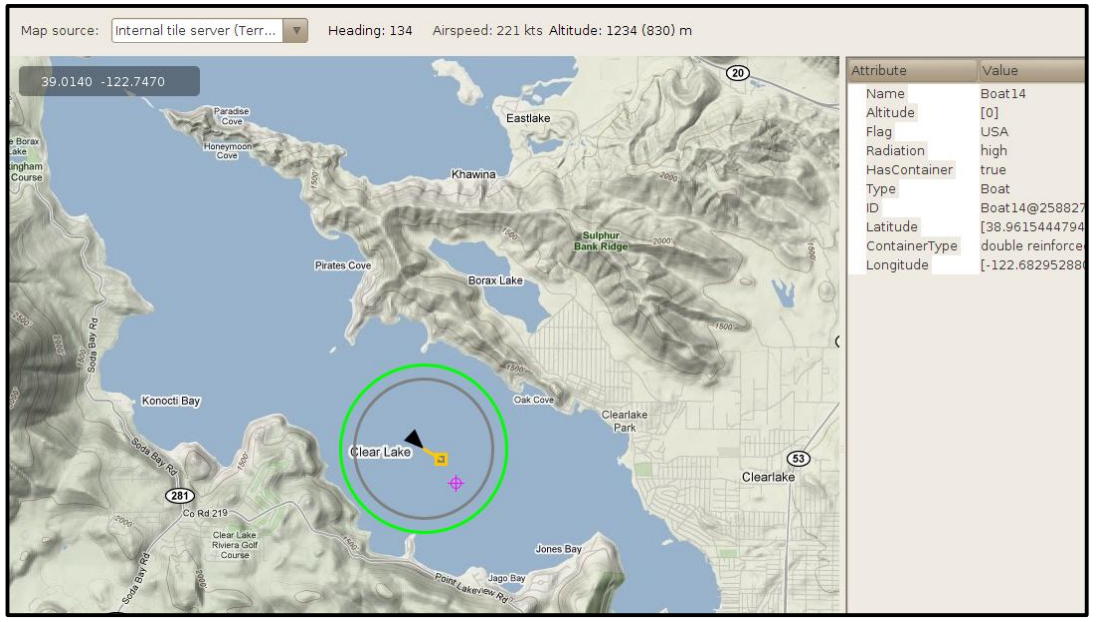

UAV with three sensor ranges displayed: wide-range camera in gray, high resolution camera in yellow and radiation sensor in green.

[1]X-Plane Laminar Research: http://www.x-plane.com

BLUI: Teaching & Testing Facilities

Four modes of instruction:

- *Teaching concepts by example* – using the object labeling facility
- *Teaching by demonstration* – using the procedure demonstration facility (positive and negatives traces of a demonstration can be given)
- *Teaching by feedback* (positive and negative feedback can be provided)
- *Testing the Student*

Note: A free text chat facility was also provided to teachers for use in case they were unable to convey instruction to Student using existing teaching tools

