Learning Language from Perceptual Context

David Chen

Supervising Professor: Raymond J. Mooney

Doctoral Dissertation Proposal

December 15, 2009

Semantics of Language

- The *meaning* of words, phrases, etc
- Learning semantics of language is one of the ultimate goals in natural language processing
- The meanings of many words are grounded in our perception of the physical world: red, ball, cup, run, hit, fall, etc. [Harnad, 1990]
- Computer representation should also be grounded in real world perception

Grounding Language

Spanish goalkeeper Casillas blocks the ball

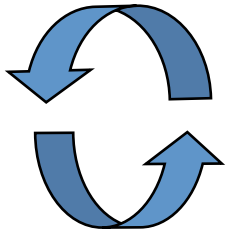

Block(*Casillas*)

Grounding Language

Spanish goalkeeper Casillas blocks the ball

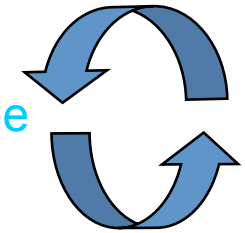

Block(*Casillas*)

Natural Language and Meaning Representation

Spanish goalkeeper Casillas blocks the ball

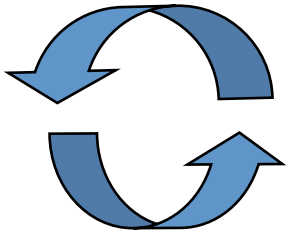

Block(*Casillas*)

Natural Language and Meaning Representation

Natural Language (NL)

Spanish goalkeeper Casillas blocks the ball

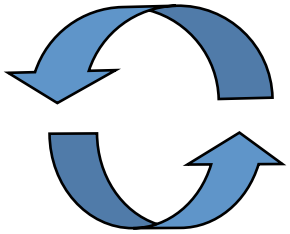

Block(*Casillas*)

NL: A language that has evolved naturally, such as English, German, French, Chinese, etc

Natural Language and Meaning Representation

Natural Language (NL)

Meaning Representation Language (MRL)

Spanish goalkeeper Casillas blocks the ball

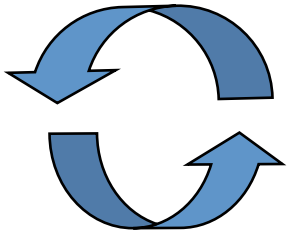

Block(*Casillas*)

NL: A language that has evolved naturally, such as English, German, French, Chinese, etc

MRL: Formal languages such as logic or any computer-executable code

Semantic Parsing and Tactical Generation

NL **MRL**

Spanish goalkeeper Casillas blocks the ball

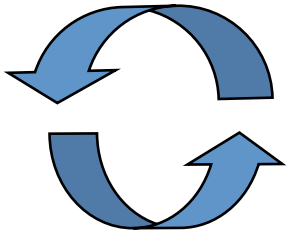

Block(*Casillas*)

Semantic Parsing (NL → MRL)

Semantic Parsing: maps a natural-language sentence to a complete, detailed semantic representation

Semantic Parsing and Tactical Generation

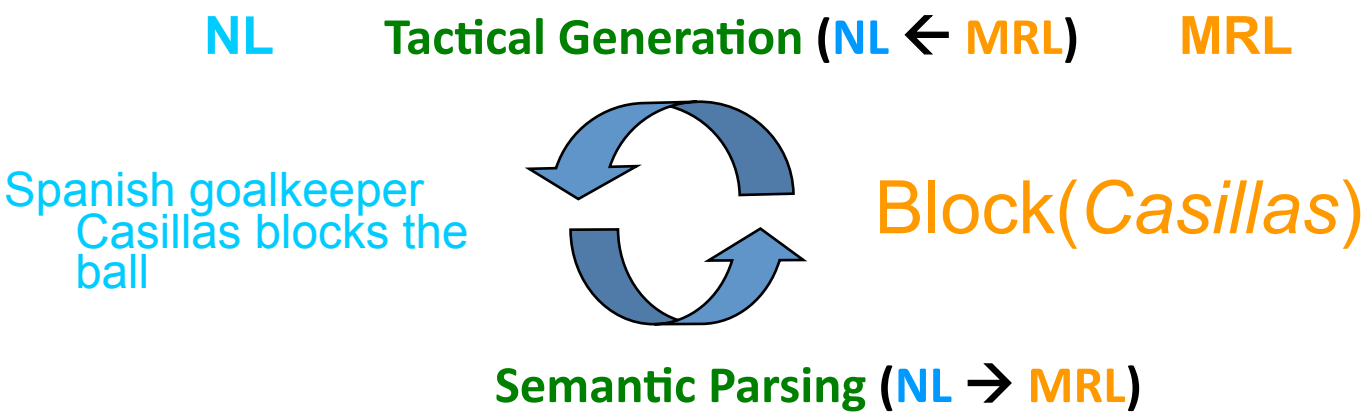

Semantic Parsing: maps a natural-language sentence to a complete, detailed semantic representation

Tactical Generation: Generates a natural-language sentence from a meaning representation.

Strategic Generation

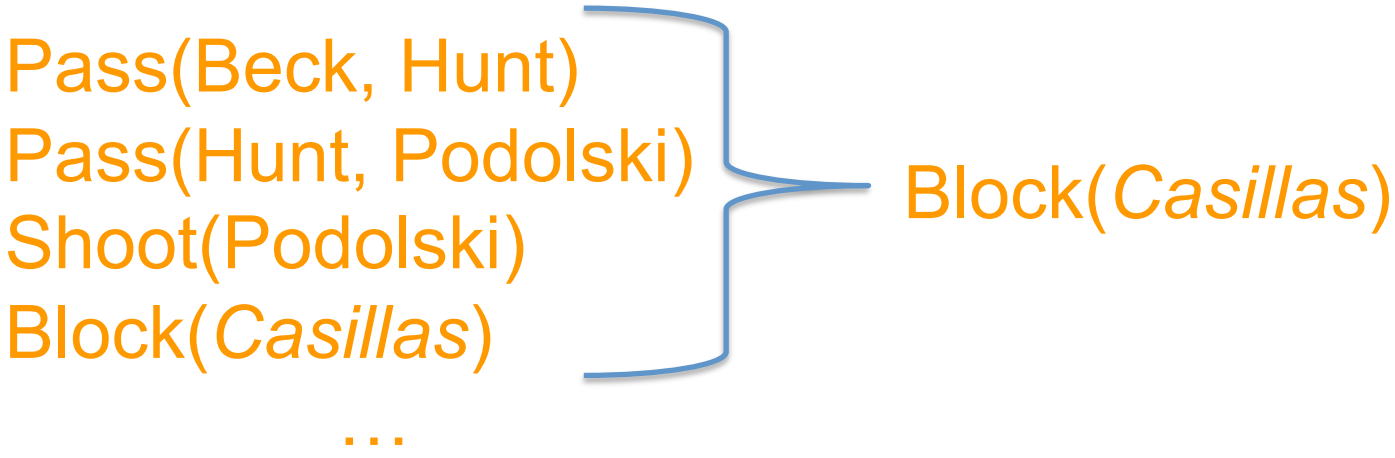

- **Strategic Generation (Content Selection):** Given a set of meaning representations, select a subset

Applications

- Natural language interface
  - Issue commands and queries in natural language
  - Computer responds with answer in natural language
- Knowledge acquisition
- Computer assisted tasks

Traditional Learning Approach

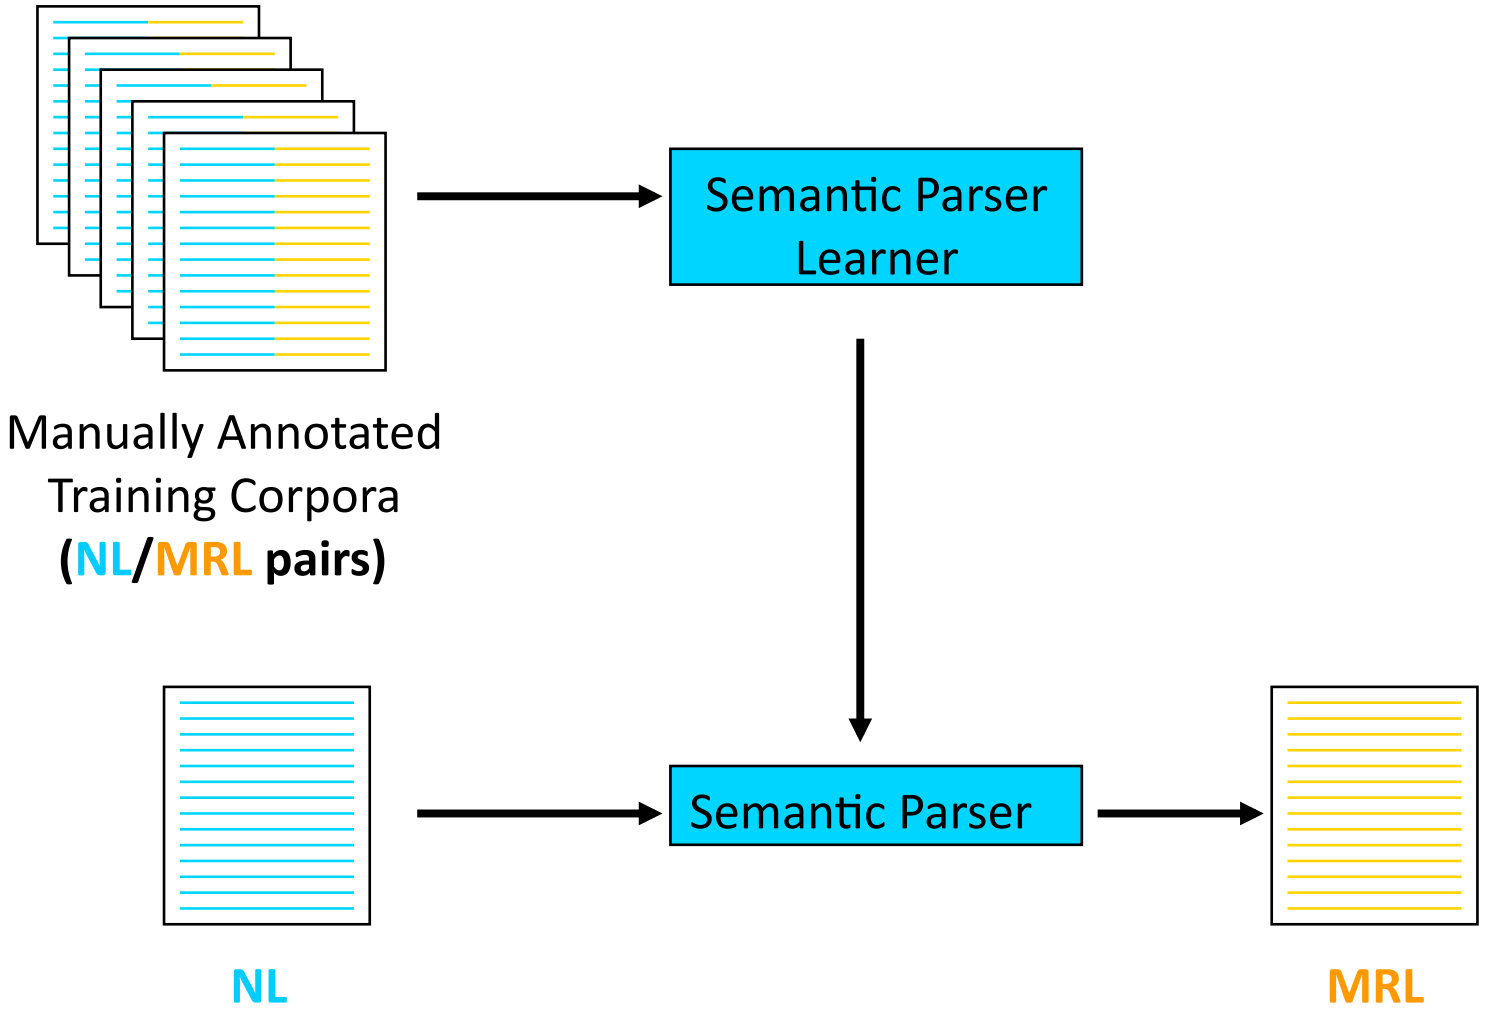

Traditional Learning Approach

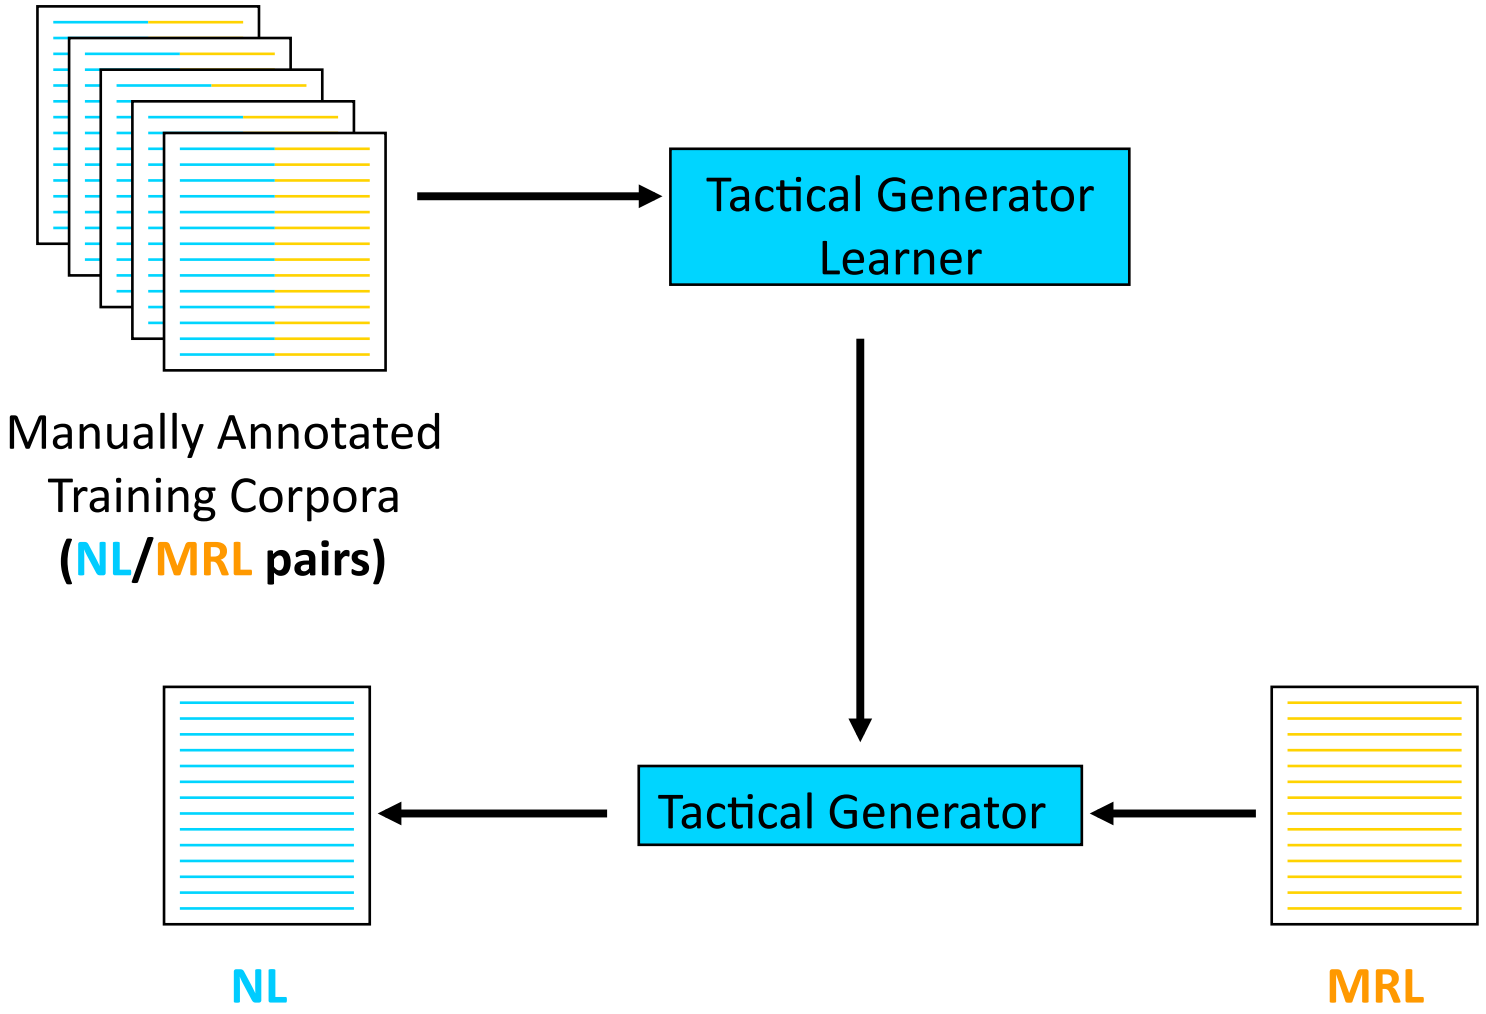

Example of Annotated Training Corpus

Natural Language (NL)	Meaning Representation Language (MRL)
Alice passes the ball to Bob	Pass(Alice, Bob)
Bob turns the ball over to John	Turnover(Bob, John)
John passes to Fred	Pass(John, Fred)
Fred shoots for the goal	Kick(Fred)
Paul blocks the ball	Block(Paul)
Paul kicks off to Nancy	Pass(Paul, Nancy)
…	…

Learning Language from Perceptual Context

- Constructing annotated corpora for language learning is difficult
- Children acquire language through exposure to linguistic input in the context of a rich, relevant, perceptual environment
- Ideally, a computer system can learn language in the same manner

Learning in Virtual Environment

- Many schools use 3D virtual environments to support language learning
  - Immersive: Surrounded by a stimulating environment
  - Social: Language learners can interact with others
  - Creative: Constructing objects as part of learning
- Online worlds including Second Life
- Different ways of learning
  - Task-based learning
  - Collaborative construction
  - Virtual tourism

Learning in Virtual Environment

- Growing video game industry
  - $9.5 billion in the US in 2007, $11.7 billion in 2008 (Entertainment Software Association annual report)
- Serious games
  - DARWARS: Military training systems
  - SimPort: Simulated construction and management of a sea port project