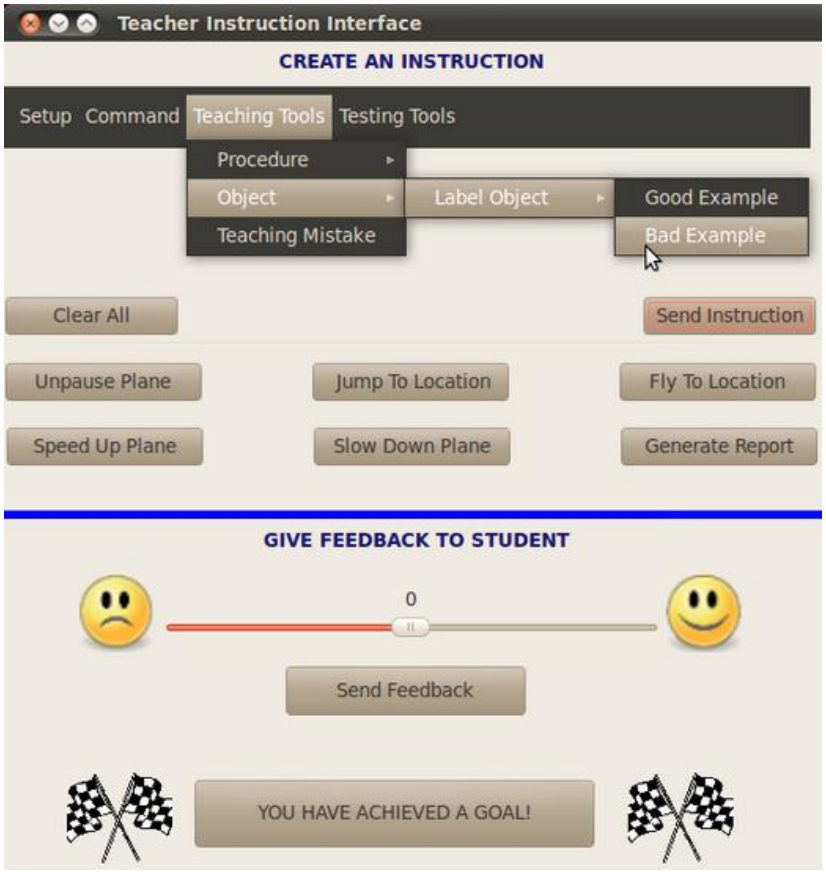

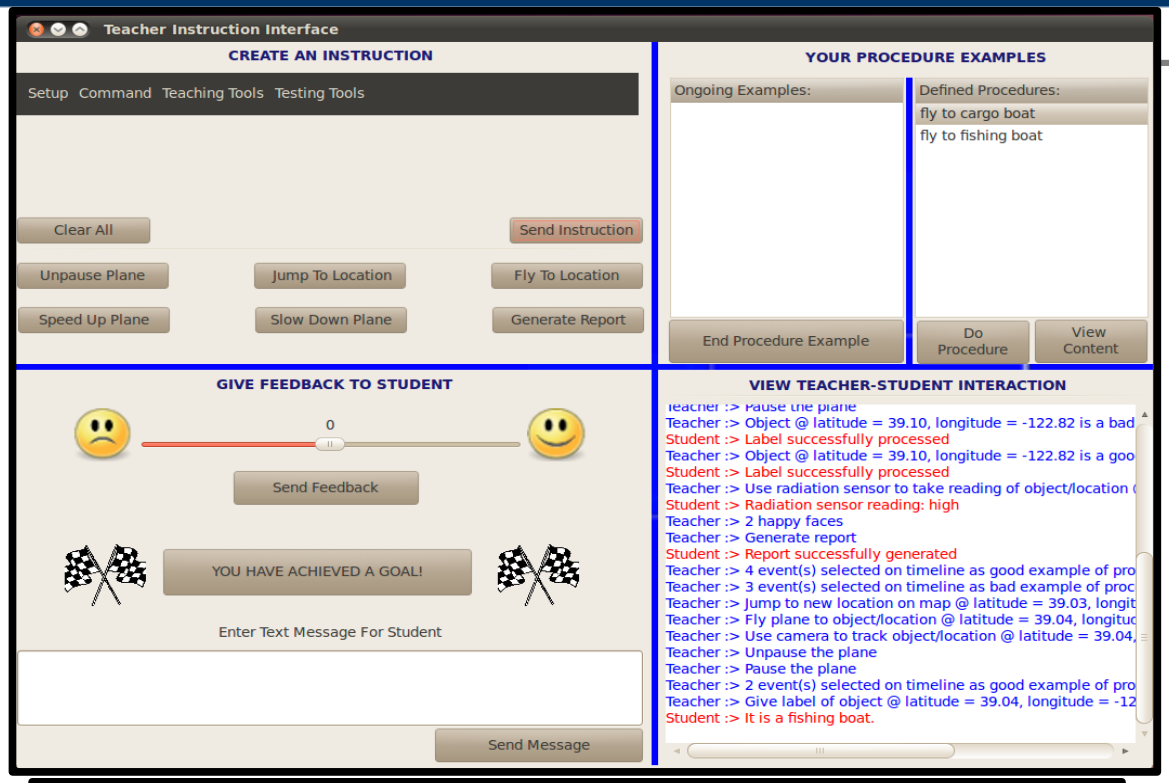

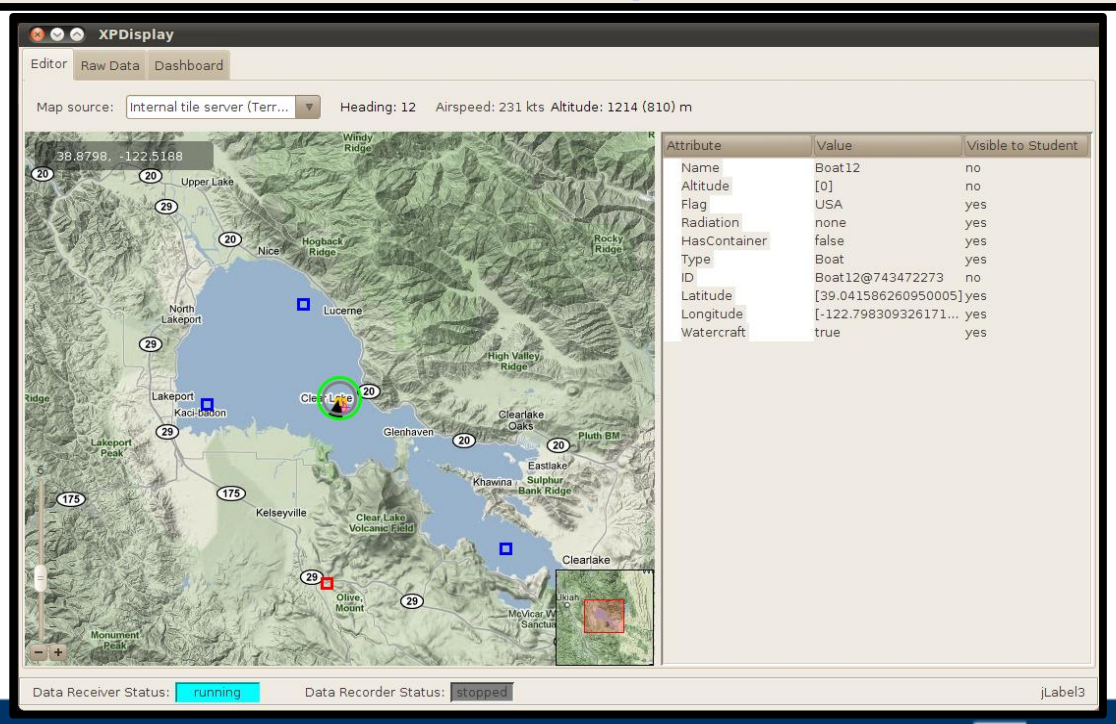

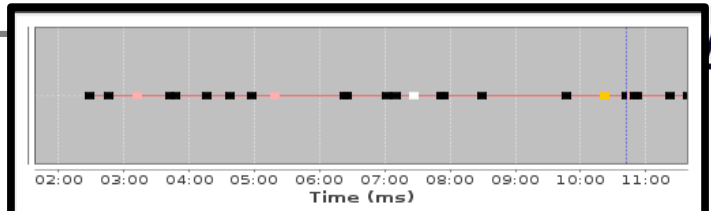

COMMAND	PARAMETERS	TIMESTAMP
Load New Scenario	File: Clear Lake 1; Instruction:...	02:28:790
Fly plane to object/location	Name: Boat11; ID: Boat11@1...	02:45:724
Use camera to track object/lo...	Name: Boat11; ID: Boat11@1...	03:12:742
Unpause Plane	Instruction: Unpause the pla...	03:41:831
Pause Plane	Instruction: Pause the plane;	03:47:256
Label Good Object Example	Name: Boat11; ID: Boat11@1...	04:16:698
Label Bad Object Example	Name: Boat11; ID: Boat11@1...	04:37:347
Fly plane to object/location	Name: Boat10; ID: Boat10@1...	04:58:005
Use camera to track object/lo...	Name: Boat10; ID: Boat10@1...	05:18:488
Unpause Plane	Instruction: Unpause the pla...	06:22:186
Pause Plane	Instruction: Pause the plane;	06:24:978
Label Bad Object Example	Name: Boat10; ID: Boat10@1...	07:01:061
Label Good Object Example	Name: Boat10; ID: Boat10@1...	07:10:230
Use radiation sensor to take ...	Name: Boat10; ID: Boat10@1...	07:26:903
Teacher Feedback	Feedback: 2 happy faces; La...	07:51:896
Generate Report	Instruction: Generate report;	07:53:792
Label Good Procedure Example	Step 1: Fly plane to object/loc...	08:29:355
Label Bad Procedure Example	Step 1: Fly plane to object/loc...	09:47:369
Reset to new location	Name: ; ID: ; Latitude: 39.031...	10:22:410
Fly plane to object/location	Name: ; ID: ; Latitude: 39.039...	10:41:940
Use camera to track object/lo...	Name: Boat12; ID: Boat12@7...	10:47:412
Unpause Plane	Instruction: Unpause the pla...	10:50:124
Pause Plane	Instruction: Pause the plane;	10:52:412
Label Good Procedure Example	Step 1: Fly plane to object/loc...	11:21:654
Give label of object	Name: Boat12; ID: Boat12@7...	11:37:735