Navigation Example

Alice: 在餐廳的地方右轉

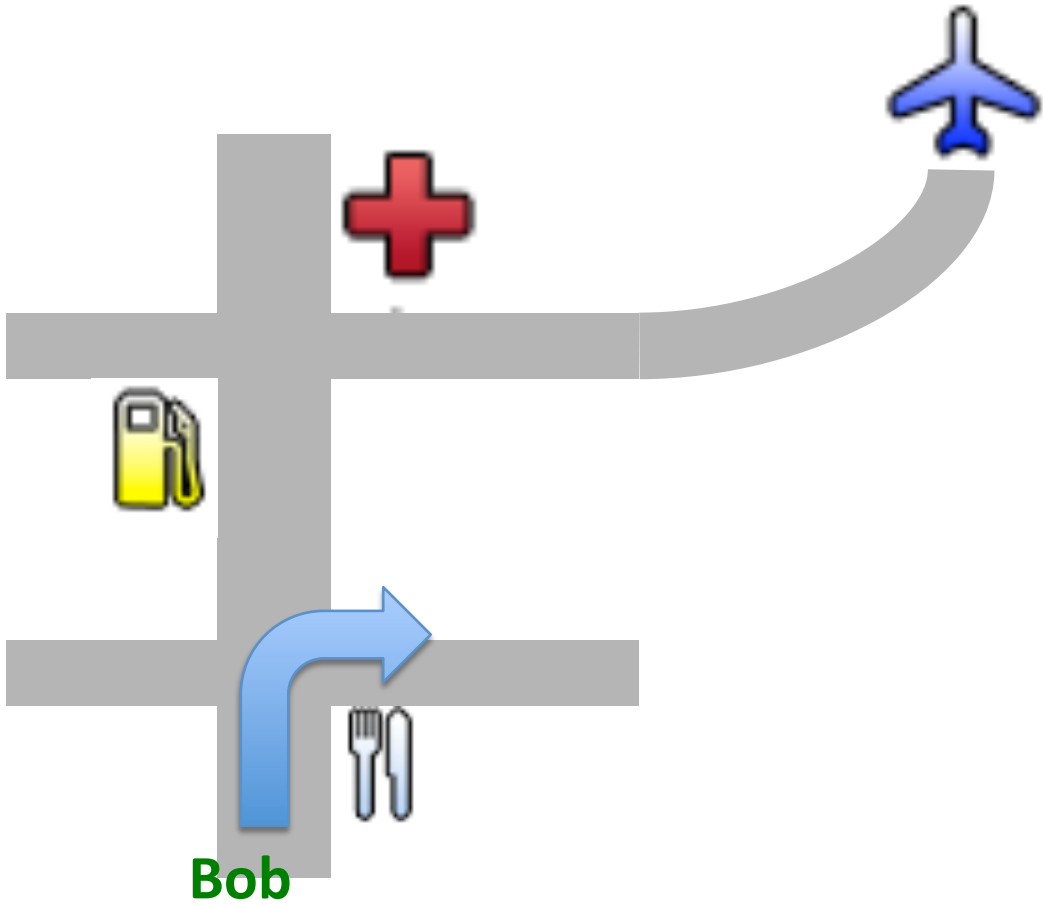

Navigation Example

Alice: 在醫院的地方右轉

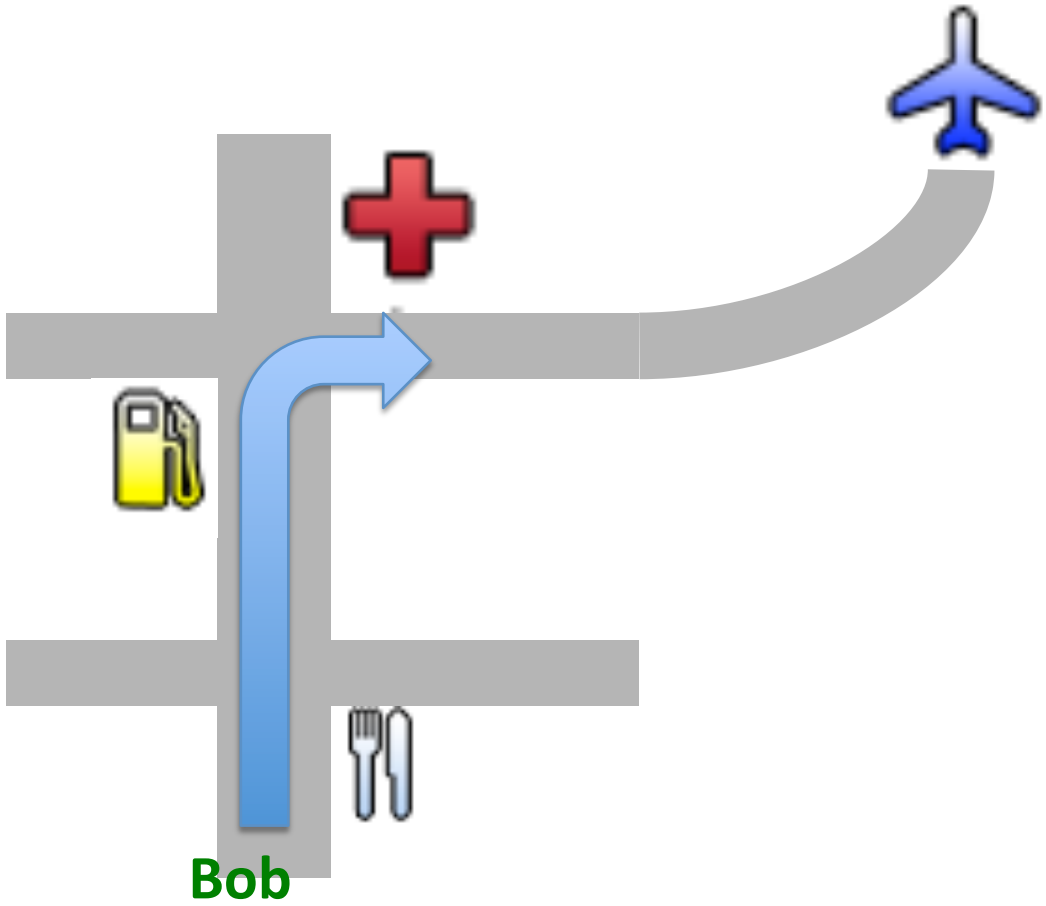

Navigation Example

Scenario 1

在餐廳的地方右轉

Scenario 2

在醫院的地方右轉

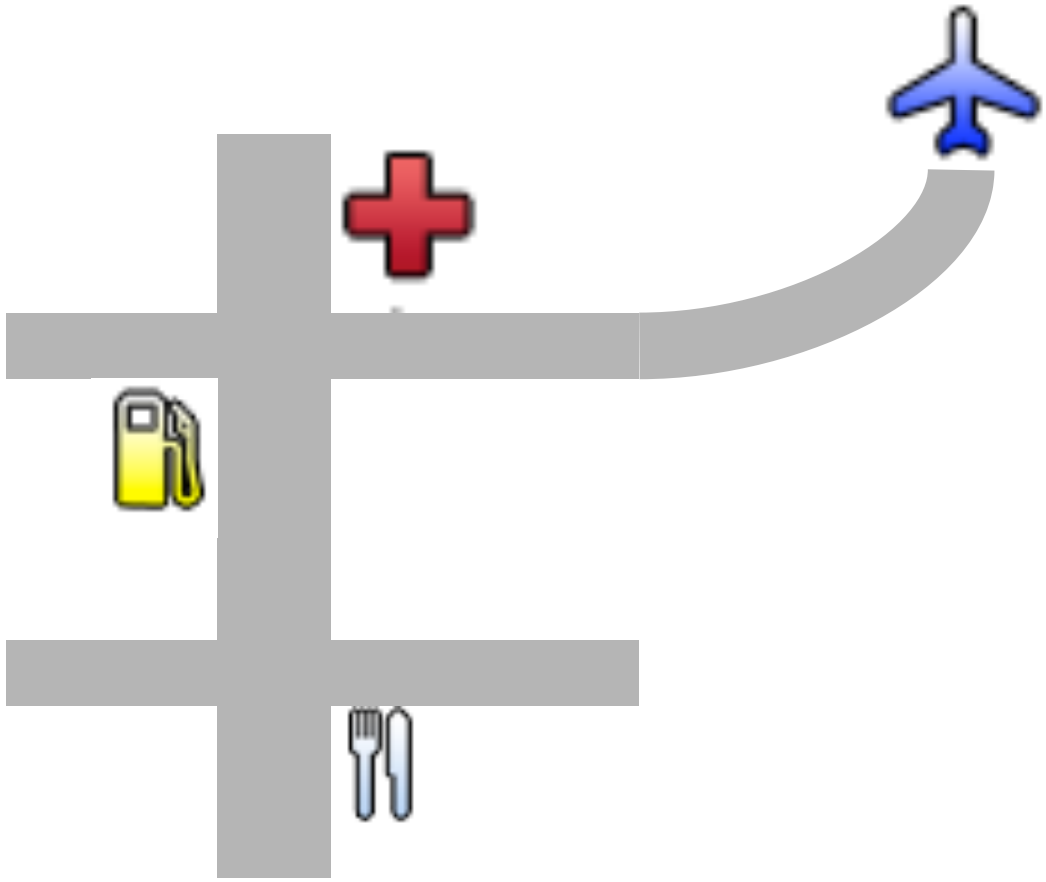

Navigation Example

Scenario 1

在　　的地方右轉

Scenario 2

在　　的地方右轉

Make a right turn

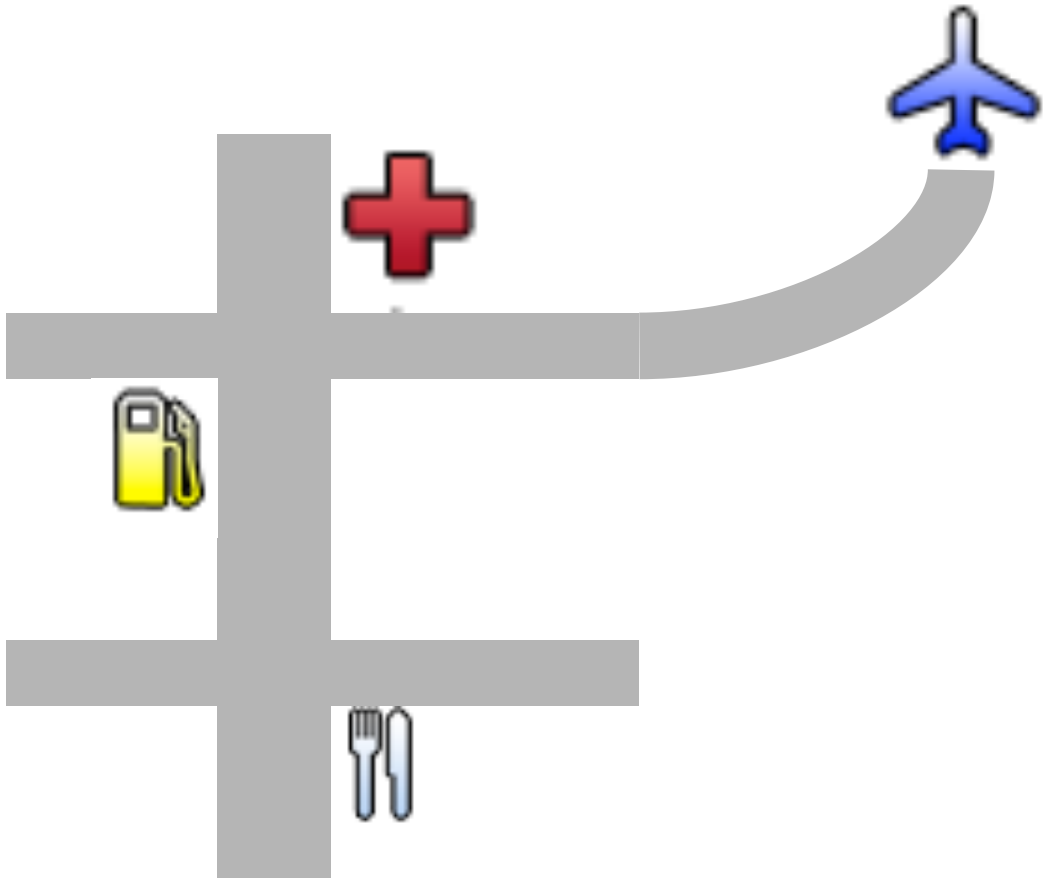

Navigation Example

Scenario 1

在餐廳的地方右轉

Scenario 2

在醫院的地方右轉

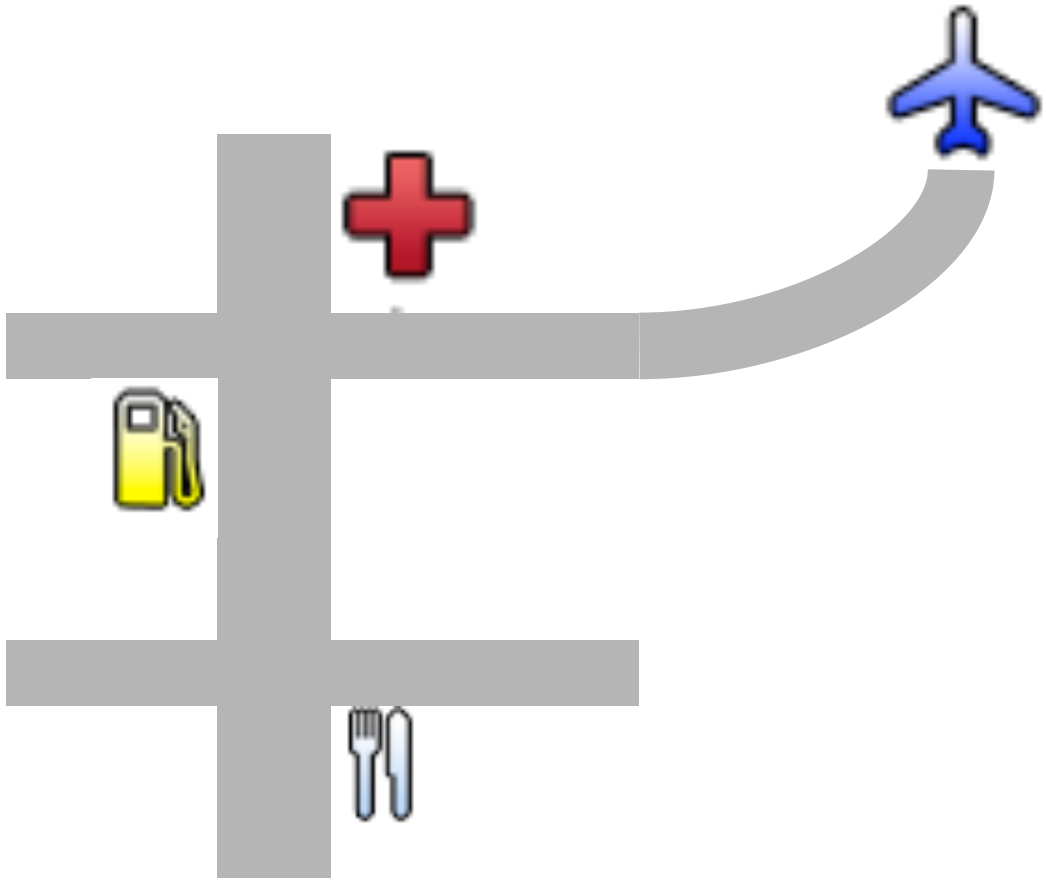

Navigation Example

Scenario 1

餐廳

Scenario 2

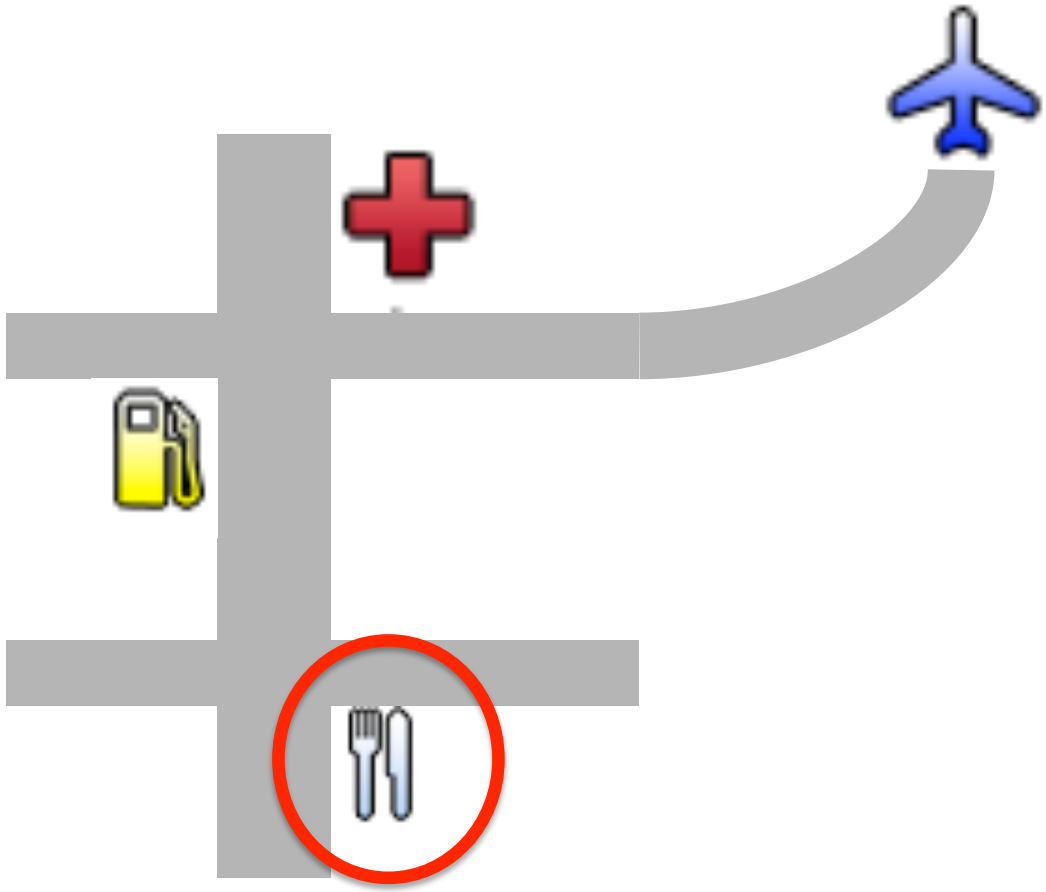

Navigation Example

Scenario 1

在餐廳的地方右轉

Scenario 2

在醫院的地方右轉

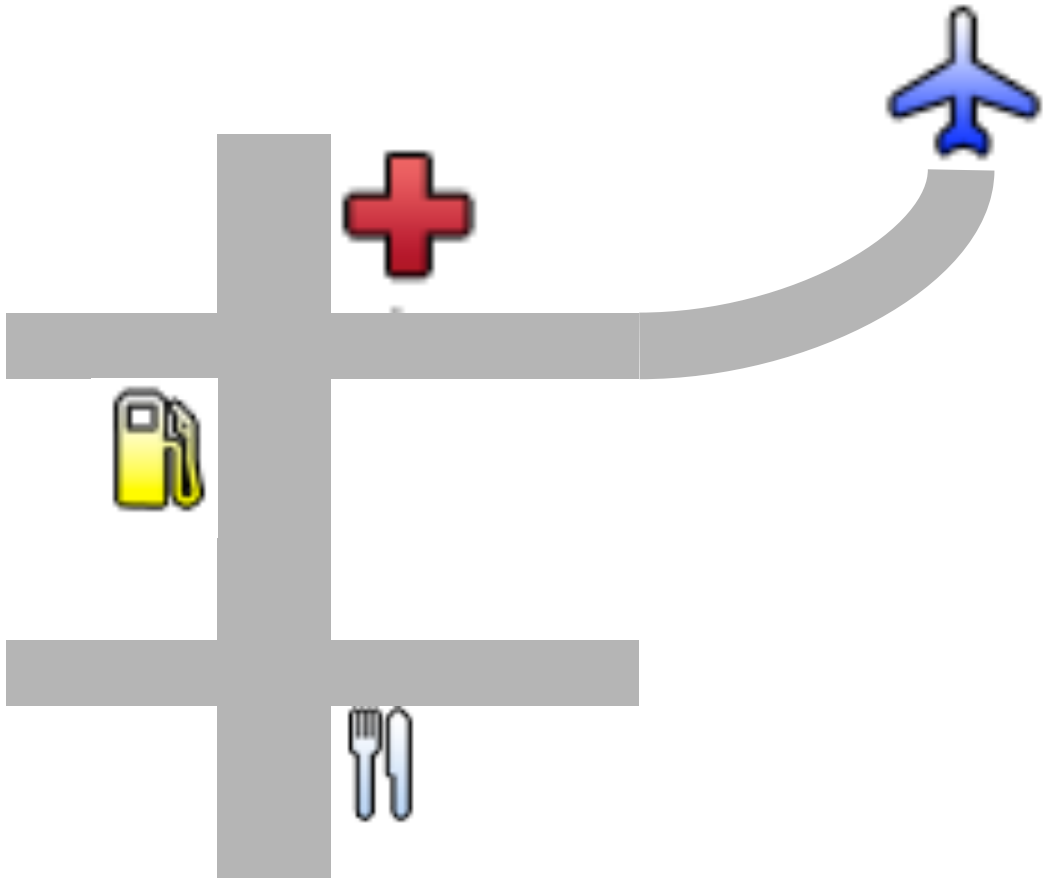

Navigation Example

Scenario 1

Scenario 2

醫院

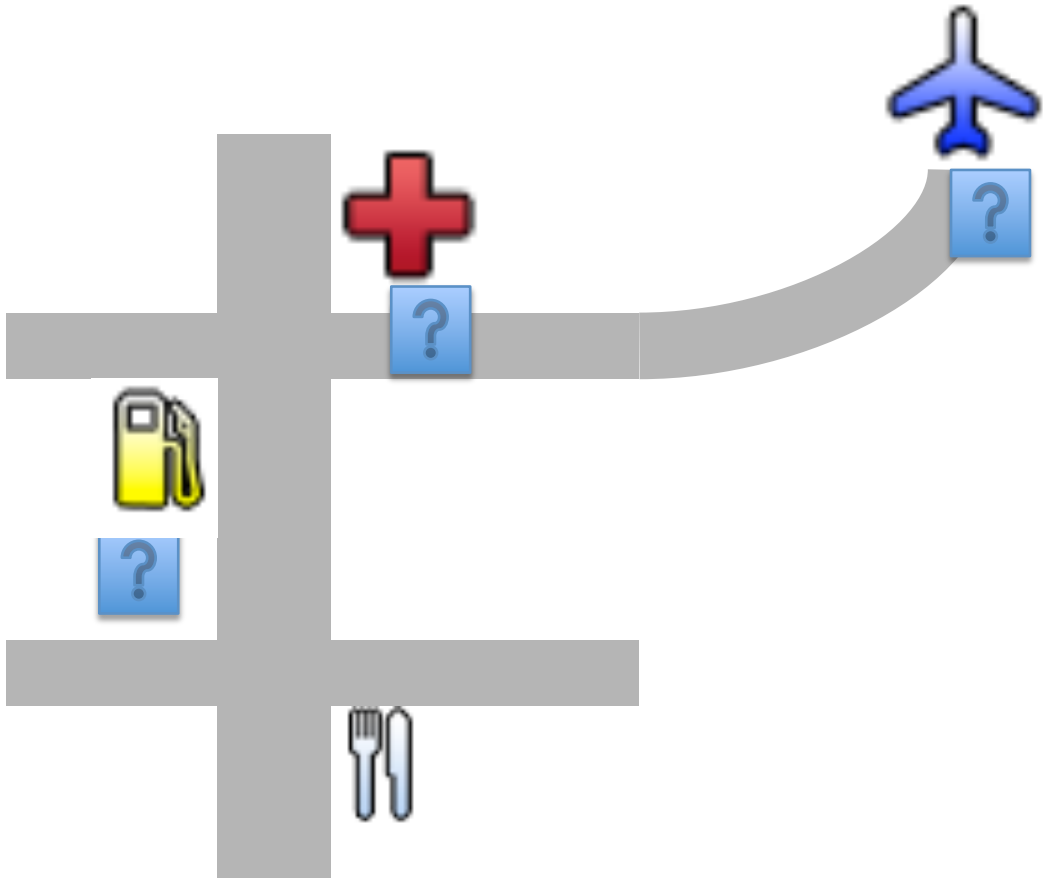

Overview

- **Background and related works**
- Completed work: Sportscasting
  - Tactical generation
  - Strategic generation
  - Human evaluation
- Proposed work: Navigation instructions
- Conclusions

Semantic Parser Learners

- Learn a function from NL to MR

NL: “Purple3 passes the ball to Purple5”

Semantic Parsing
(NL → MR)

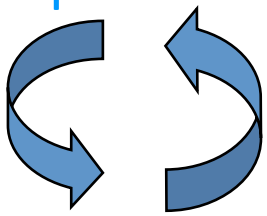

Tactical Generation
(MR → NL)

MR: Pass (Purple3, Purple5)

- We experiment with two semantic parser learners
  - –WASP [Wong & Mooney, 2006; 2007]
  - –KRISP [Kate & Mooney, 2006]

WASP: **W**ord **A**lignment-based **S**emantic **P**arsing

- Uses statistical machine translation techniques
  - Synchronous context-free grammars (SCFG) [Wu, 1997; Melamed, 2004; Chiang, 2005]
  - Word alignments [Brown et al., 1993; Och & Ney, 2003]
- Capable of both semantic parsing and tactical generation

KRISP: **K**ernel-based **R**obust **I**nterpretation by **S**emantic **P**arsing

- Productions of MR language are treated like semantic concepts
- SVM classifier is trained for each production with string subsequence kernel
- These classifiers are used to compositionally build MRs of the sentences
- More resistant to noisy supervision but incapable of tactical generation

KRISPER: KRISP with **E**M-like **R**etraining

- Extension of KRISP that learns from ***ambiguous*** supervision [Kate & Mooney, 2007]
- Uses an iterative EM-like method to gradually converge on a correct meaning for each sentence.

Overview

- Background and related works
- **Completed work: Sportscasting**
  - Tactical generation
  - Strategic generation
  - Human evaluation
- Proposed work: Navigation instructions
- Conclusions

Tractable Challenge Problem: Learning to Be a Sportscaster

- **Goal**: Learn from realistic data of natural language used in a representative context while avoiding difficult issues in computer perception (i.e. speech and vision).
- **Solution**: Learn from textually annotated traces of activity in a simulated environment.
- **Example**: Traces of games in the RoboCup simulator paired with textual sportscaster commentary.

RoboCup Simulation League

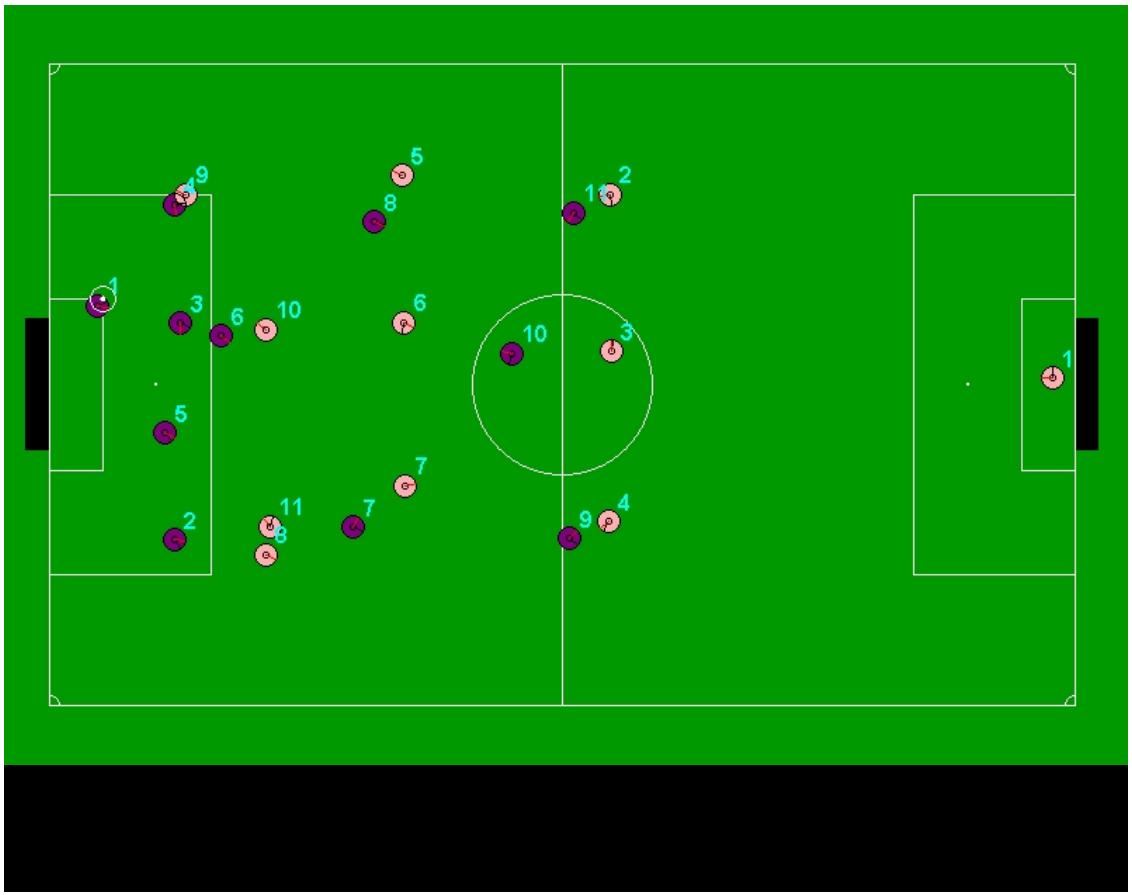

RoboCup Simulation League

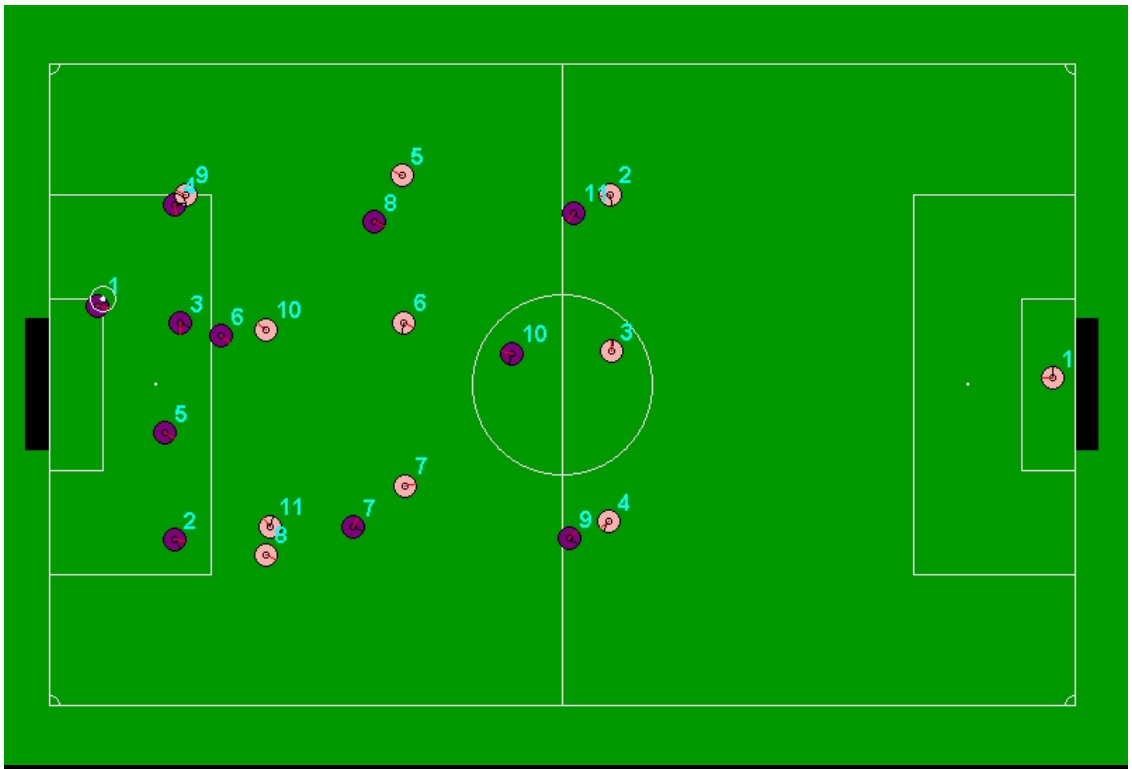

Learning to Sportscast

- Learn to sportscast by observing sample human sportscasts
- Build a function that maps between natural language (NL) and meaning representation (MR)
  - NL: Textual commentaries about the game
  - MR: Predicate logic formulas that represent events in the game

Robocup Sportscaster Trace

Natural Language Commentary	Meaning Representation
Purple goalie turns the ball over to Pink8	badPass (Purple1, Pink8) turnover (Purple1, Pink8)
Purple team is very sloppy today Pink8 passes the ball to Pink11	kick (Pink8) pass (Pink8, Pink11) kick (Pink11)
Pink11 looks around for a teammate	kick (Pink11) ballstopped
Pink11 makes a long pass to Pink8	kick (Pink11) pass (Pink11, Pink8)
Pink8 passes back to Pink11	kick (Pink8) pass (Pink8, Pink11)

Robocup Sportscaster Trace

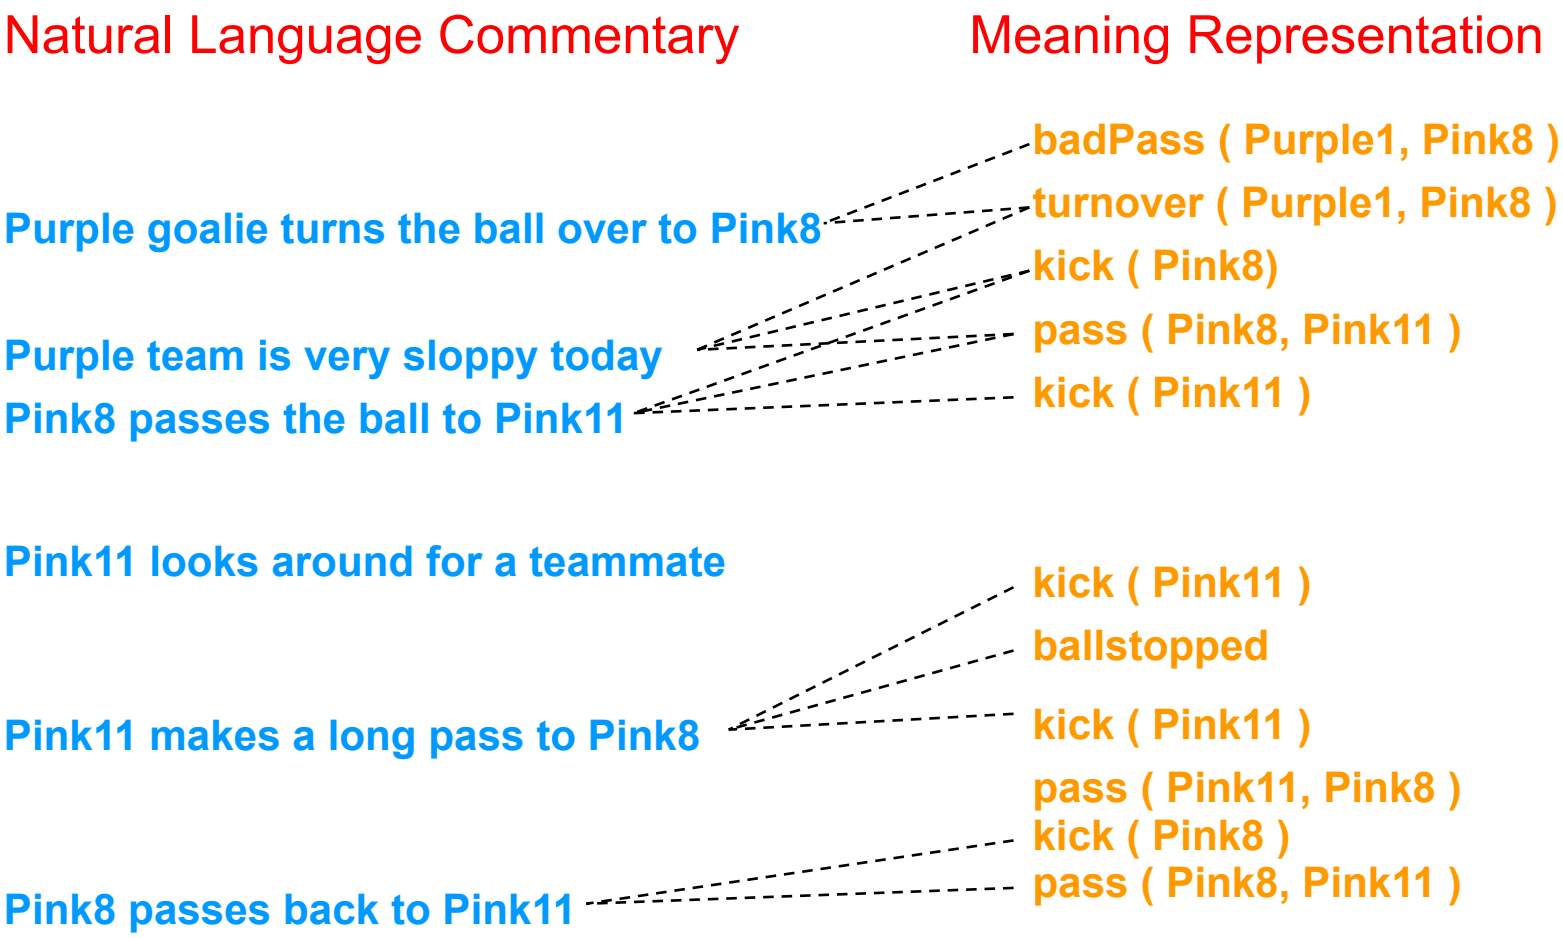

Robocup Sportscaster Trace

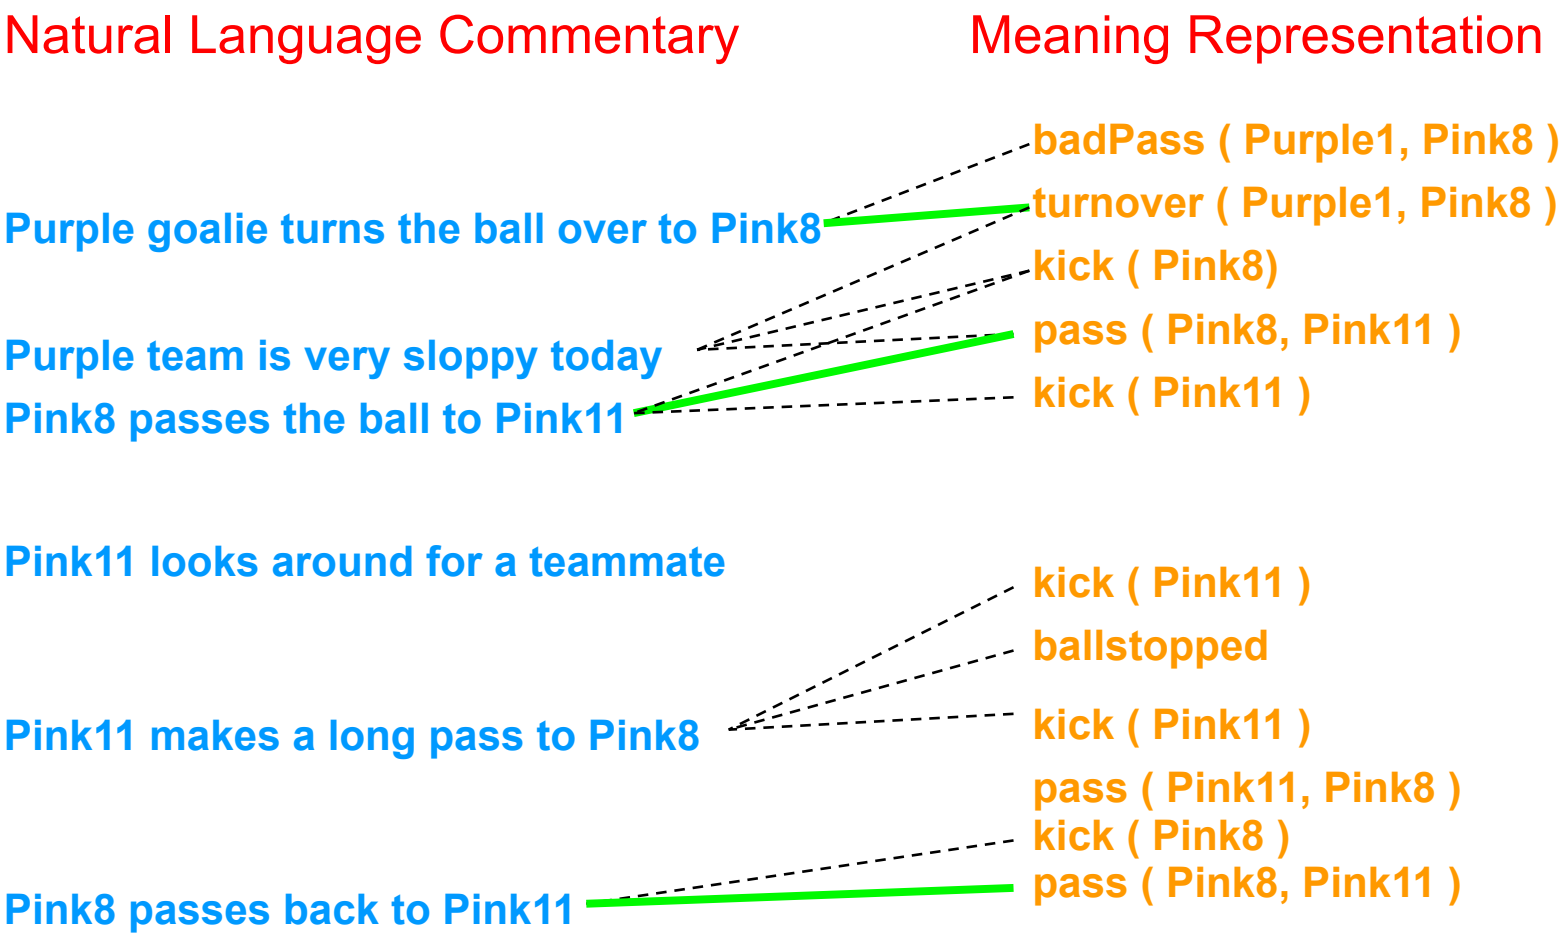

Robocup Sportscaster Trace

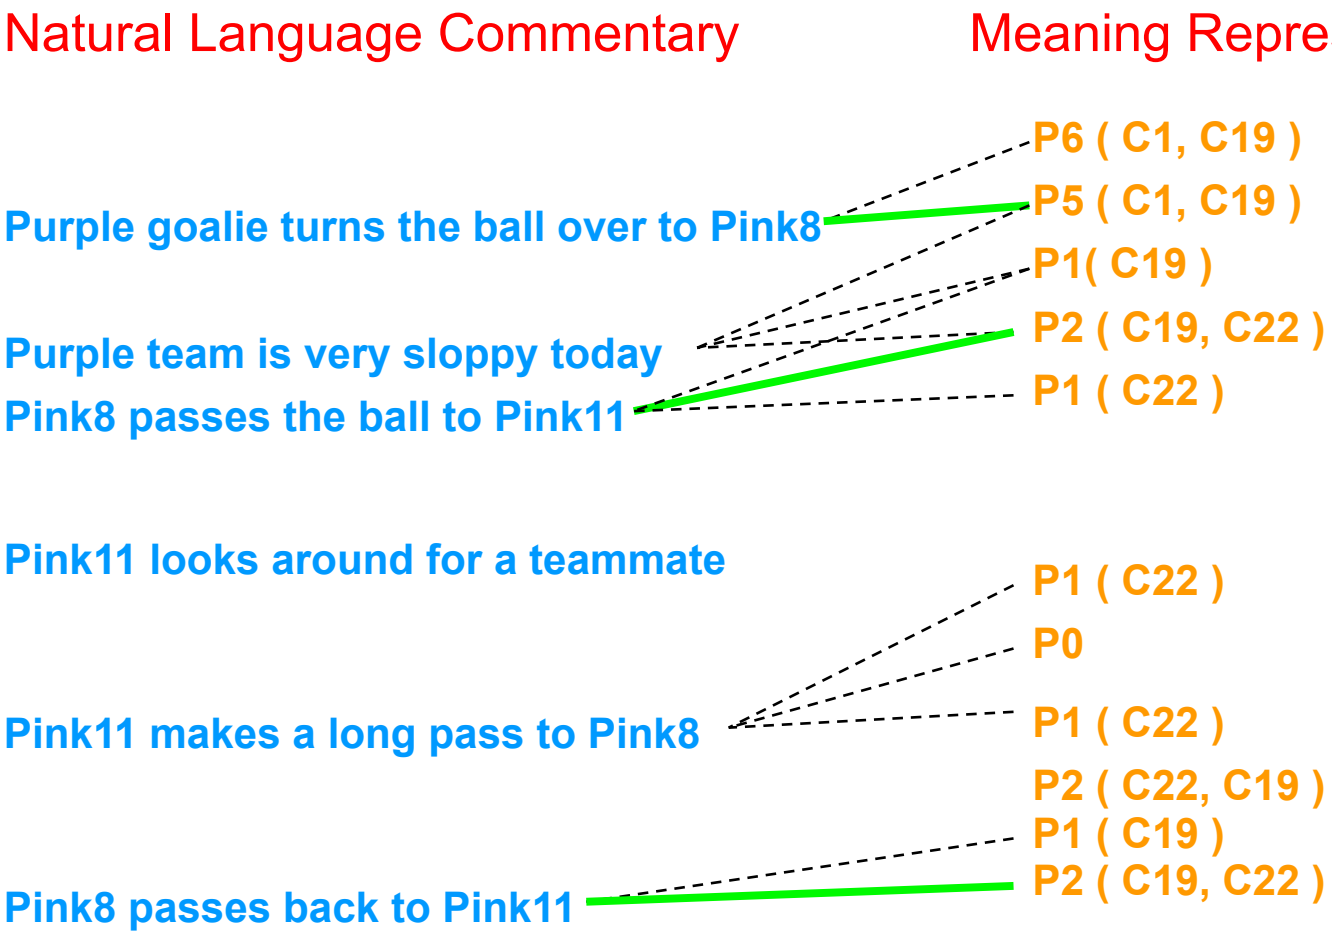

Robocup Data

- Collected human textual commentary for the 4 Robocup championship games from 2001-2004.
  - Avg # events/game = 2,613
  - Avg # English sentences/game = 509
  - Avg # Korean sentences/game = 499
- Each sentence matched to all events within previous 5 seconds.
  - Avg # MRs/sentence = 2.5 (min 1, max 12)
- Manually annotated with correct matchings of sentences to MRs (for evaluation purposes only).

Overview

- Background and related works
- Completed work: Sportscasting
  - **Tactical generation**
  - Strategic generation
  - Human evaluation
- Proposed work: Navigation instructions
- Conclusions

Tactical Generation

- Learn how to generate NL from MR
- Example:

  Pass(Pink2, Pink3) → "Pink2 kicks the ball to Pink3"

- Two steps
  1. Disambiguate the training data
  2. Learn a language generator

WASPER

- WASP with EM-like retraining to handle ambiguous training data.
- Same augmentation as added to KRISP to create KRISPER.

KRISPER-WASP

- First train KRISPER to disambiguate the data
- Then train WASP on the resulting unambiguously supervised data.

WASPER-GEN

- Determines the best matching based on **generation** (MR→NL).
- Score each potential NL/MR pair by using the currently trained WASP^{-1} generator.
- Compute NIST MT score [NIST report, 2002] between the generated sentence and the potential matching sentence.