BLUI: Teacher's View

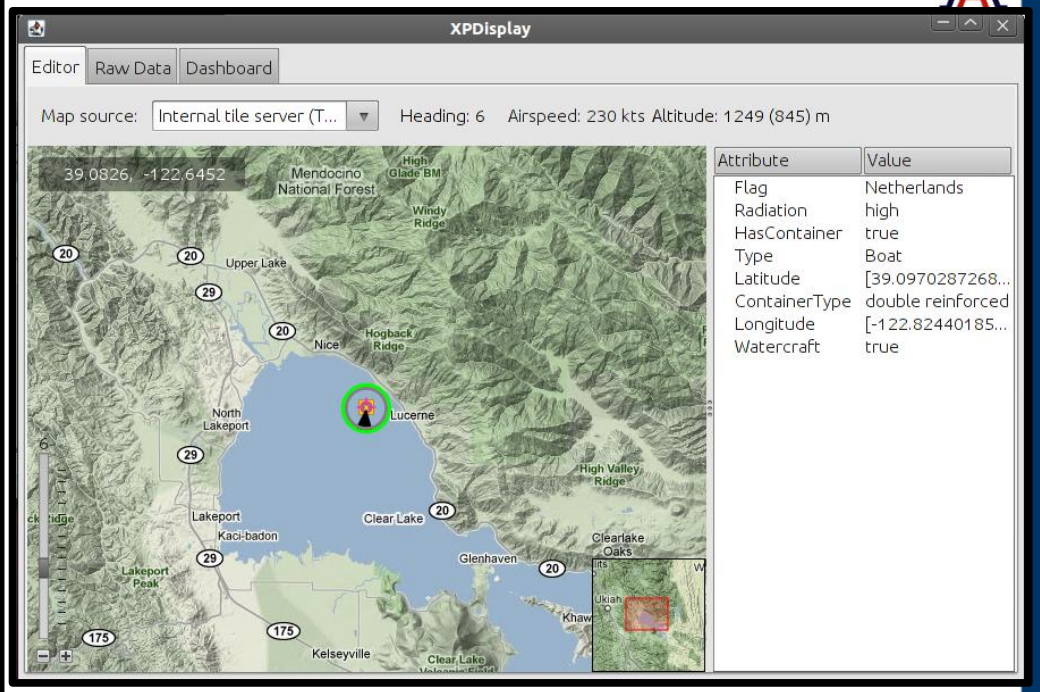

BLUI: Student's View

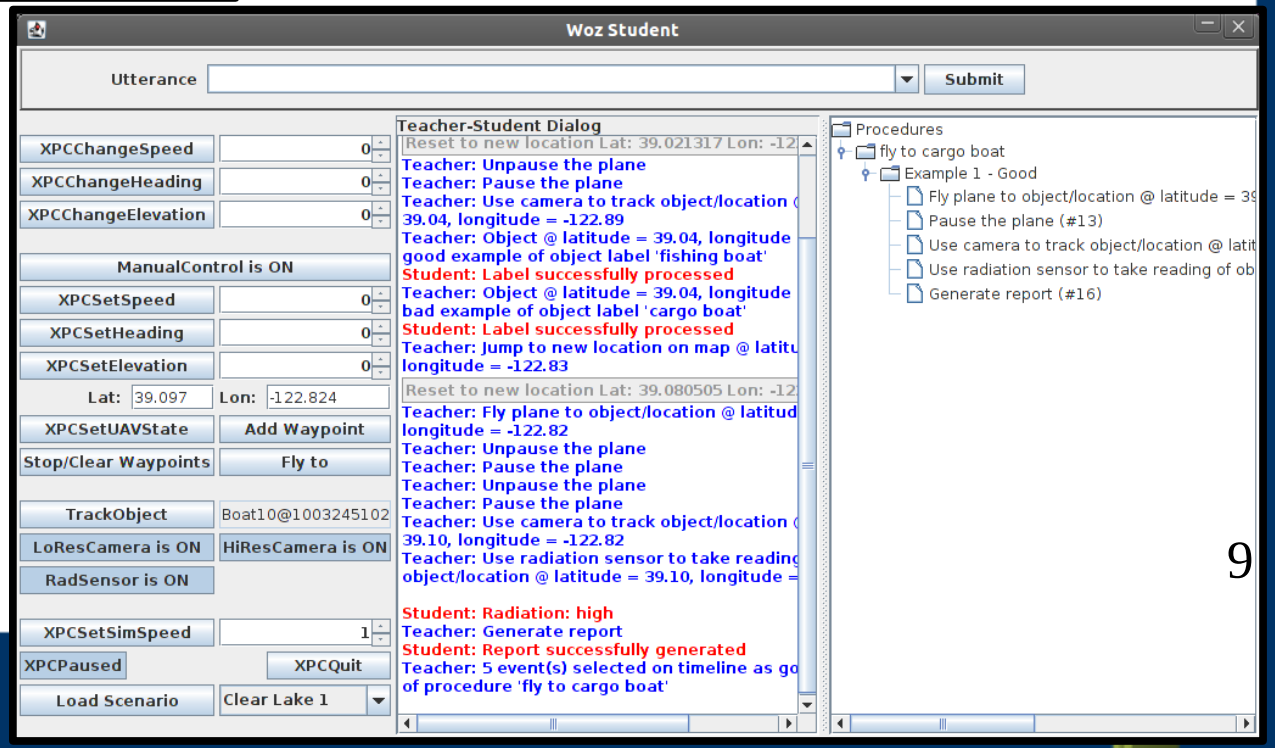

Wizard of Oz (WoZ) Behavioral Study

- We want to learn *how* humans would teach *if* they believed that they were interacting with a capable electronic student
- We perform an exploratory study using a Wizard of Oz paradigm
  - Human teacher participant believes he/she is interacting with a capable electronic student, who in reality is being controlled by another human (without the teacher's knowledge)

BLUI WoZ Study

- 44 non-expert human participants (UA students)
- Teaching task: teach Student to identify all cargo boats in a specified body of water. Once a cargo boat has been identified, the Student must take its radiation sensor reading and generate a report.
  - Teach concepts – cargo and fishing boats
  - Teach procedure – use radiation sensor only on cargo boats and generate a report of the readings
- Each participant spent at least 20 minutes interacting with simulated electronic student

BLUI WoZ Study: Overview of Results

- Human teaching patterns:
  - Evidence of bootstrapping in teaching
  - Testing becomes more important as teaching session progresses
  - Teach-test-feedback is very common
  - Implicit object labeling
  - Implicit procedure definition
  - Ill-defined procedure boundaries
  - Consistent naming conventions