BLEU/NIST scores

Target: Purple2 quickly passes to Purple3

Candidate: Purple2 passes to Purple3

1-grams: Purple2, passes, to, Purple3

2-grams: Purple2 passes, passes to, to Purple3

3-grams: Purple2 passes to, passes to Purple3

4-gram: Purple2 passes to Purple3

BLEU/NIST scores

Target: Purple2 quickly passes to Purple3

Candidate: Purple2 passes to Purple3

4/4 1-grams: Purple2, passes, to, Purple3

2-grams: Purple2 passes, passes to, to Purple3

3-grams: Purple2 passes to, passes to Purple3

4-gram: Purple2 passes to Purple3

BLEU/NIST scores

Target: Purple2 quickly passes to Purple3

Candidate: Purple2 passes to Purple3

4/4 1-grams: Purple2, passes, to, Purple3

2/3 2-grams: Purple2 passes, passes to, to Purple3

3-grams: Purple2 passes to, passes to Purple3

4-gram: Purple2 passes to Purple3

BLEU/NIST scores

Target: Purple2 quickly passes to Purple3

Candidate: Purple2 passes to Purple3

4/4 1-grams: Purple2, passes, to, Purple3

2/3 2-grams: Purple2 passes, passes to, to Purple3

1/2 3-grams: Purple2 passes to, passes to Purple3

4-gram: Purple2 passes to Purple3

BLEU/NIST scores

Target: Purple2 quickly passes to Purple3

Candidate: Purple2 passes to Purple3

4/4 1-grams: Purple2, passes, to, Purple3

2/3 2-grams: Purple2 passes, passes to, to Purple3

1/2 3-grams: Purple2 passes to, passes to Purple3

0/1 4-gram: Purple2 passes to Purple3

BLEU/NIST scores

Target: Purple2 quickly passes to Purple3

Candidate: Purple2 passes to Purple3

4/4 1-grams: Purple2, passes, to, Purple3

2/3 2-grams: Purple2 passes, passes to, to Purple3

1/2 3-grams: Purple2 passes to, passes to Purple3

0/1 4-gram: Purple2 passes to Purple3

BLEU: $\sqrt[4]{\frac{4}{4} \times \frac{2}{3} \times \frac{1}{2} \times \frac{0}{1}} = 0$

BLEU/NIST scores

Target: Purple2 quickly passes to Purple3

Candidate: Purple2 passes to Purple3

4/4 1-grams: Purple2, passes, to, Purple3

2/3 2-grams: Purple2 passes, passes to, to Purple3

1/2 3-grams: Purple2 passes to, passes to Purple3

0/1 4-gram: Purple2 passes to Purple3

BLEU: $\sqrt[4]{\frac{4}{4} \times \frac{2}{3} \times \frac{1}{2} \times \frac{0}{1}} = 0$ NIST: $\frac{4}{4} + \frac{2}{3} + \frac{1}{2} + \frac{0}{1} = 2.167$