Teachers begin session by defining object concepts

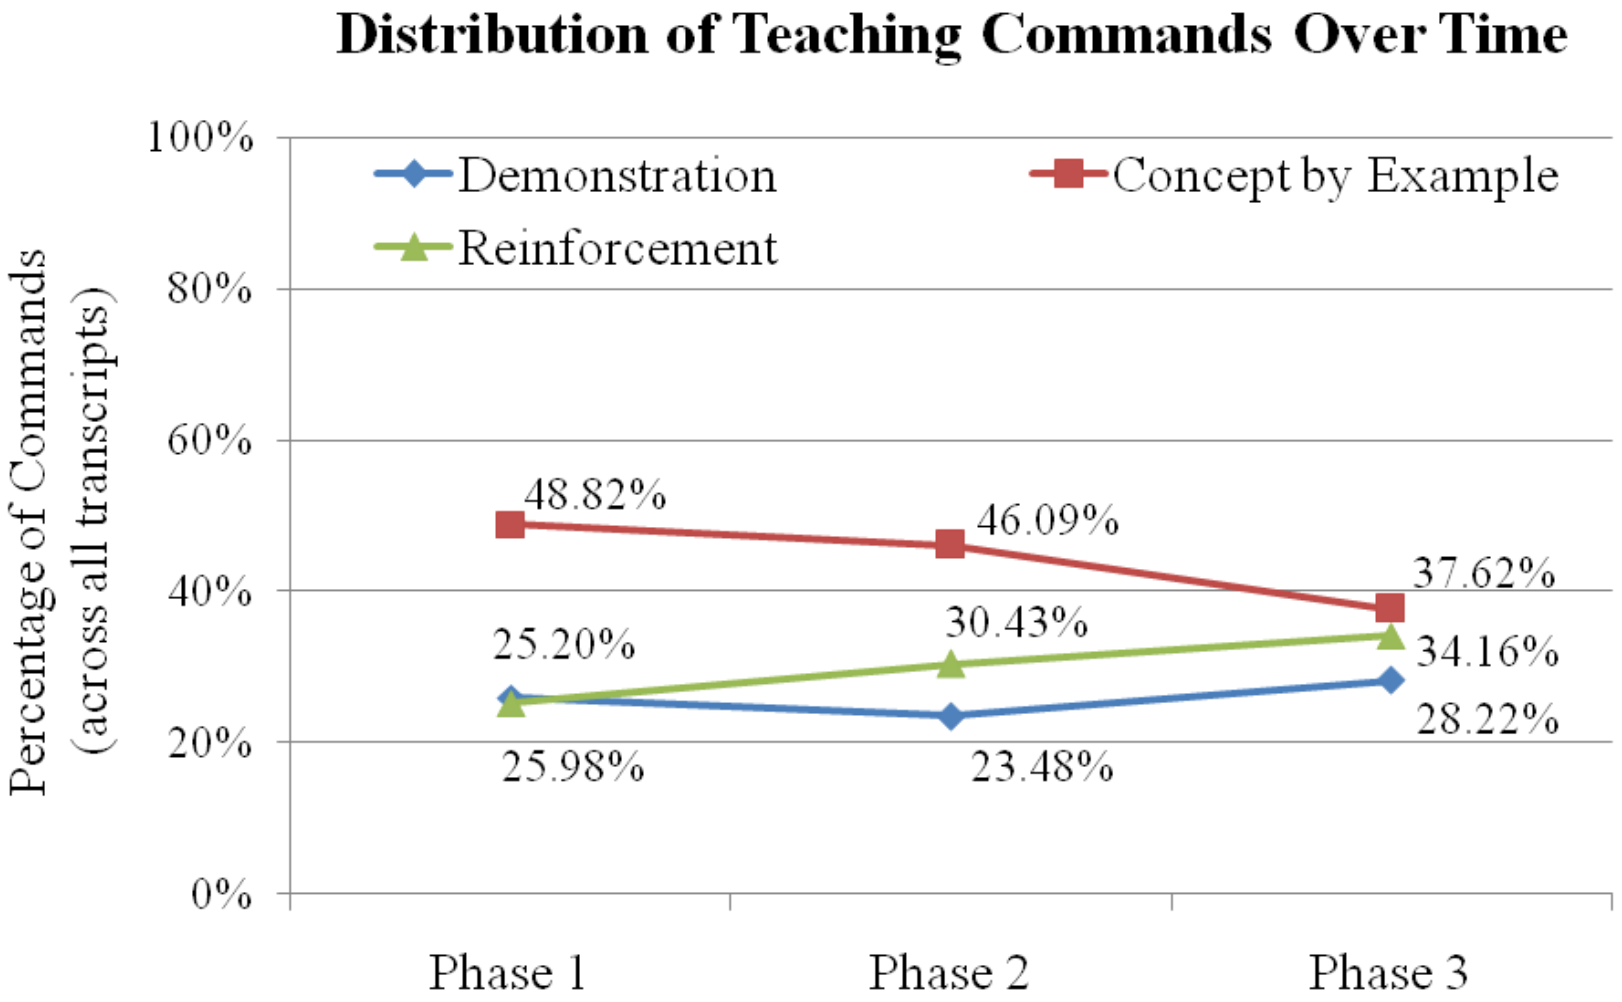

Note: All 44 teaching session data was split into 3 equal time phases

Testing becomes more important as teaching session progresses

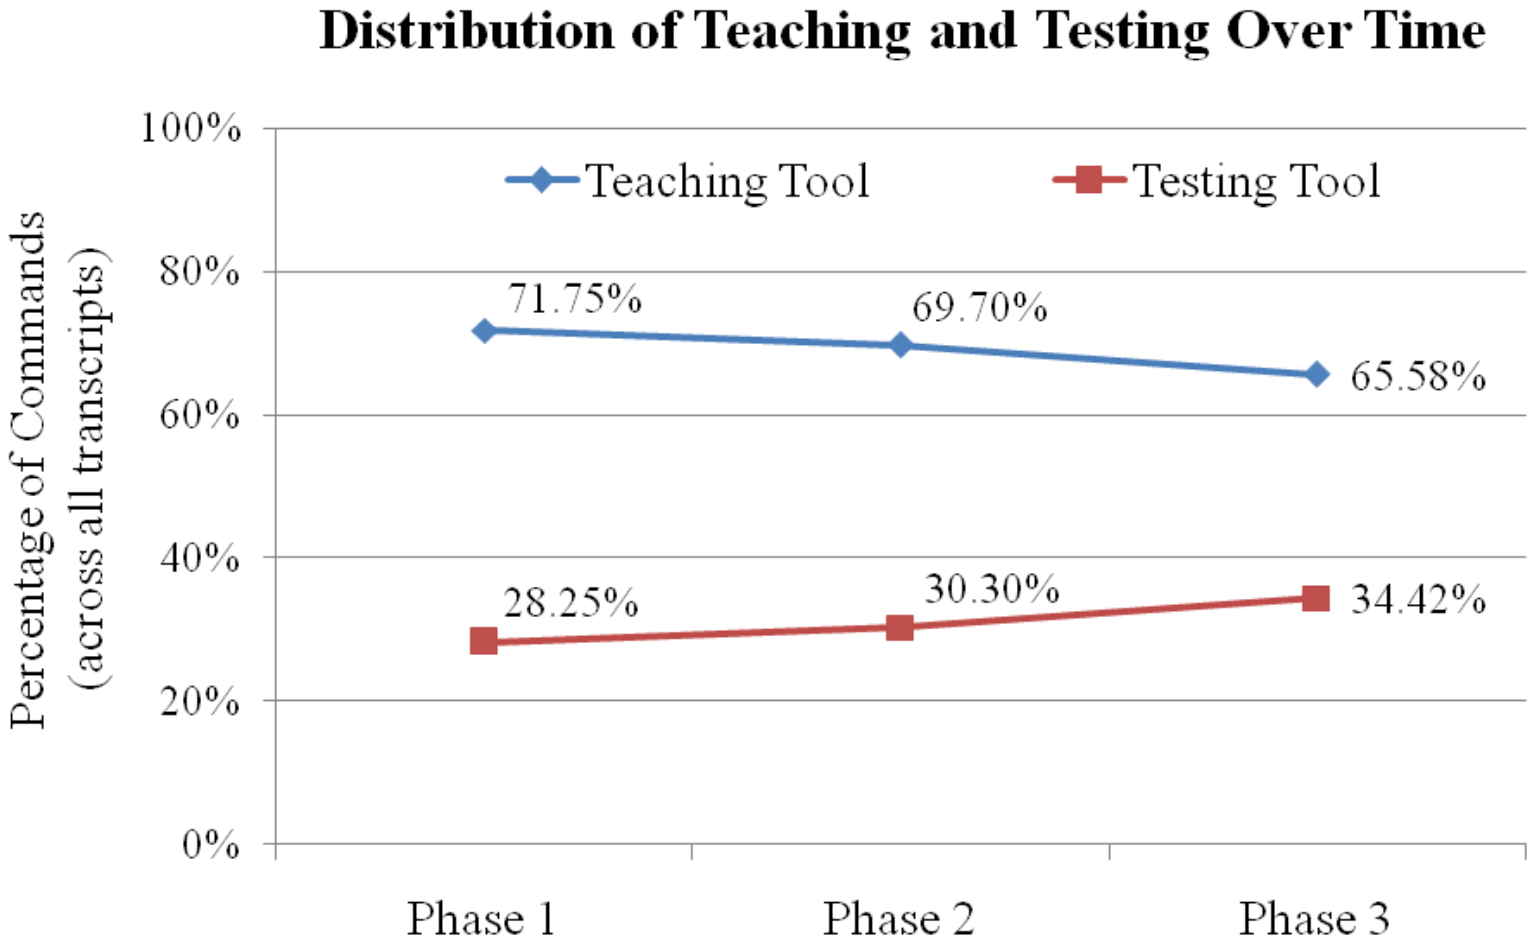

Note: All 44 teaching session data was split into 3 equal time phases

Teaching-Testing-Feedback is common pattern

Teaching – Testing – Feedback Loop

130: T: Start good example of procedure 'fly to cargo boat'

131: T: Fly to object at lat 38.62, long. -120.12

...

159: T: End example of procedure 'fly to cargo boat'

160: T: Perform procedure 'fly to cargo boat' near lat. 39.10, long. -122.82

...

164: S:Radiation sensor reading: high

165: T: You achieved goal 'find cargo boat'

166: T: 1 happy face

Start procedure demonstration

...procedure steps

End procedure demonstration

Test procedure comprehension in a new scenario location

Positive feedback provided

Patterns in Object Labeling

```
12: 12:32 T: Object @lat. = 39.04, long. = -122.89
    is a good example of object label 'Cargo Boat'
    (Object name = Boat11)
```

Explicit Object Labeling

```
39: 43:29 T: Start good example of procedure 'Cargo Boat'
40: 43:38 T: Fly plane to object/location @ lat= 39.10,
             lon= -122.82 (Object name = Boat10)
43: 44:52 T: End example of procedure 'Cargo Boat'
```

Implicit Object Labeling

Patterns in Procedure Definitions

```
125: 05:10 T: Start good example of procedure 'fly to cargo boat'
127: 05:17 T: Fly plane to lat = 39.10, lon = -122.82
    (…UAV heading towards destination…)