System Overview

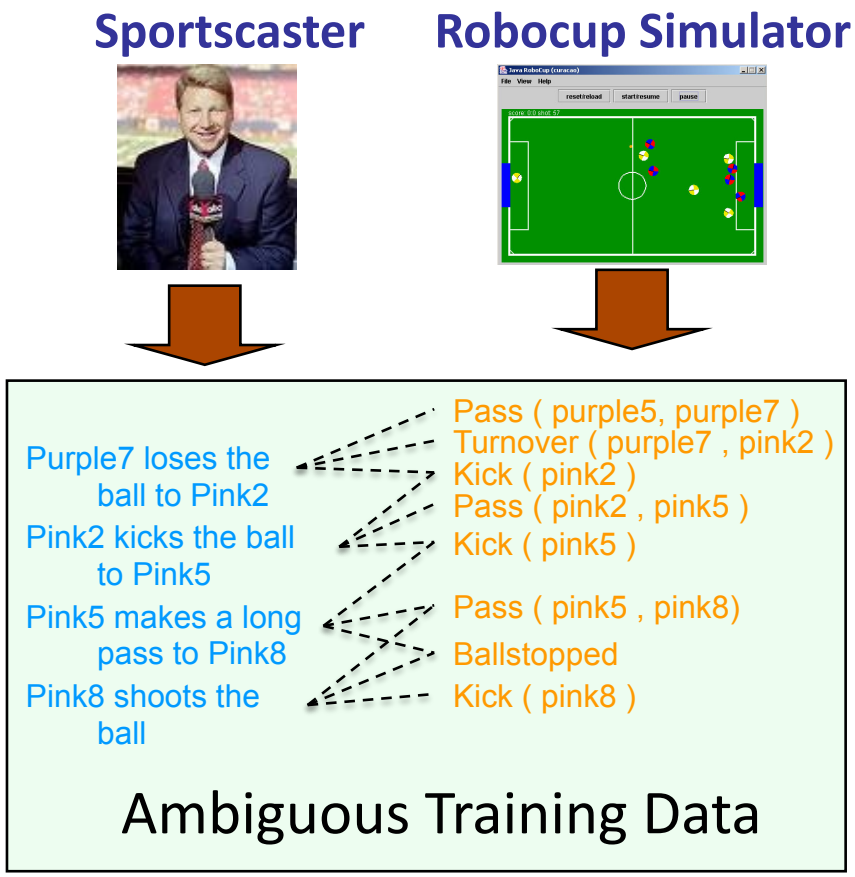

System Overview

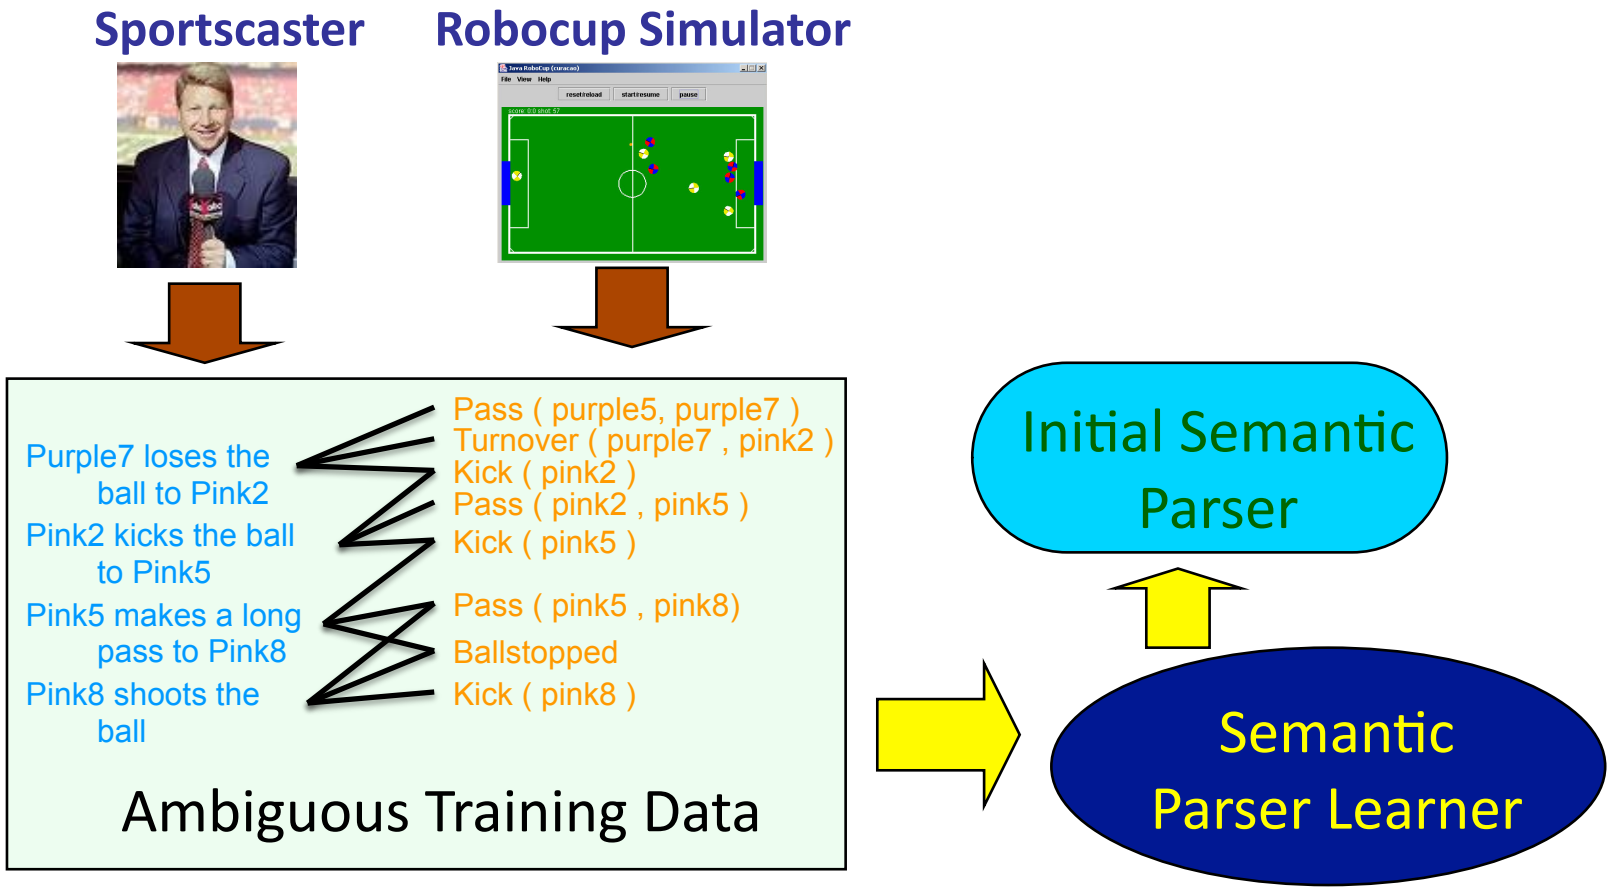

System Overview

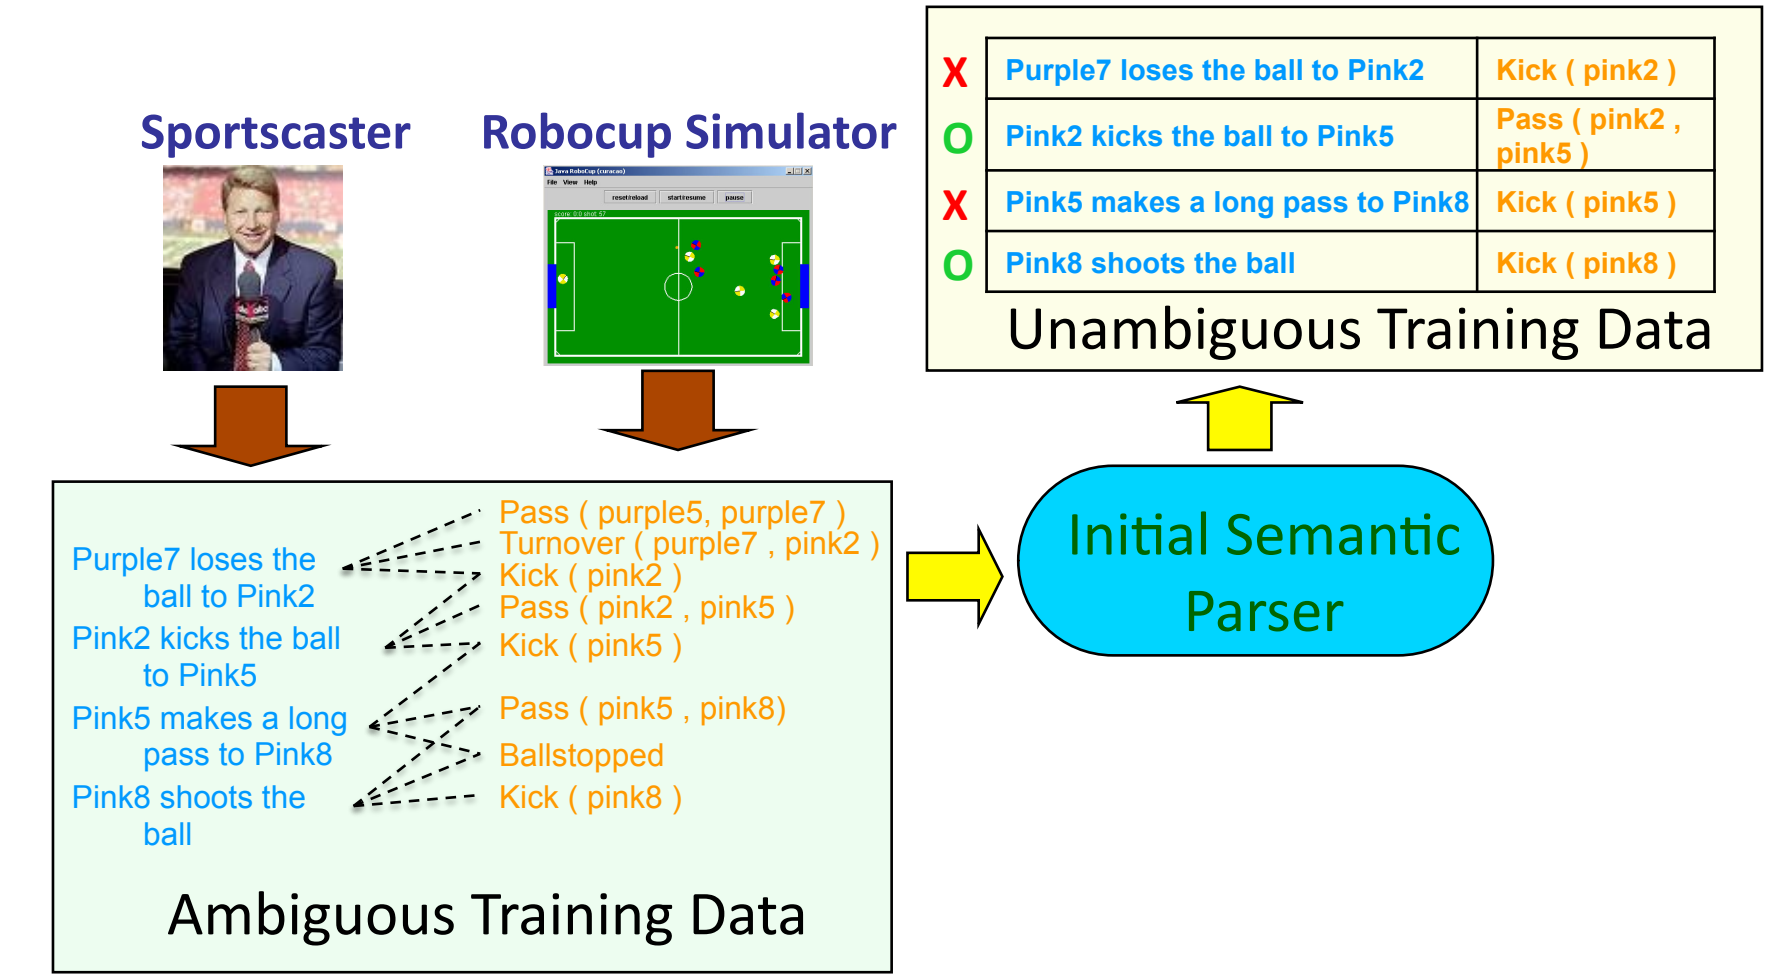

System Overview

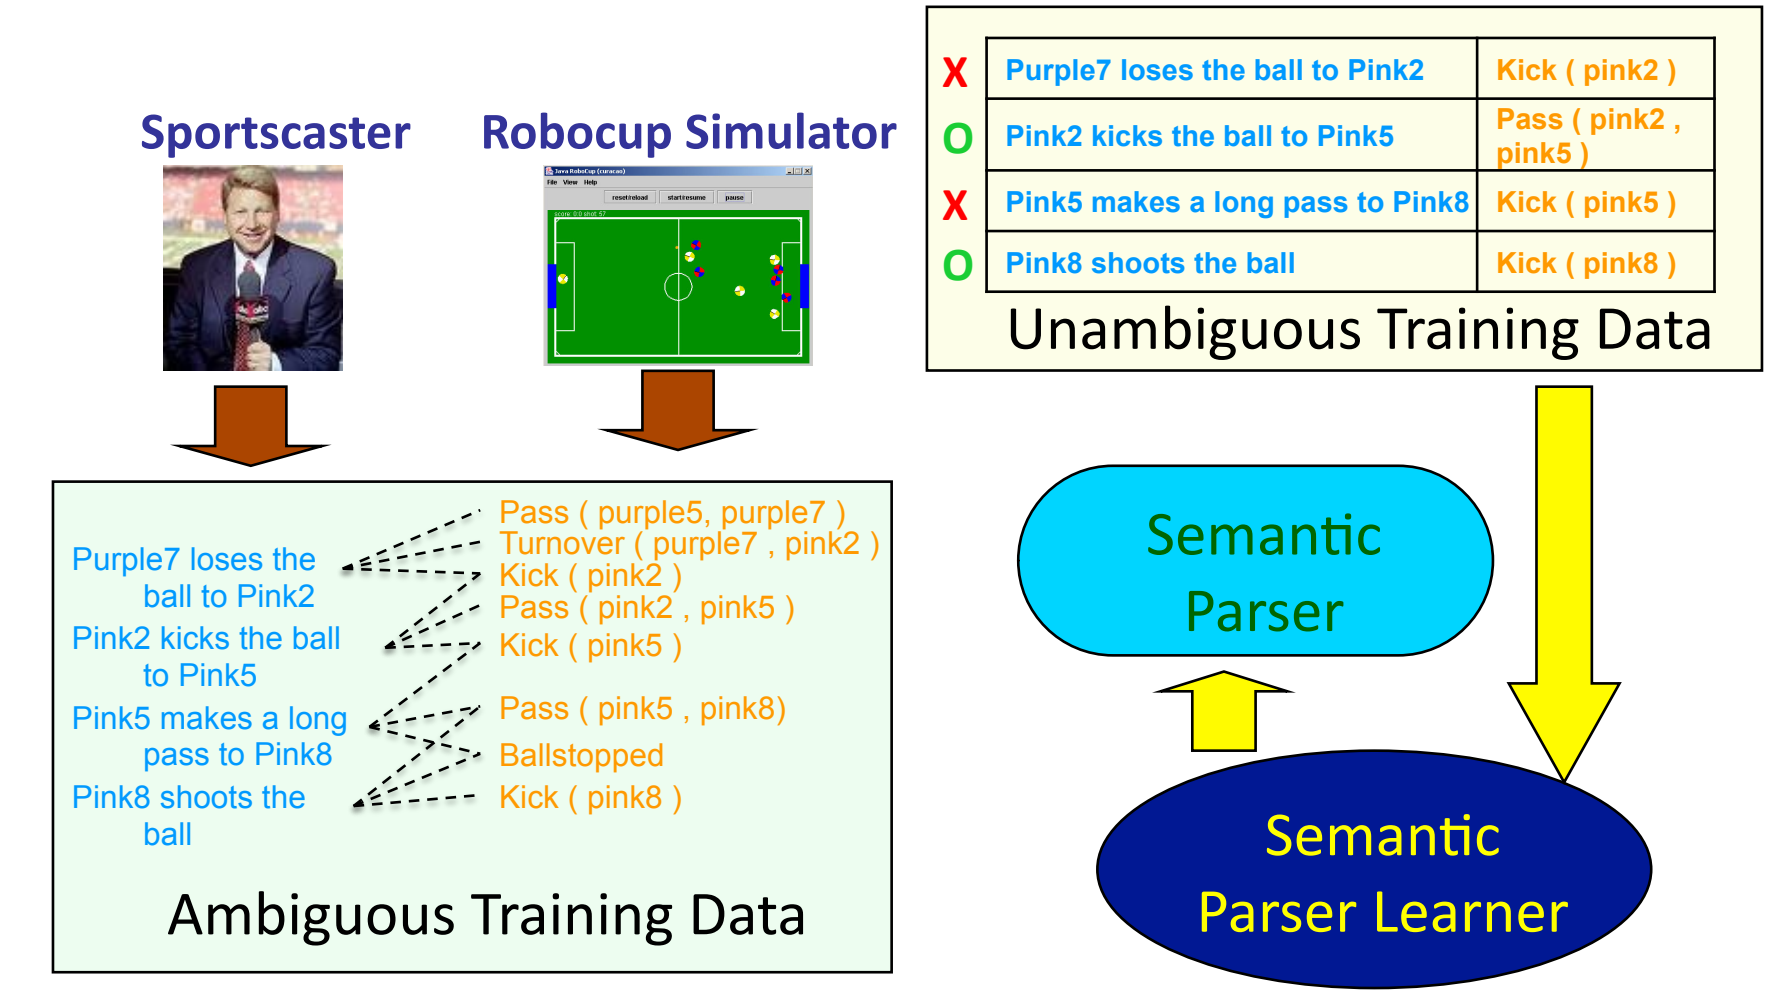

System Overview

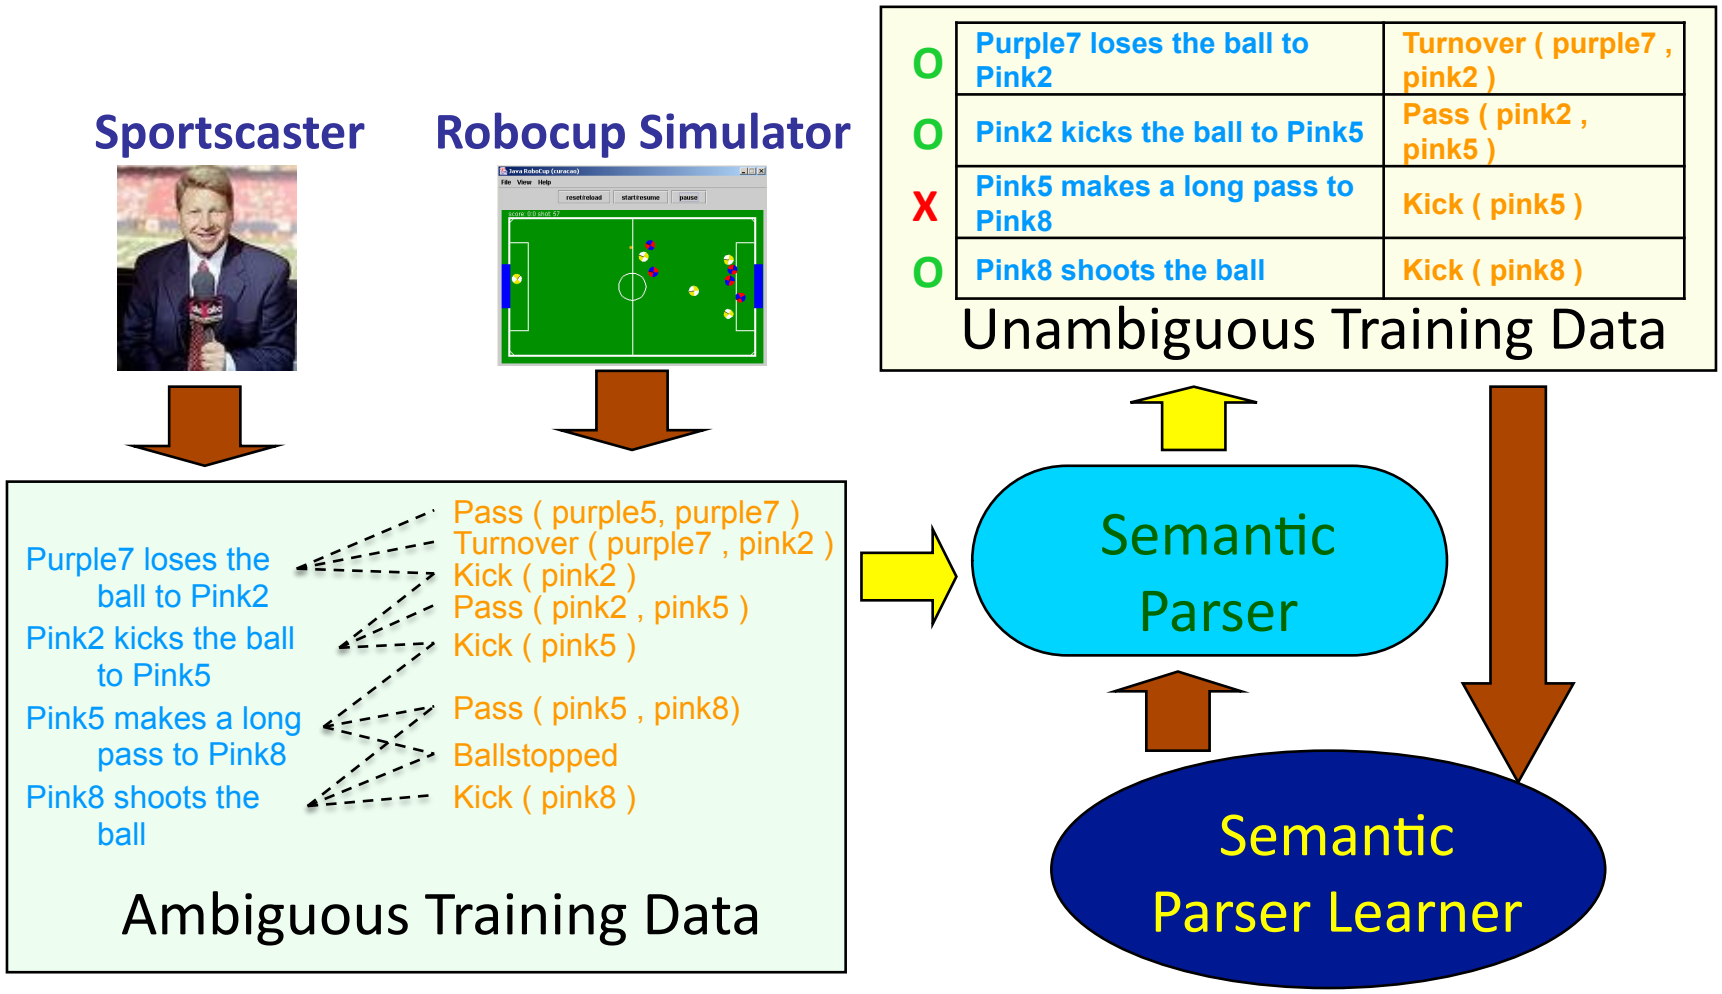

System Overview

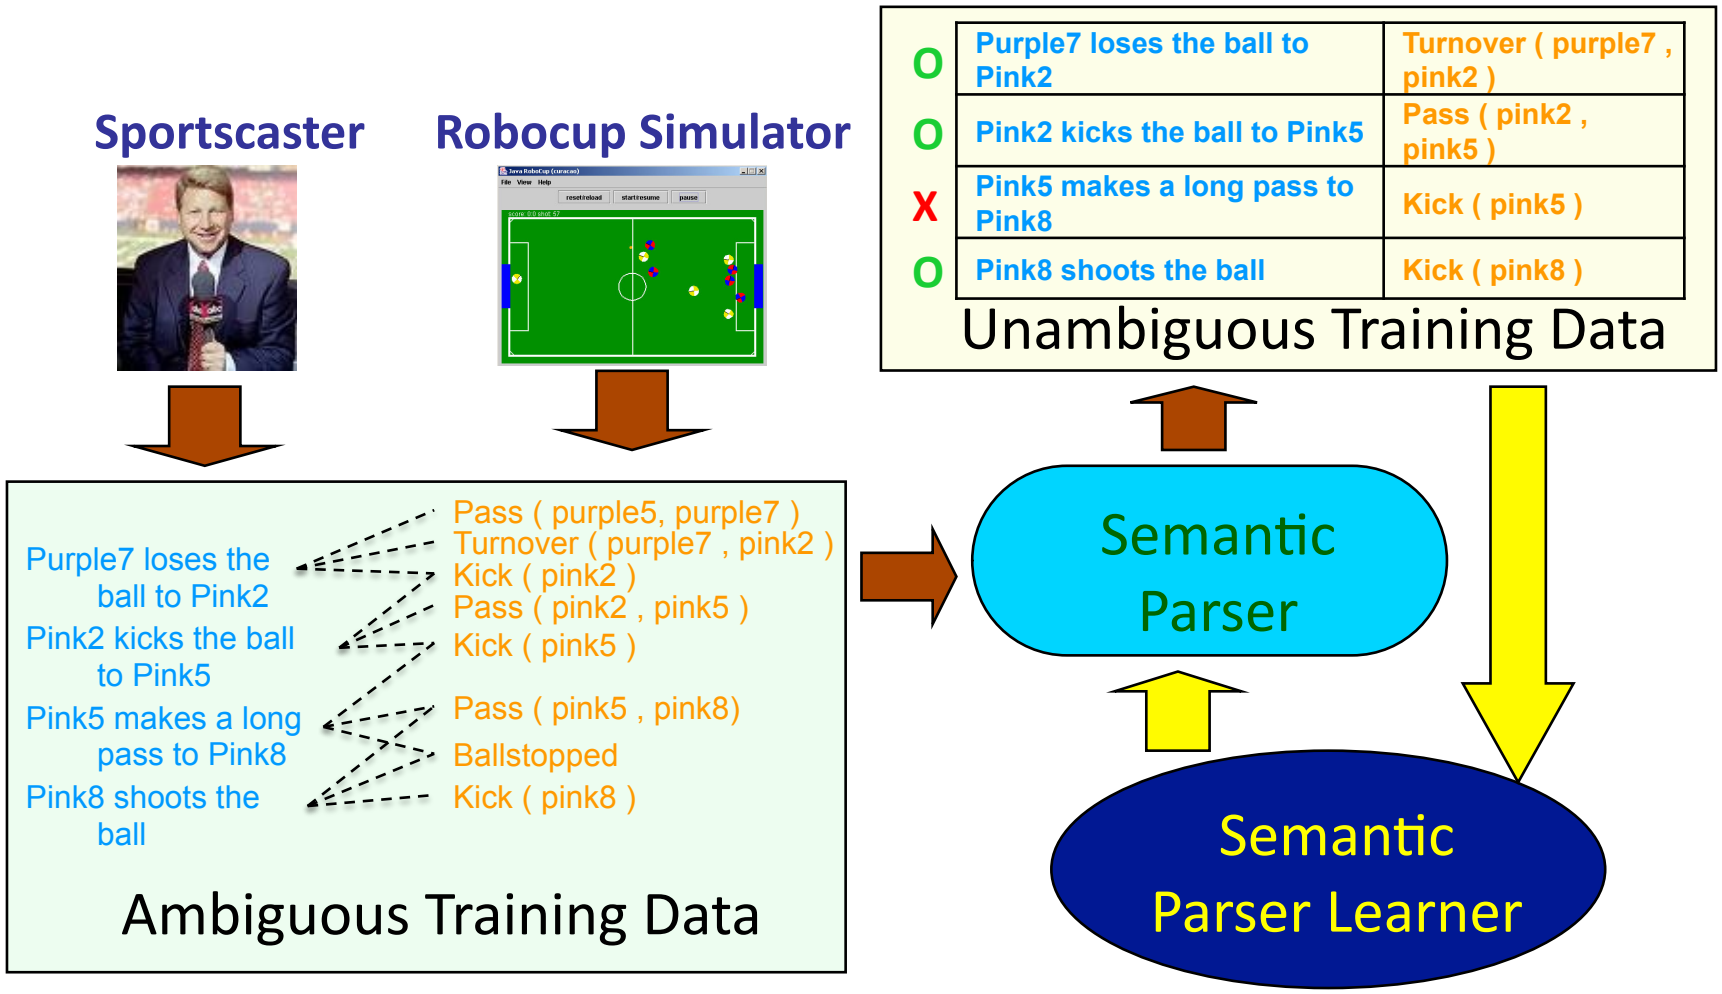

System Overview

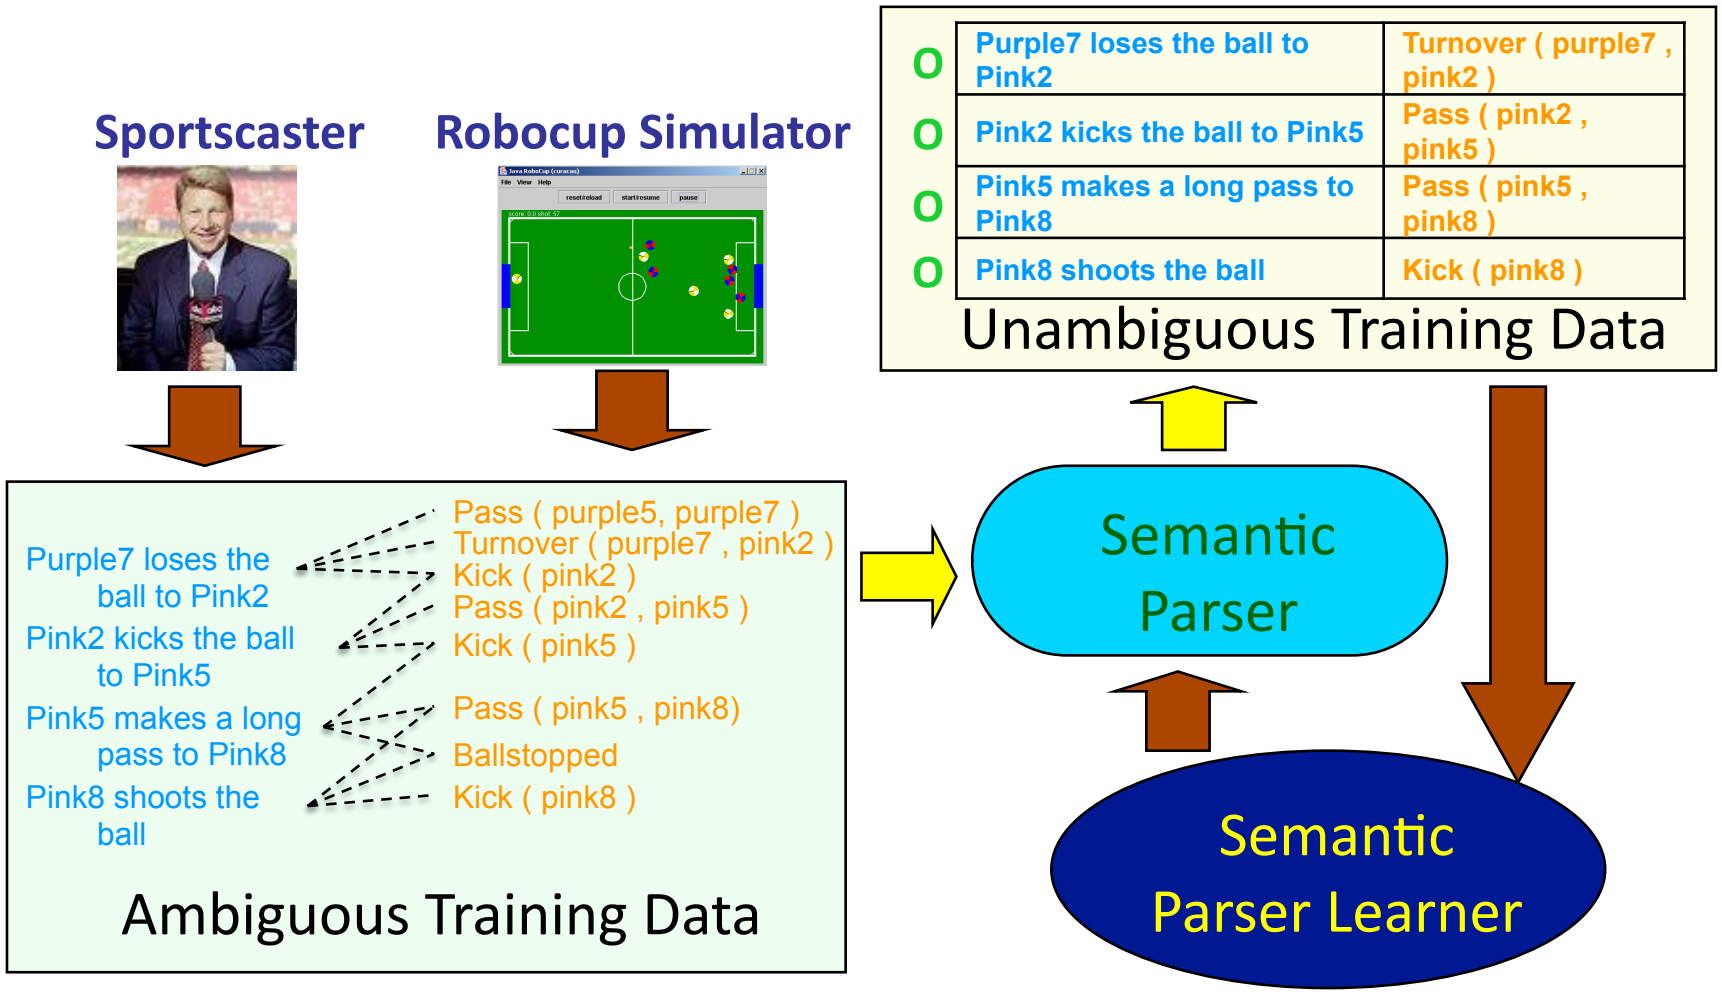

KRISPER and WASPER

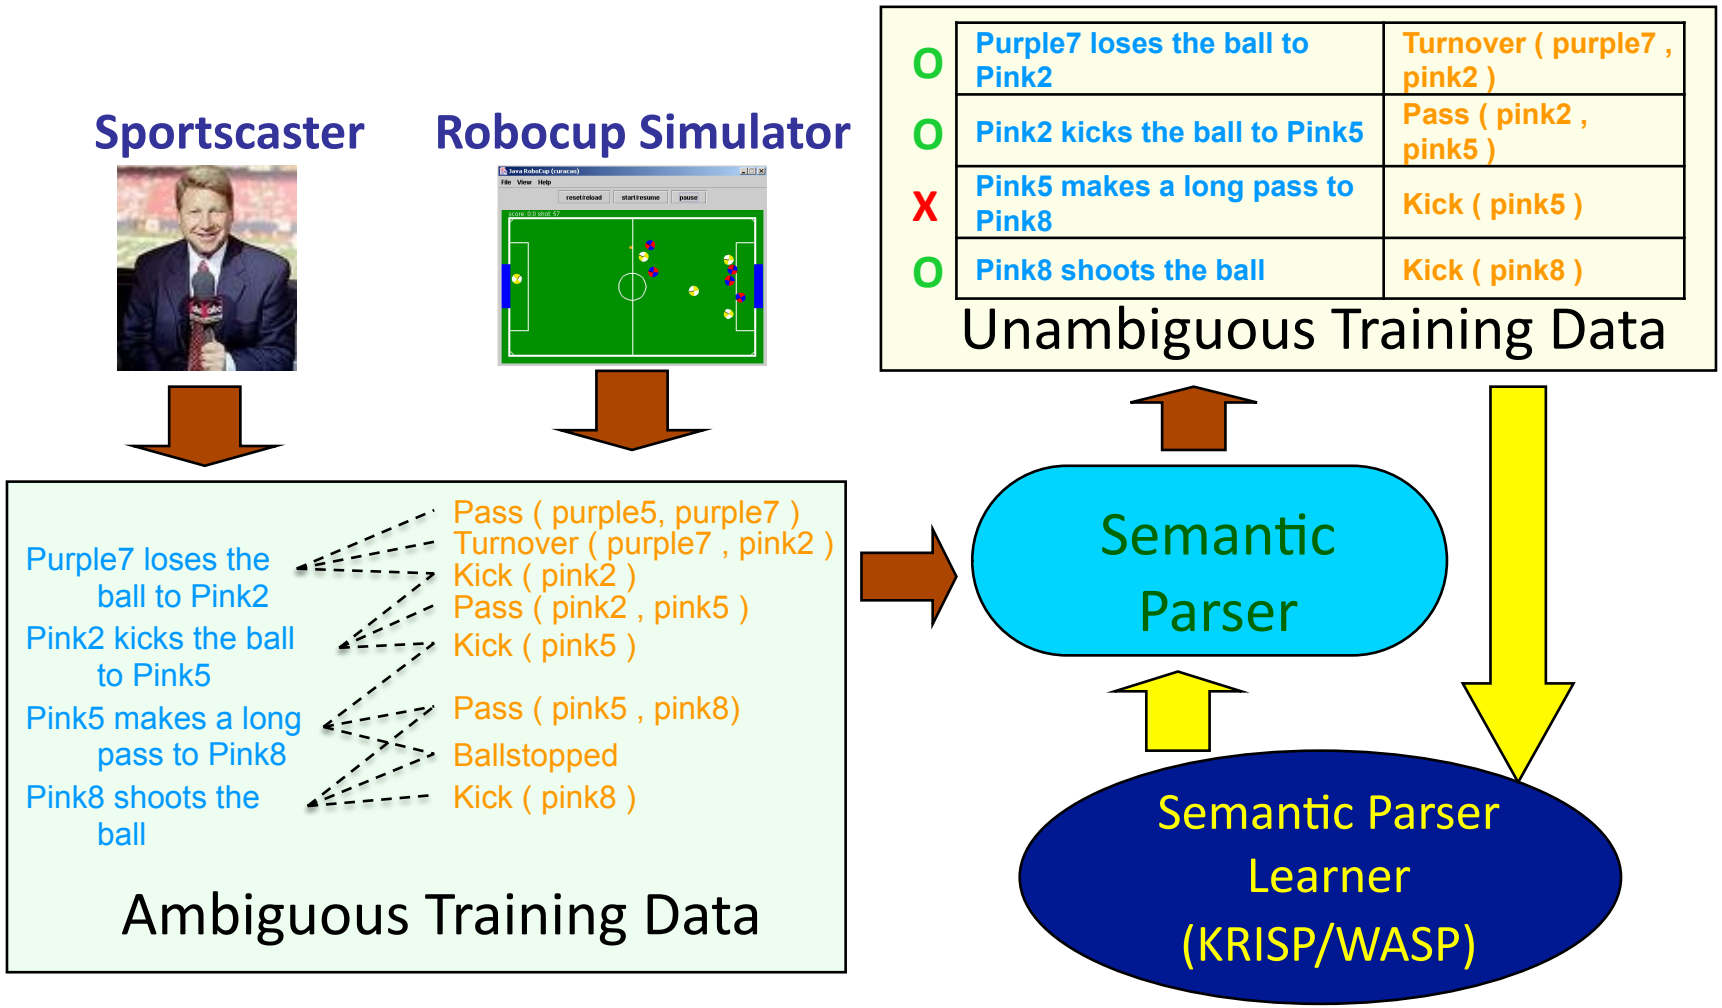

WASPER-GEN

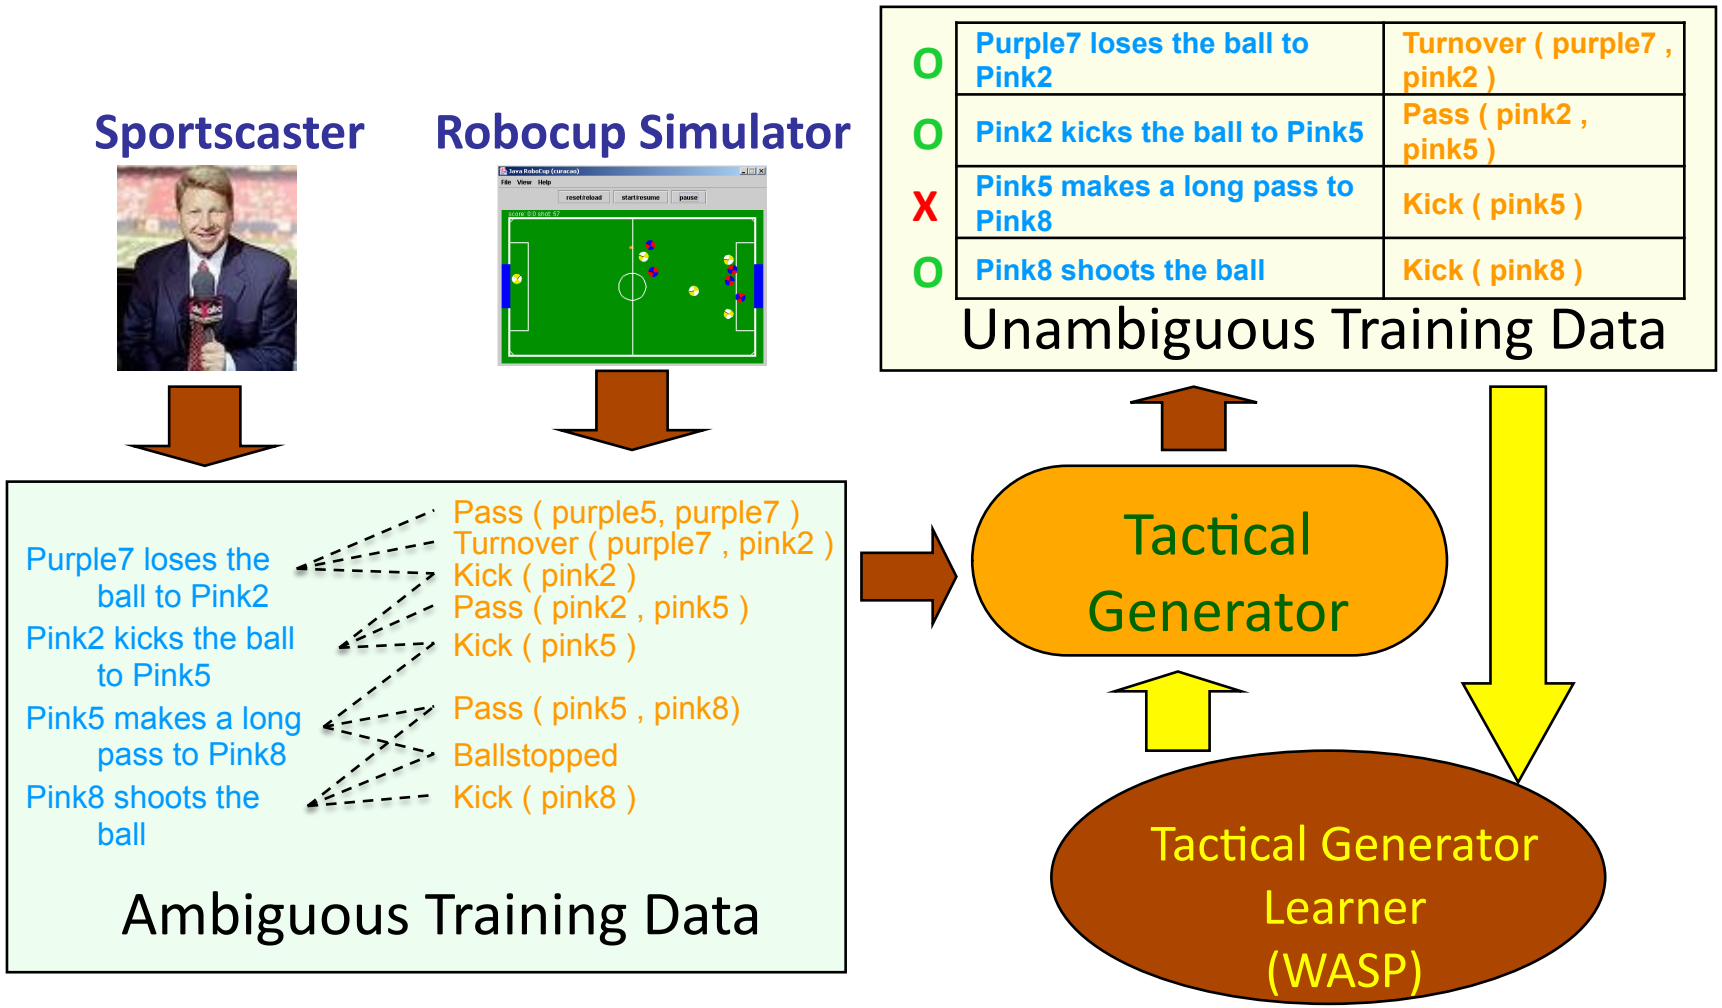
Sportscaster
Robocup Simulator
O
Purple7 loses the ball to Pink2
Turnover (purple7 , pink2)
O
Pink2 kicks the ball to Pink5
Pass (pink2 , pink5)
X
Pink5 makes a long pass to Pink8
Kick (pink5)
O
Pink8 shoots the ball
Kick (pink8)
Unambiguous Training Data
Purple7 loses the ball to Pink2
Pink2 kicks the ball to Pink5
Pink5 makes a long pass to Pink8
Pink8 shoots the ball
Pass (purple5, purple7)
Turnover (purple7 , pink2)
Kick (pink2)
Pass (pink2 , pink5)
Kick (pink5)
Pass (pink5 , pink8)
Ballstopped
Kick (pink8)
Ambiguous Training Data
Tactical Generator
Tactical Generator Learner (WASP)

Matching

- 4 Robocup championship games from 2001-2004.
  - Avg # events/game = 2,613
  - Avg # English sentences/game = 509
- Leave-one-game-out cross-validation
- Metric:
  - **Precision**: % of system's annotations that are correct
  - **Recall**: % of gold-standard annotations produced
  - **F-measure**: Harmonic mean of precision and recall

Matching Results

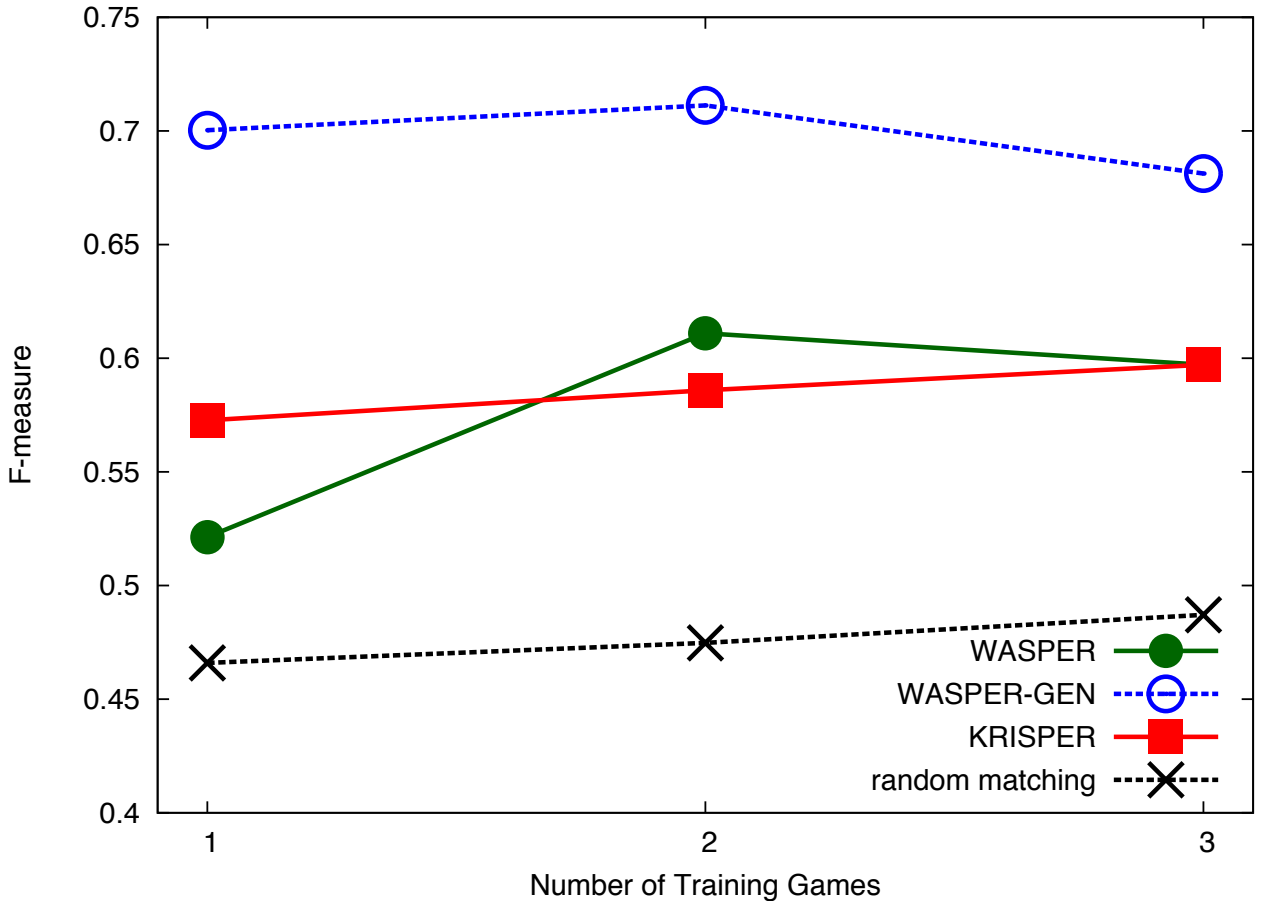

Tactical Generation

- Measure how accurately NL generator produces English sentences for chosen MRs in the test games.
- Use gold-standard matches to determine the correct sentence for each MR that has one.
- Leave-one-game-out cross-validation
- Metric:
  - **BLEU score:** [Papineni et al, 2002], N=4

Tactical Generation Results

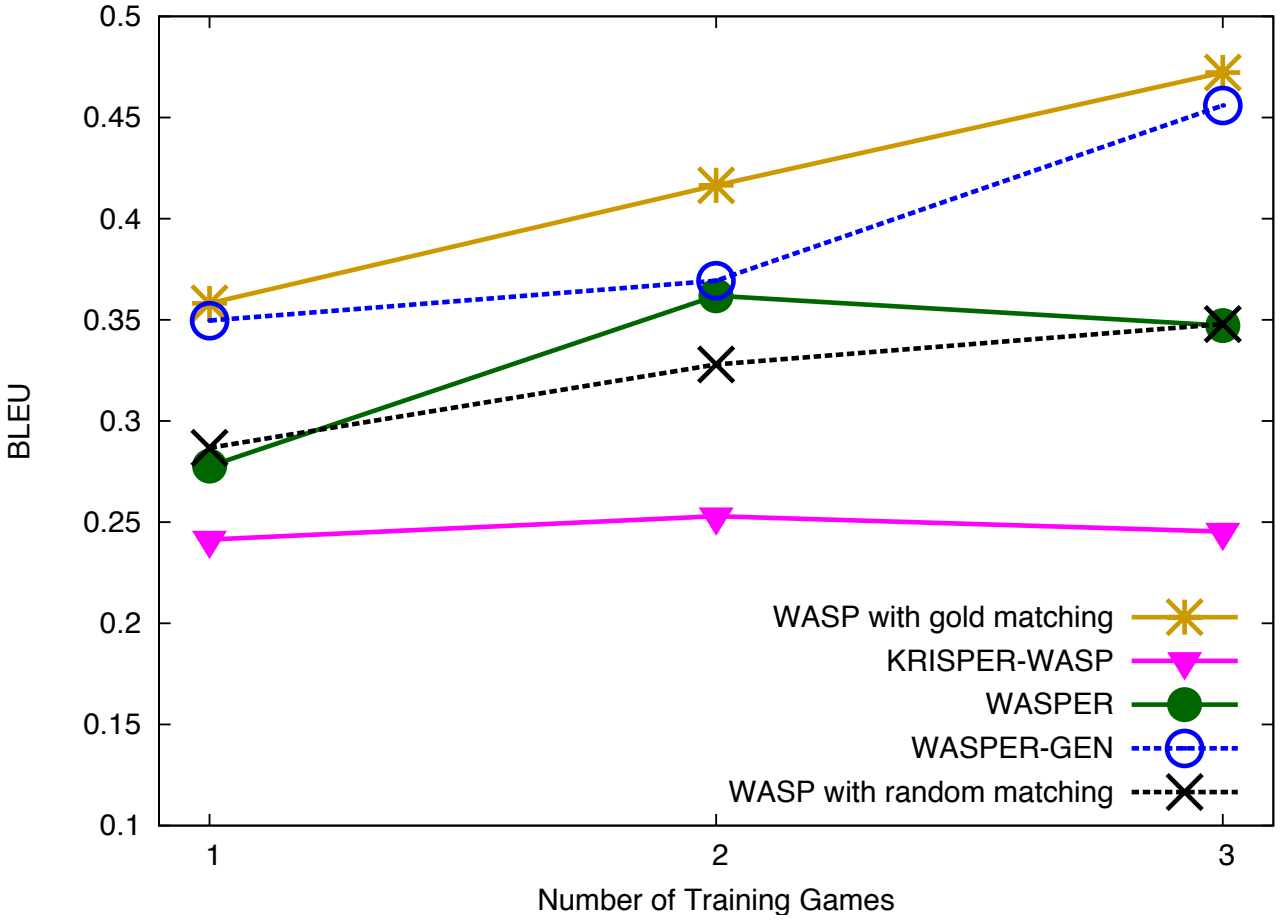

Overview

- Background and related works
- Completed work: Sportscasting