130: 06:45 T: Use camera to track object @ lat= 39.10,
              lon = -122.82 (Object name = Boat10)
    (…UAV reached destination…)
131: 06:47 T: Pause the plane
134: 07:12 T: Turn on radiation sensor
136: 07:20 T: Use radiation sensor to take reading of object @
               lat = 39.10, lon = -122.82 (Object name = Boat10)
137: 07:36 T: Unpause the plane
143: 08:06 T: End example of procedure 'fly to cargo boat'
```

Well-defined procedure boundary

```
50: 29:542 T: Fly plane to lat = 39.04, lon = -122.80
   (...UAV heading towards destination...)
63: 29:603 T: Turn on high resolution camera
66: 29:609 T: Use camera to track object @ lat = 39.04,
              lon = -122.80 (Object name = Boat12)
   (...UAV nearing destination...)
72: 29:807 T: Start bad example of proc. 'Fly to cargo boat'
73: 29:815 T: Unpause the plane
77: 29:821 T: Pause the plane
   (...UAV reached destination...)
79: 29:831 T: End example of procedure 'Fly to cargo boat'
```

Important commands excluded from procedure boundary specification

Ill-defined procedure boundary

Patterns in Procedure Definitions (cont.)

```
137: 05:17 T: Fly plane to lat = 39.10, lon = -122.82
    (…UAV heading towards destination…)
140: 06:45 T: Use camera to track object @ lat= 39.10,
              lon = -122.82 (Object name = Boat12)
    (…UAV reached destination…)
141: 06:47 T: Pause the plane
144: 07:12 T: Turn on radiation sensor
146: 07:20 T: Use radiation sensor to take reading of object @
               lat = 39.10, lon = -122.82 (Object name = Boat12)
```