  - Tactical generation
  - **Strategic generation**
  - Human evaluation
- Proposed work: Navigation instructions
- Conclusions

Strategic Generation

- Generation requires not only knowing ***how*** to say something (tactical generation) but also ***what*** to say (strategic generation).
- For automated sportscasting, one must be able to effectively choose which events to describe.

Example of Strategic Generation

pass (purple7 , purple6)

ballstopped

kick (purple6)

pass (purple6 , purple2)

ballstopped

kick (purple2)

pass (purple2 , purple3)

kick (purple3)

badPass (purple3 , pink9)

turnover (purple3 , pink9)

Example of Strategic Generation

pass (purple7 , purple6)

ballstopped

kick (purple6)

pass (purple6 , purple2)

ballstopped

kick (purple2)

pass (purple2 , purple3)

kick (purple3)

badPass (purple3 , pink9)

turnover (purple3 , pink9)

Strategic Generation

- For each event type (e.g. pass, kick) estimate the probability that it is described by the sportscaster.
- Requires correct NL/MR matching
  - Use estimated matching from tactical generation
  - Iterative Generation Strategy Learning

Iterative Generation Strategy Learning (IGSL)

- Directly estimates the likelihood of an event being commented on
- Self-training iterations to improve estimates
- Uses events not associated with any NL as negative evidence

IGSL Example

Natural Language Commentary	Meaning Representation
Purple7 passes to Purple 3	Ballstopped Kick (purple7) Pass (purple7 , purple3) Ballstopped Kick (pink3)
Purple3 loses the ball to Pink4	Turnover (purple3 , pink4)
Pink4 passes out to Pink8	Kick (pink4) Pass (pink4 , pink8)

IGSL Example

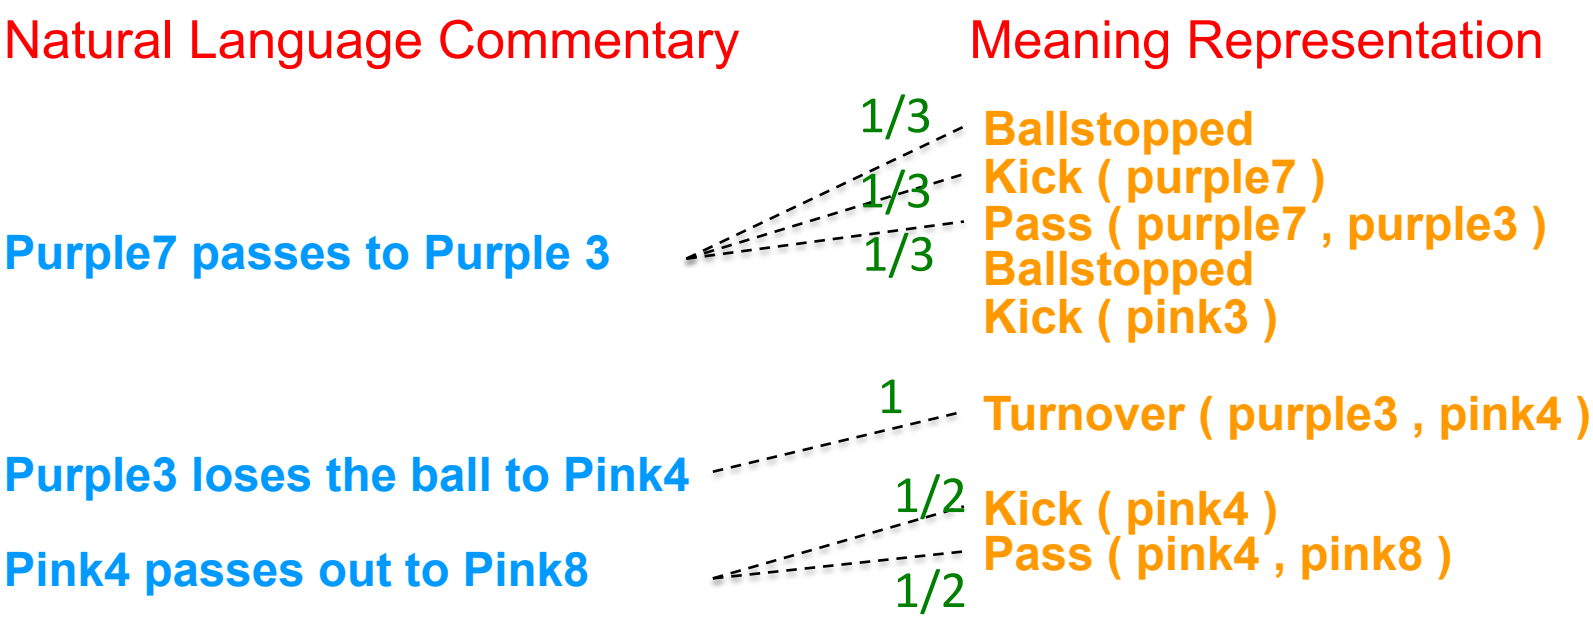

IGSL Example

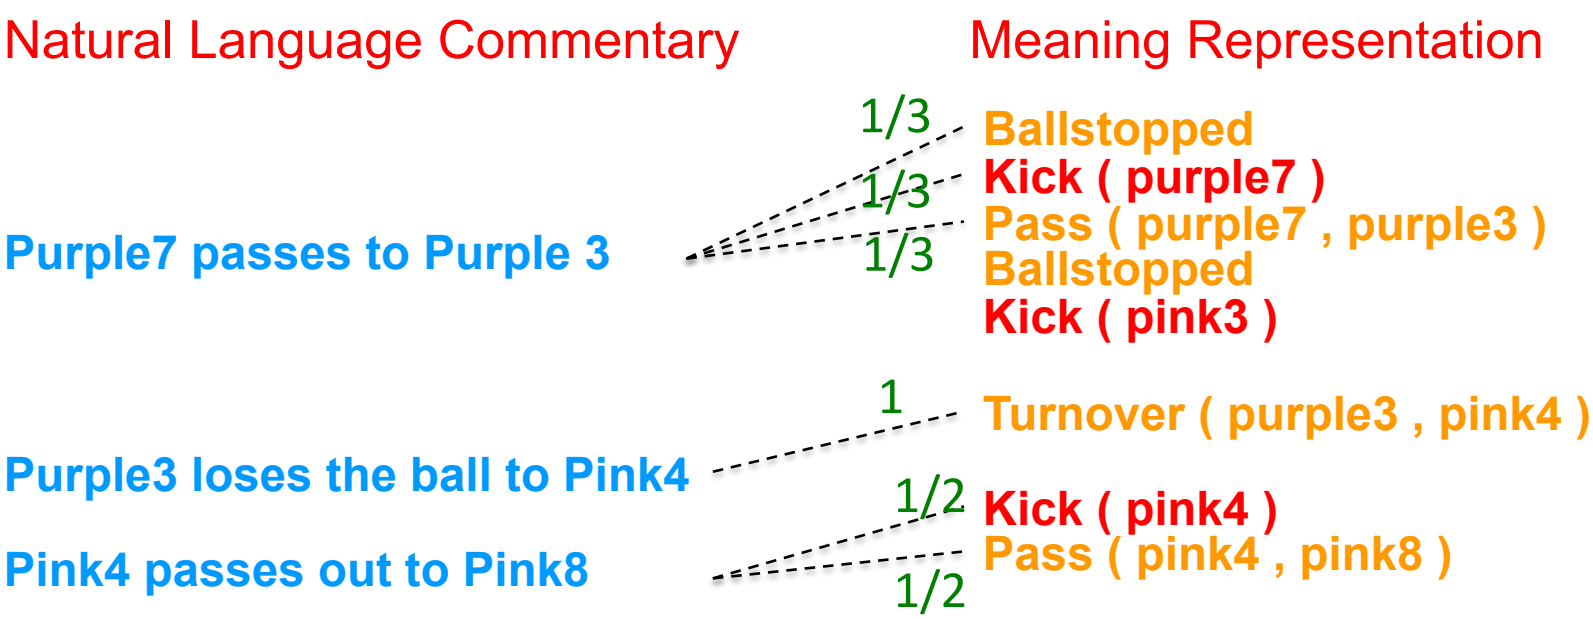

P (Kick) = (1/3 + 1/2) / 3 = 0.278

IGSL Example

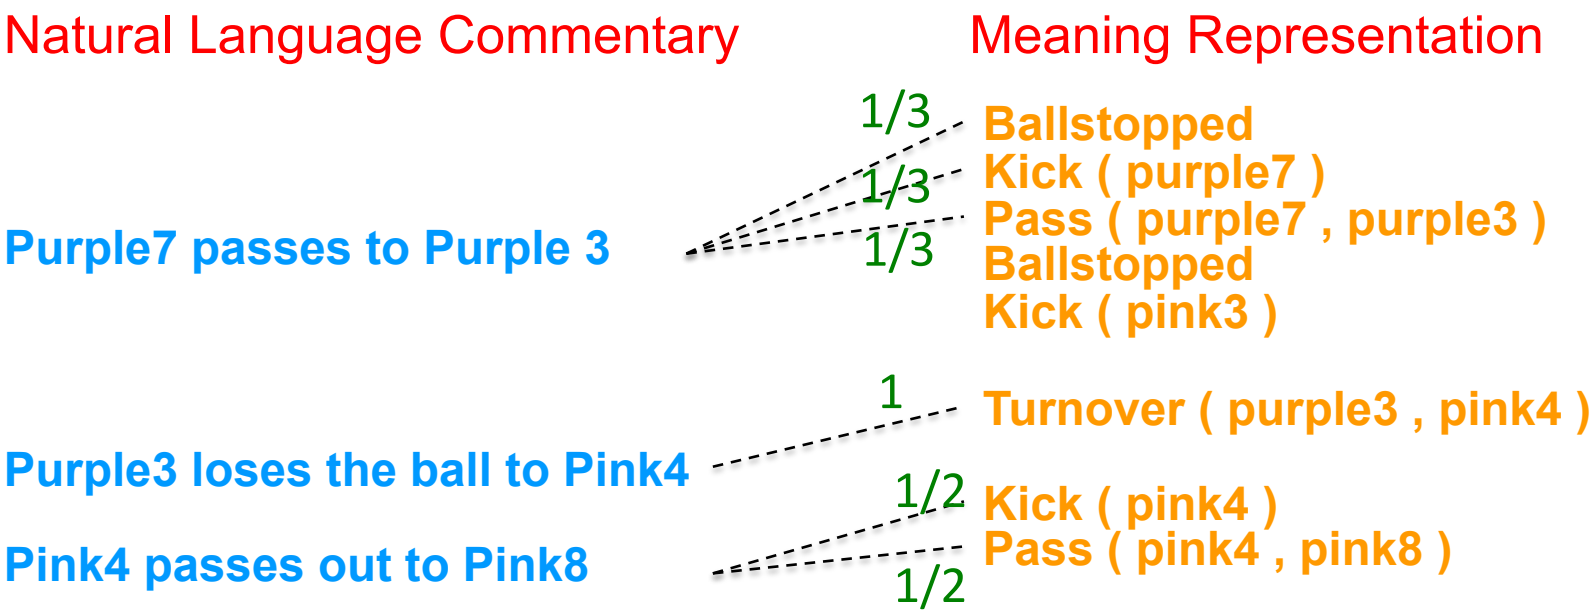

P (Kick) = (1/3 + 1/2) / 3 = 0.278
P (Pass) = (1/3 + 1/2) / 2 = 0.417
P (Turnover) = (1) / 1 = 1
P (Ballstopped) = (1/3) / 2 = 0.17

IGSL Example

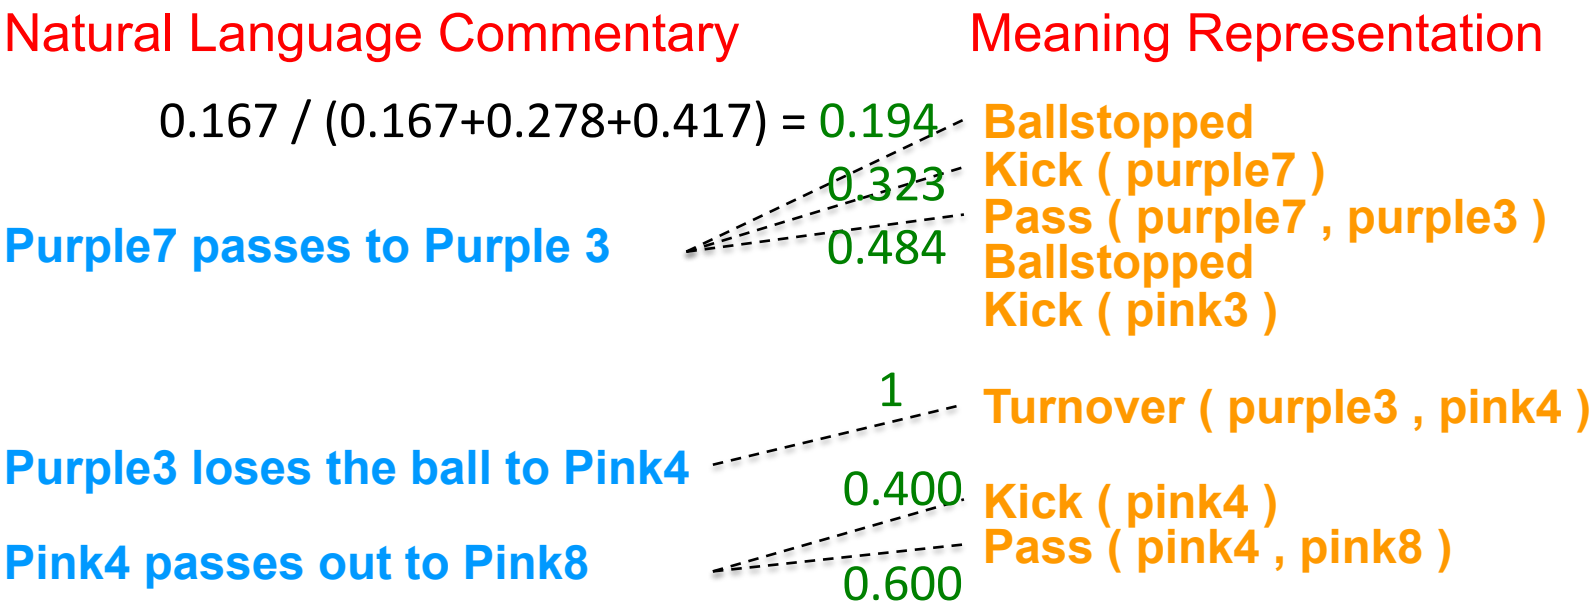

P (Kick) = (1/3 + 1/2) / 3 = 0.278
P (Pass) = (1/3 + 1/2) / 2 = 0.417
P (Turnover) = (1) / 1 = 1
P (Ballstopped) = (1/3) / 2 = 0.17

IGSL Example

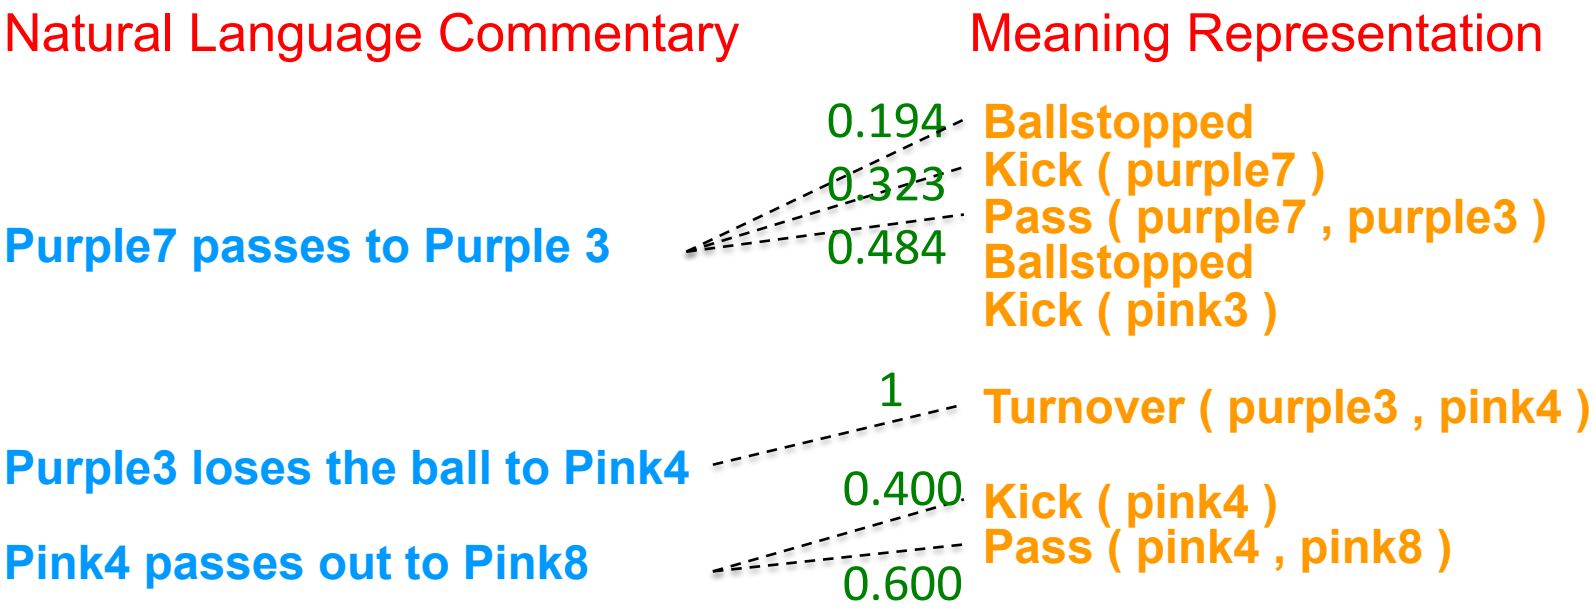

P (Kick) = (0.323 + 0.400) / 3 = 0.241
P (Pass) = (0.484 + 0.600) / 2 = 0.542
P (Turnover) = (1) / 1 = 1
P (Ballstopped) = (0.194) / 2 = 0.097

IGSL Example

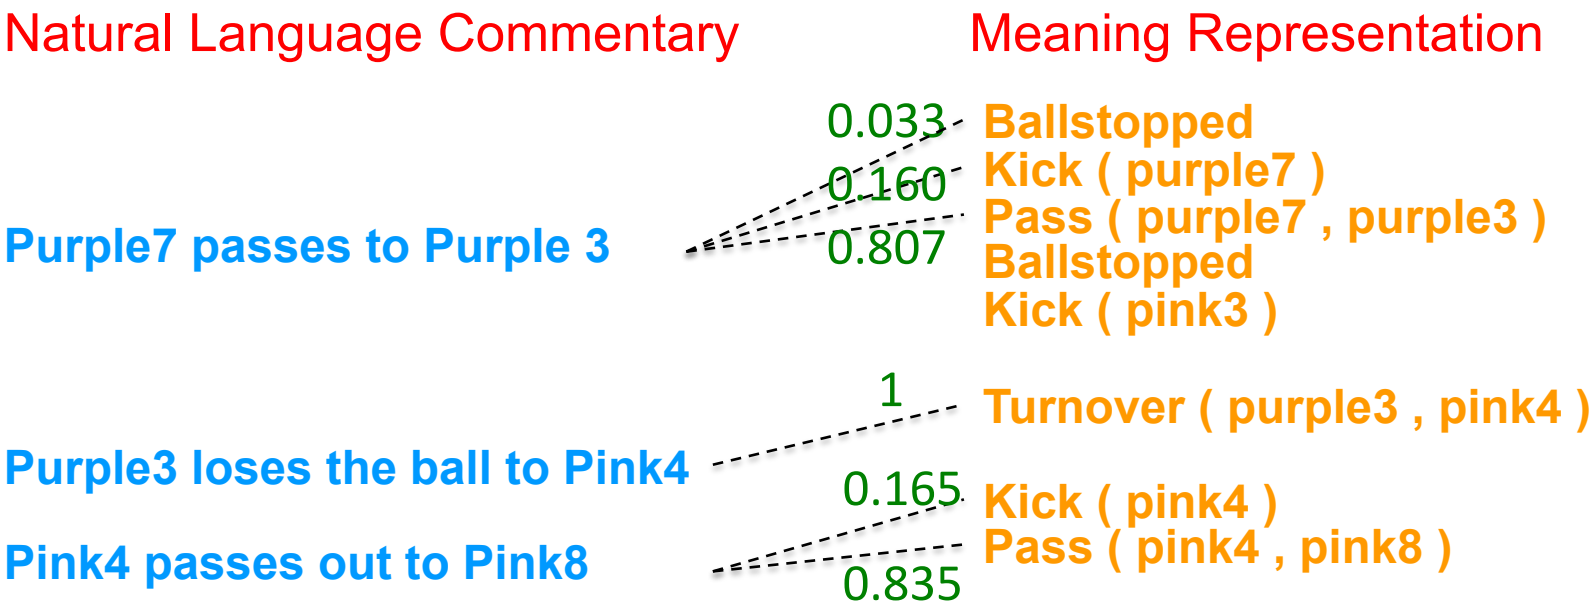

P (Kick) = 0.148
P (Pass) = 0.748
P (Turnover) = (1) / 1 = 1
P (Ballstopped) = (1/3) / 2 = 0.030

Strategic Generation Performance

- Evaluate how well the system can predict which events a human comments on
- Metric:
  - **Precision**: % of system's annotations that are correct
  - **Recall**: % of gold-standard annotations correctly produced
  - **F-measure**: Harmonic mean of precision and recall

Strategic Generation Results

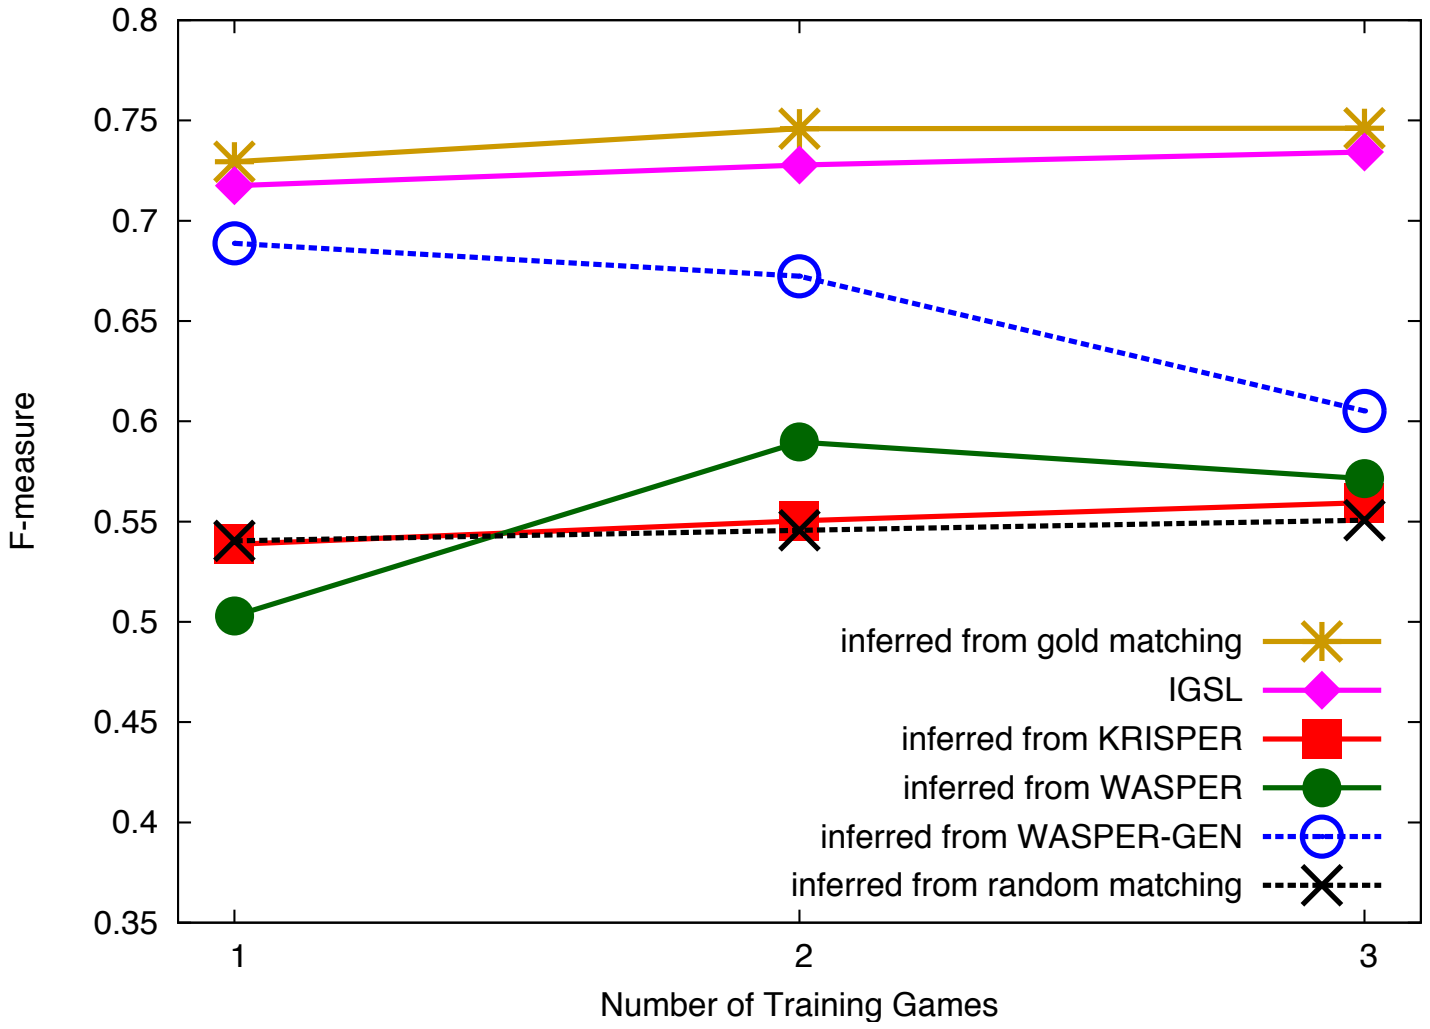

Overview

- Background and related works
- Completed work: Sportscasting
  - Tactical generation
  - Strategic generation
  - **Human evaluation**
- Proposed work: Navigation instructions
- Conclusions

Human Evaluation

- Used Amazon's Mechanical Turk to recruit human judges (~40 judges per video)
- 8 commented game clips
  - 4 minute clips randomly selected from each of the 4 games
  - Each clip commented once by a human, and once by the machine
- Presented in random counter-balanced order
- Judges were not told which ones were human or machine generated

Human Evaluation

Score	English Fluency	Semantic Correctness	Sportscasting Ability
5	Flawless	Always	Excellent
4	Good	Usually	Good
3	Non-native	Sometimes	Average
2	Disfluent	Rarely	Bad
1	Gibberish	Never	Terrible

Demo Clip

- Game clip commentated using WASPER-GEN with IGSL, since this gave the best results for generation.
- FreeTTS was used to synthesize speech from textual output.
- English: http://www.youtube.com/watch?v=L_MIRS7NBpU
- Korean: http://www.youtube.com/watch?v=Dur9K5AiK8Y

Human Evaluation

	Syntax	Semantic	Overall	Human?
2001 Human	3.74	3.59	3.15	20.59%
2001 Machine	3.89	3.81	3.61	40.00%
2002 Human	4.13	4.58	4.03	42.11%
2002 Machine	3.97	3.74	3.29	11.76%
2003 Human	3.54	3.73	2.61	13.51%
2003 Machine	3.89	4.26	3.37	19.30%
2004 Human	4.03	4.17	3.54	20.00%
2004 Machine	4.13	4.38	4.0	56.25%
Average Human	3.86	4.03	3.34	24.31%
Average Machine	3.94	4.03	3.48	26.76%

Overview

- Background and related works
- Completed work: Sportscasting
  - Tactical generation
  - Strategic generation
  - Human evaluation
- **Proposed work: Navigation instructions**
- Conclusions

Referential vs. Functional Meanings

- Referential meanings
  - Describe objects and events in the world
  - Completed work on learning to sportscast
- Functional meanings
  - Aim to achieve some actions in the world
  - Proposed work on learning navigation instructions

Challenge on Generating Instructions in Virtual Environments (GIVE)

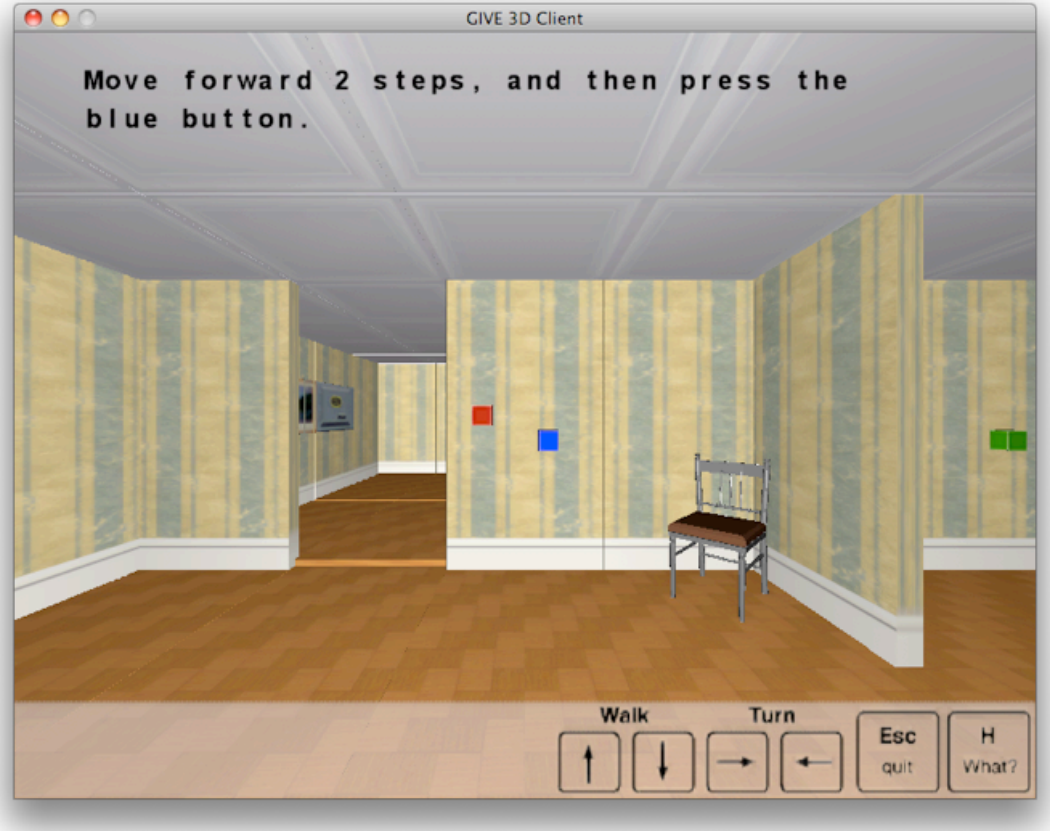

http://www.give-challenge.org/research/

Learning Approach

- Passive learning
  - Observes human instructor guiding a human follower
- Interactive learning as follower
  - Tries to follow human instructions
- Interactive learning as instructor
  - Generates instructions to guide human follower

Navigation Task

- Two participants: instructor and follower
- Given: starting location and destination
- Instructor: Give directions for navigating
- Follower: Follows direction
- Success if follower reaches the intended destination
- Data contains ~800 instructions for 3 virtual environments [MacMahon et al., 2006]

Evaluations

- Task completion
  - Did the follower reach the destination?
- Efficiency
  - How long and how many steps did it take to complete the task
- Partial correctness
  - How much of the task did the follower complete

Challenges

- Many different instructions for the same task
  - Describe different actions
  - Use different parameters
  - Different ways to describe the same parameters
- Hidden MRs
  - Needs to infer the MR from observed actions
- Exponential number of possible MRs