Implicit procedure definition

Consistent Naming Conventions

- Human teacher participants used meaningful naming conventions when providing labels for object concepts and procedures
  - Names derived from vocabulary of the task domain
    - *cargo boat', 'fish boat', 'fishing boat', 'fly to cargo boat', 'scan boat'*
  - Identifying verb phrases versus noun phrases can help identify when procedure definition facility was used for object labeling
    - *'fly to cargo boat'* versus *'cargo boat'*

Most teachers are unstructured in their teaching

We categorized our 44 teachers based on the organization of Teacher-Student interaction transcripts

Structured teachers (16%)

- Used the interface's object labeling facility to teach object concepts – no implicit object labeling
- Used the procedure demonstration facility to define procedures – well-defined procedure boundaries
- Tested only on previous lessons

Semi-structured teachers (50%)

- Tested on previous lessons
- Explicit and implicit object labeling

Free-style teachers (34%)

- Testing before teaching
- Explicit and Implicit labeling
- Ill-defined procedure boundaries

What we learned…

- ***Humans can teach by demonstration, concept definitions and feedback***, which is good news because these are the modes of interaction from which ML algorithms can learn
  - Teachers rarely used the free text chat facility to instruct the Student
- When the Student "acted smart" and competent, the ***majority of teachers were pretty sloppy and unorganized***.
- However, despite the unstructured teaching style of most teachers, ***patterns in teaching do emerge and may be used to automatically extract teacher intentions***

Next Step: Translate NIMs into ML Algorithms

Natural Instruction Methods (NIMs)

Machine Learning Algorithms

- Teachers interchange modes of interaction without notification
- Often times instruction is implicit

- Precision
- Structure

Automatic labeling/learning systems from natural instruction

Complete end-to-end system

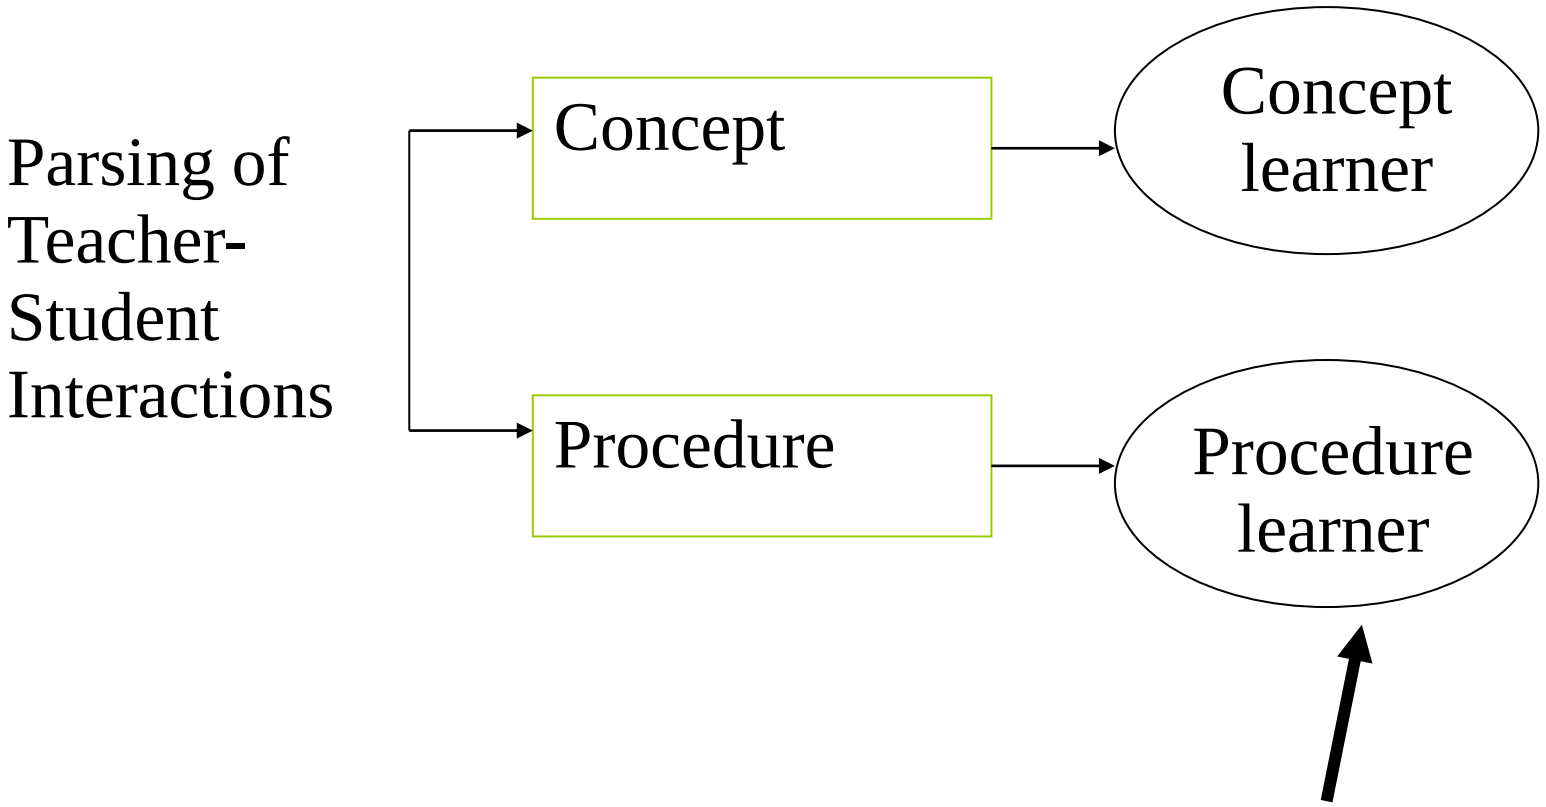

Automatic Transcript Annotation

What we can do now:

Label	Labeling Heuristic	Script Successful
Explicit Object Label	Teaching command defining good or bad object example using the *object labeling construct*	√
Implicit Object Label	Teaching command defining good or bad object example using the *procedure construct*	X
Explicit Procedure Definition	Teaching command defining a procedure using the *procedure construct*	√
Implicit Procedure Definition	Repetition of similar set of action commands (camera track, take radiation reading or picture, generate report) in different scenario locations	X
Test	Testing command (ask Student to perform a previously taught procedure or give the label of a world object)	√
Feedback	Command evaluating Student's performance (1-3 *happy* or *frowny* faces)	√
Goal Specification	Command labeling a goal that Student has achieved	√
Setting Up for Teaching	All setup commands (change in location, speed, altitude or sensor settings) preceding a teaching command (object labeling or procedure definition)	√ *
Setting Up for Testing	All setup commands preceding a testing command	√

Automatic labeling/learning systems from Natural Instruction

What we need to do next:

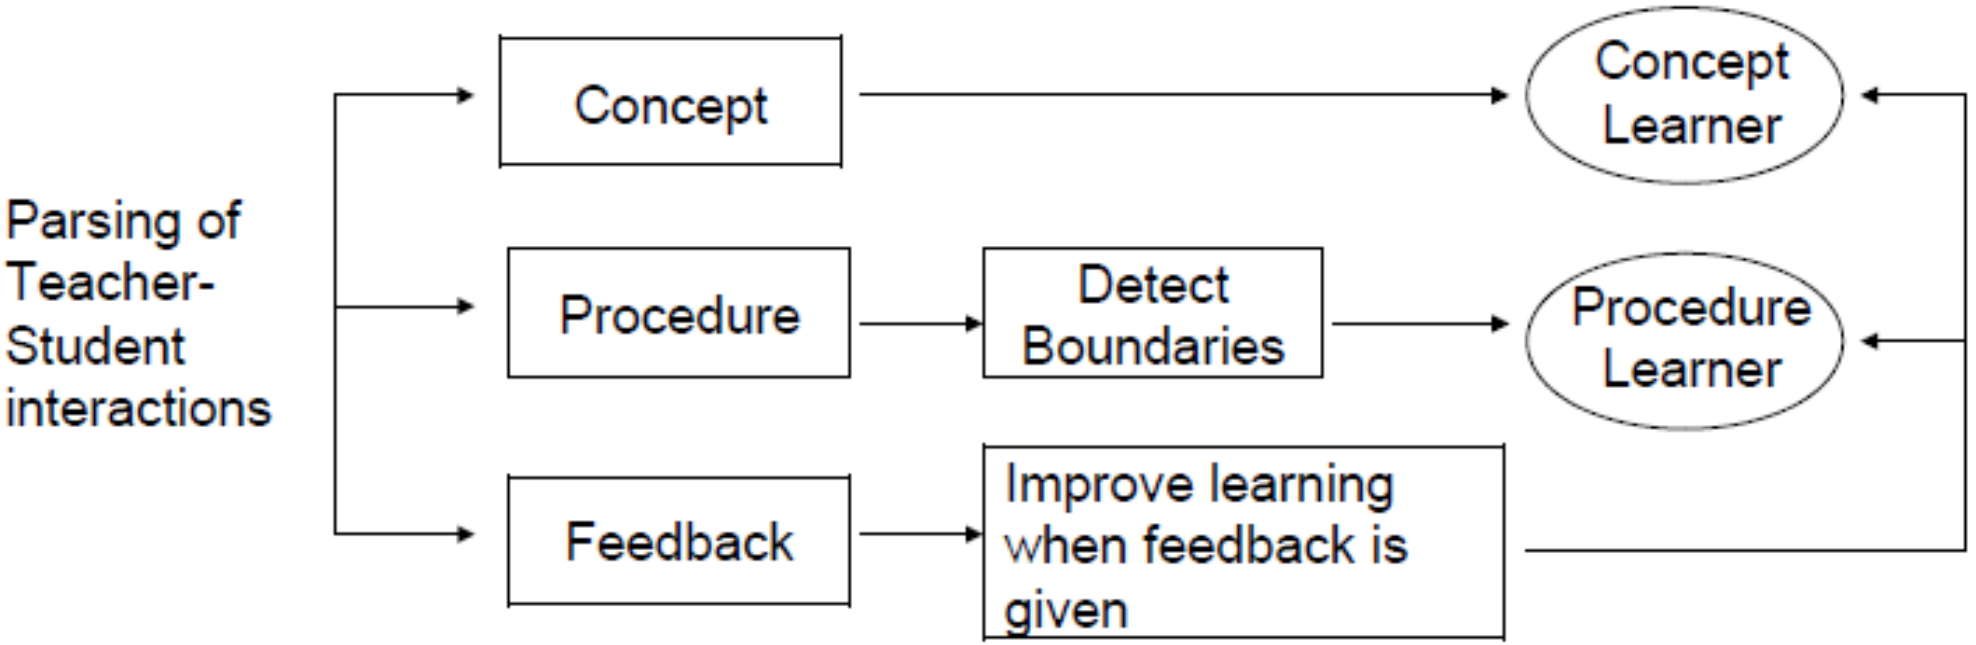

Detect Concept and Procedure Definitions (Explicit and Implicit)

STILL A LOT OF WORK TO BE DONE!