Environment

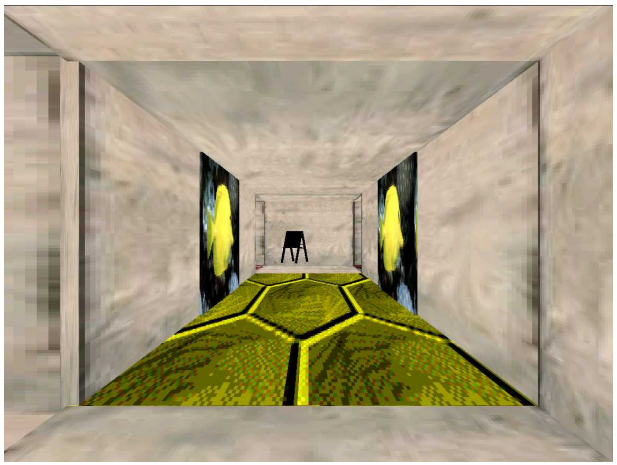

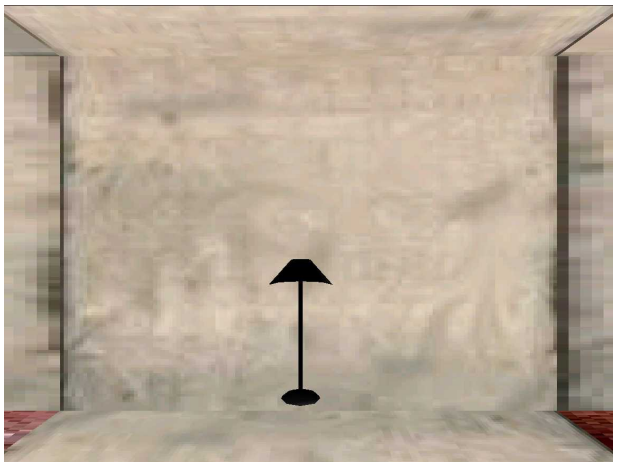

Environment

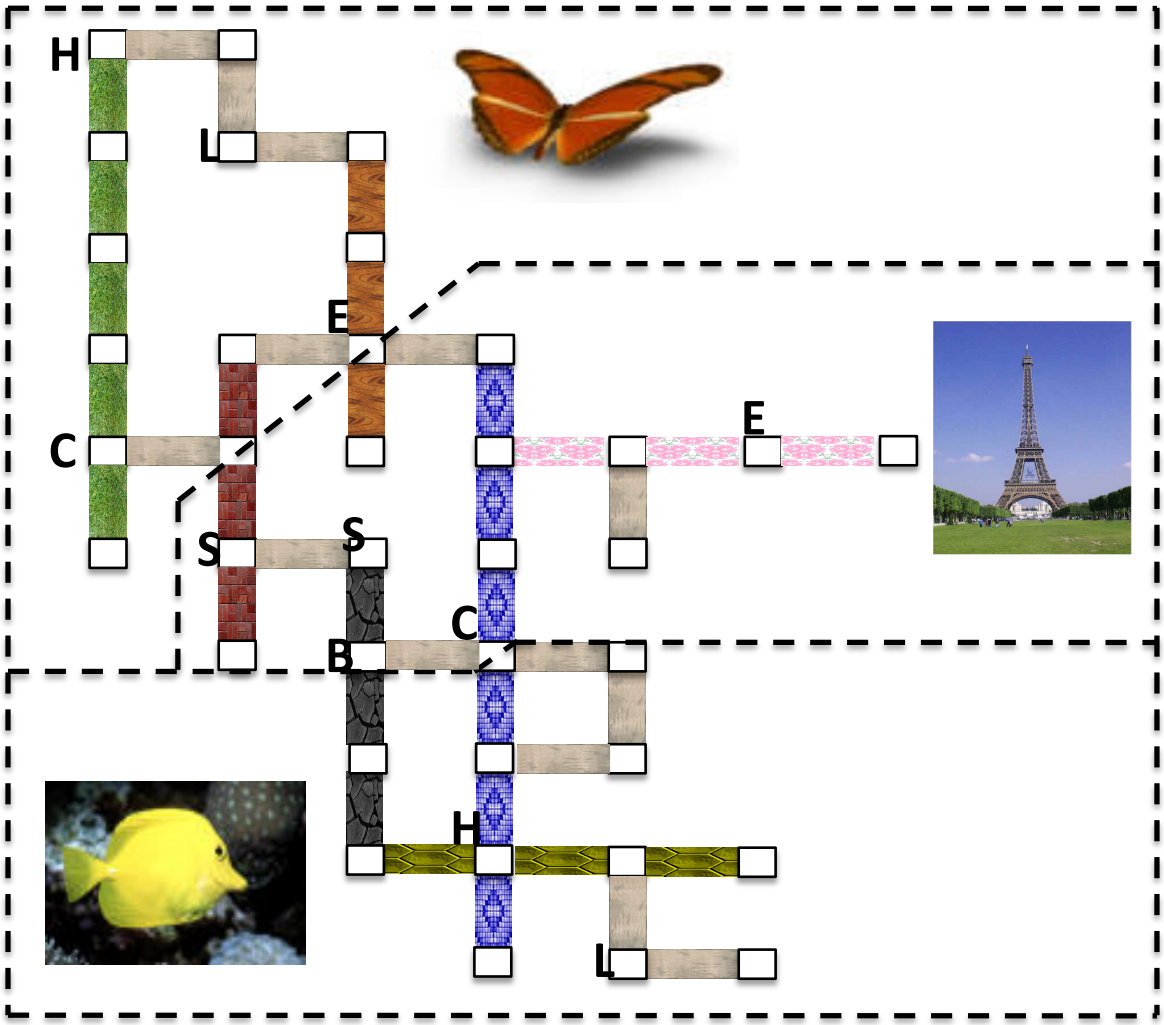

H – Hat Rack

L – Lamp

E – Easel

S – Sofa

B – Barstool

C - Chair

Example of Instructions

Start: 3, End: 4

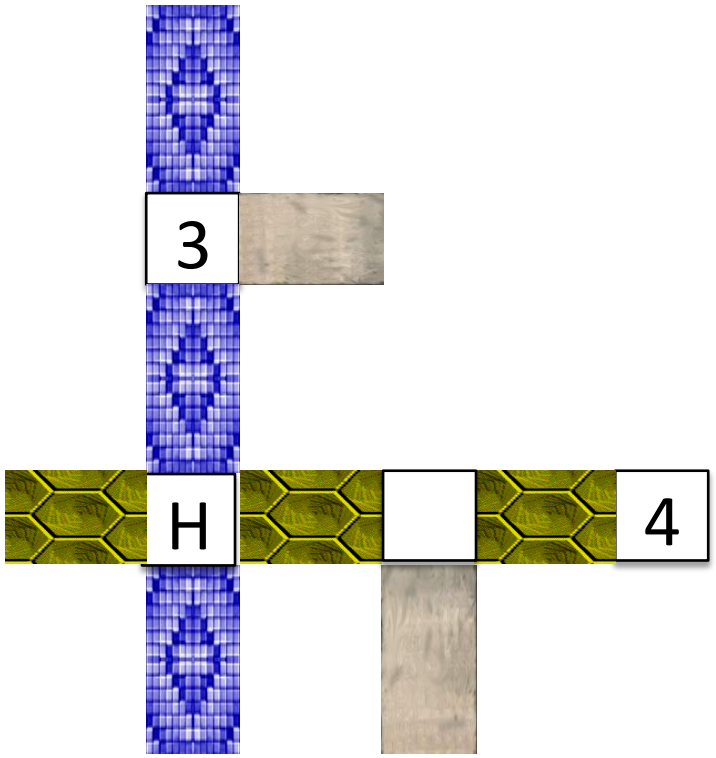

- Take your first left. Go all the way down until you hit a dead end.
- Go towards the coat hanger and turn left at it. Go straight down the hallway and the dead end is position 4.
- Walk to the hat rack. Turn left. The carpet should have green octagons. Go to the end of this alley. This is p-4.
- Walk forward once. Turn left. Walk forward twice.

Example of Instructions

Start: 3, End: 4

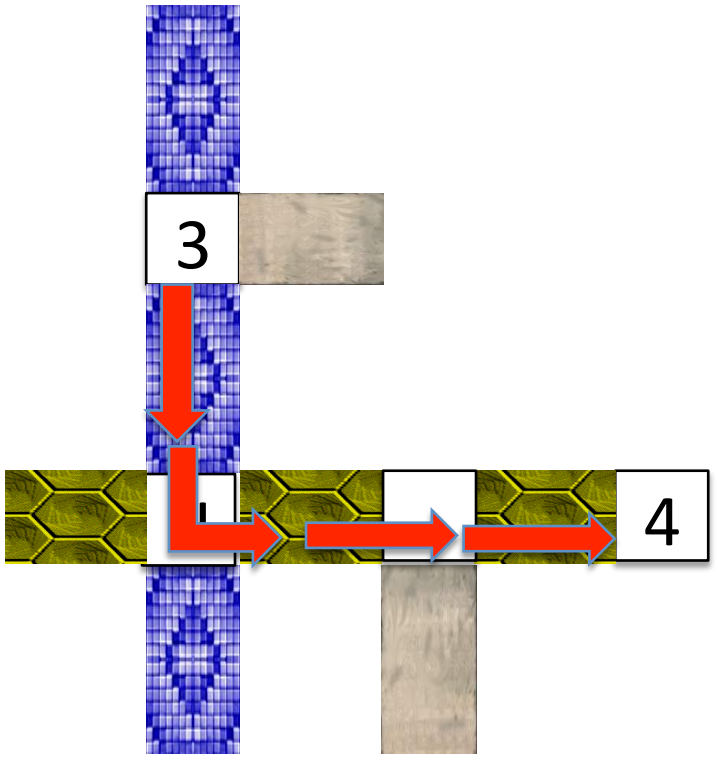

Observed primitive actions:
Forward, Left, Forward, Forward

- Take your first left. Go all the way down until you hit a dead end.
- Go towards the coat hanger and turn left at it. Go straight down the hallway and the dead end is position 4.
- Walk to the hat rack. Turn left. The carpet should have green octagons. Go to the end of this alley. This is p-4.
- Walk forward once. Turn left. Walk forward twice.

Overall System

Instructions	… Walk forward twice …
MR	… Travel(Distance = 2) …
Primitive actions	… Forward, Forward …

Overall System

Instructions	… Walk forward twice …
MR	… Travel(Distance = 2) …
Primitive actions	… Forward, Forward …

Overall System

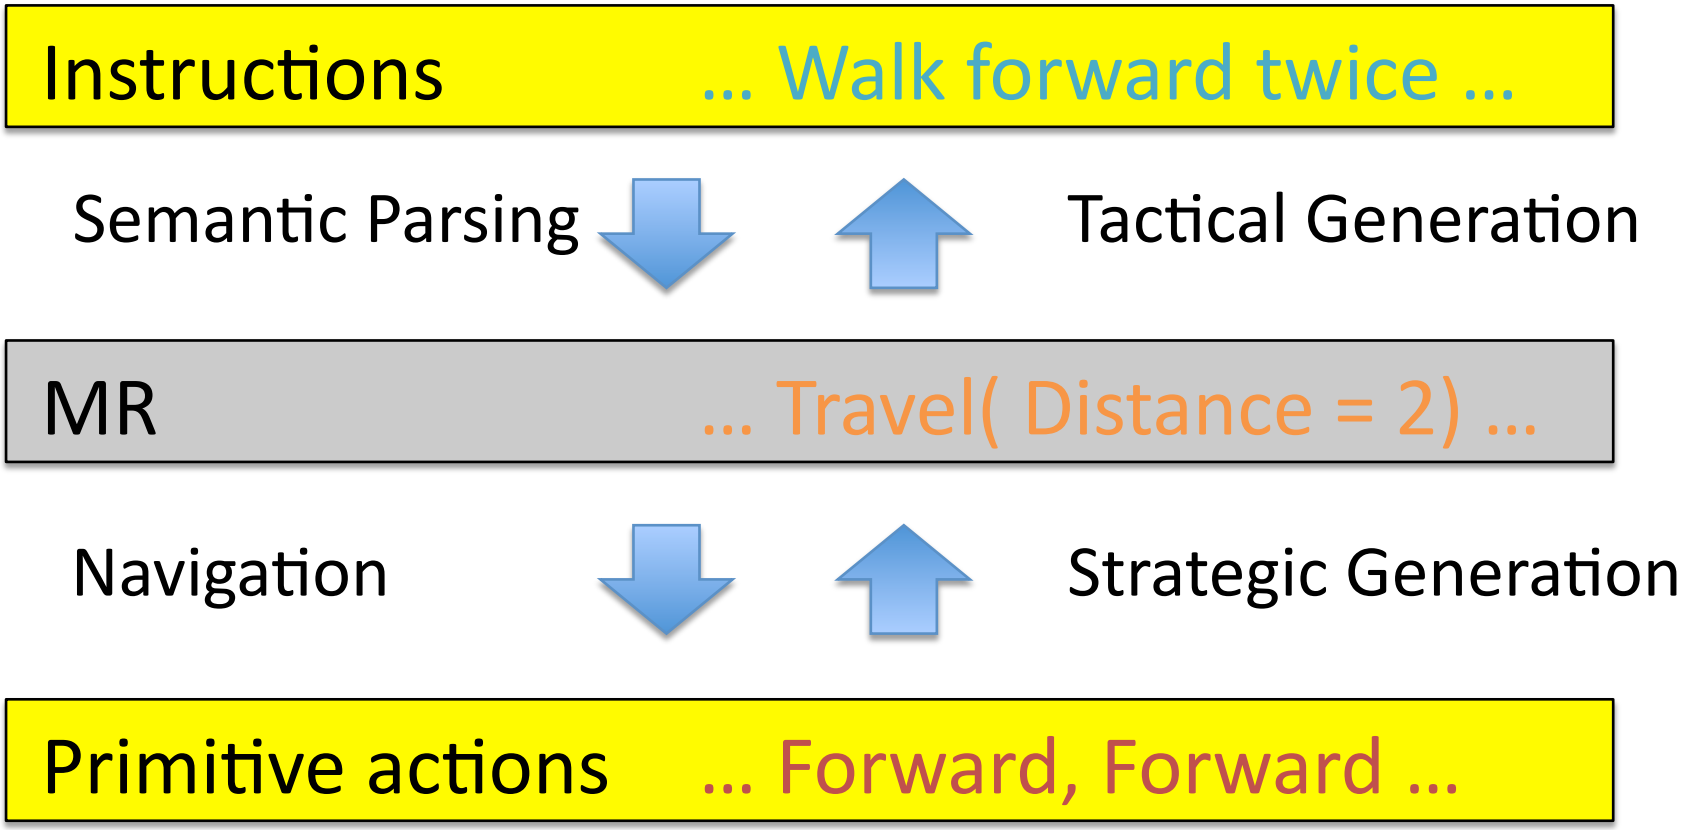

Overall System

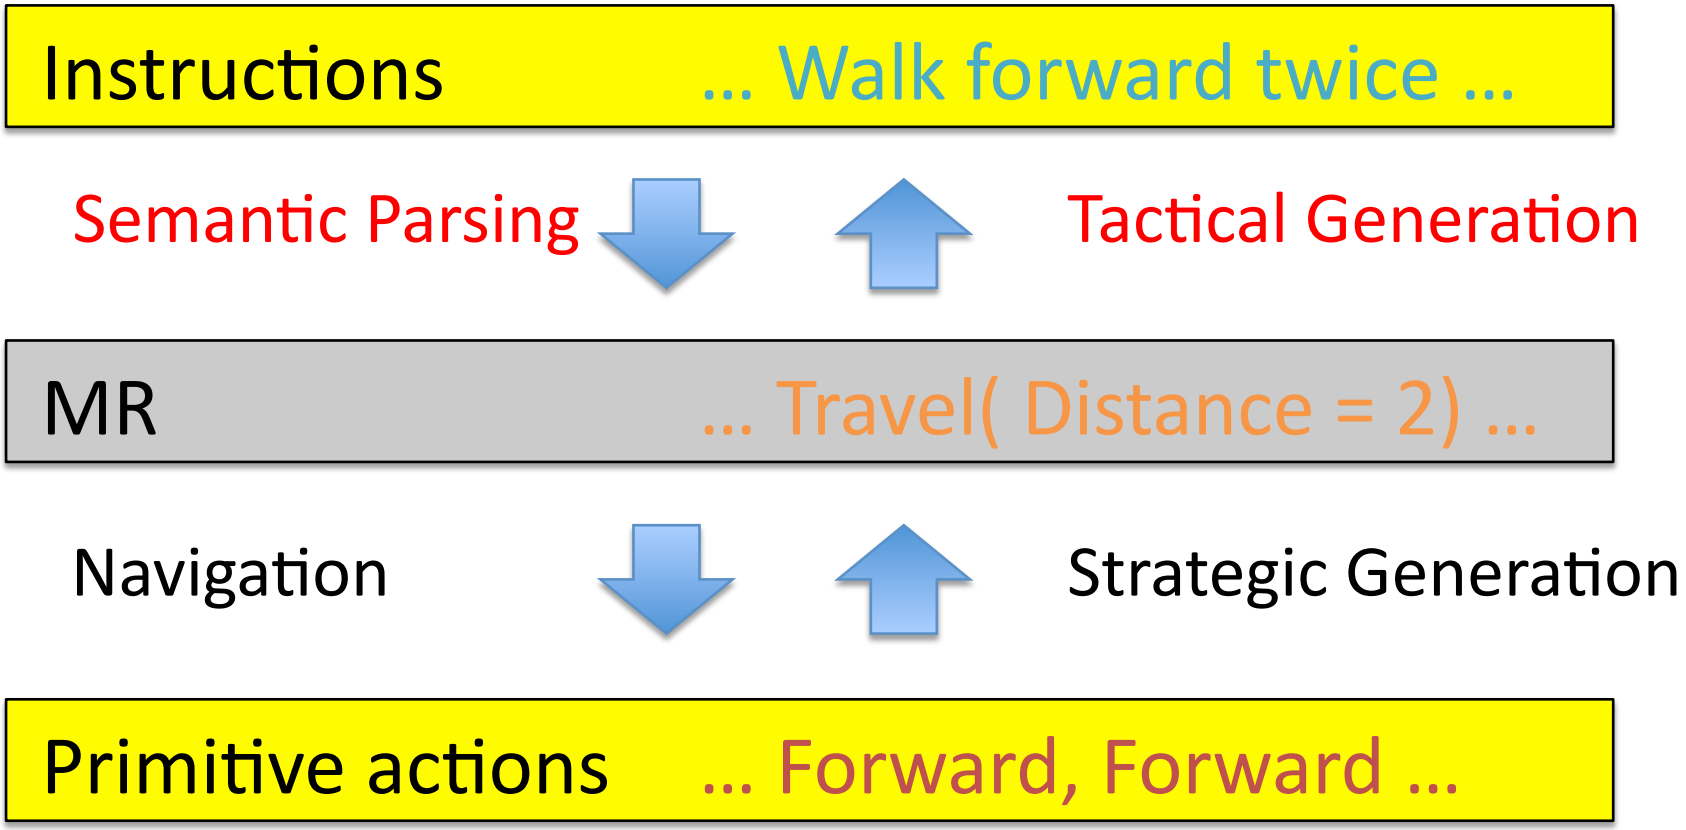

Modeling the Instruction Parsing Process

- Use semantic parser to produce a set of good MRs from the instructions
- Use the navigation component to verify which of these MRs result in the correct actions
- Refine the MRs if none of them are correct

Training Initial Semantic Parser

- Construct the *most specific MR*
- Overestimates the details

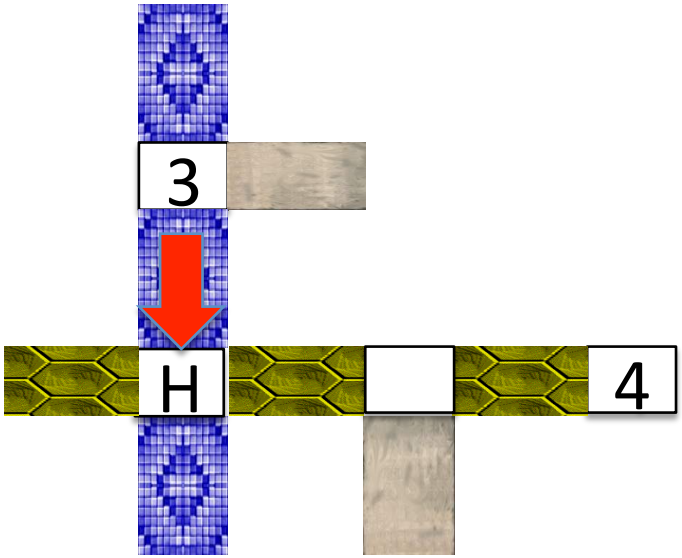


```
Travel(
  Precondition=(Right=Wall,
       Left=Concrete Hall, Front=Blue Hall,
       Back=Blue Hall),
  Distance=1,
  Until=(Intersection(Order=1,
     Current Path=Blue Hallway,
     Cross Path=Yellow Hallway),
     Hat Rack),
  Postcondition=(Right=Yellow Hall,
    Left=Yellow Hall,Front=Blue Hall,
    Back=Blue Hall)
)
```

Refinement

- Modifies an MR until it produces the correct actions
- First remove any parts that do not appear in the most specific MR
- Then systematically add parts of the most specific MR
- Prefers the least amount of modification
  - Want a MR closer to the original parse
- Prefers shortest MR
  - Avoid superfluous connections

Interactive Learning

- The system can participate in the navigation task as instructor or follower
- Feedback from human partner helps fix errors in understanding
- Reweigh the rules that led to the positive or negative feedback

Overview

- Background and related works
- Completed work: Sportscasting
  - Tactical generation
  - Strategic generation
  - Human evaluation
- Proposed work: Navigation instructions
- **Conclusions**

Conclusion

- Current language learning work uses expensive, annotated training data.
- We have developed a language learning system that can learn from language paired with an ambiguous perceptual environment.
- We have evaluated it on the task of learning to sportscast simulated RoboCup games.
- The proposed future work aims to solve the problem of learning how to give and receive navigational instructions in a virtual world

Backup Slides

KRISPER

1. Assume every possible meaning for a sentence is correct

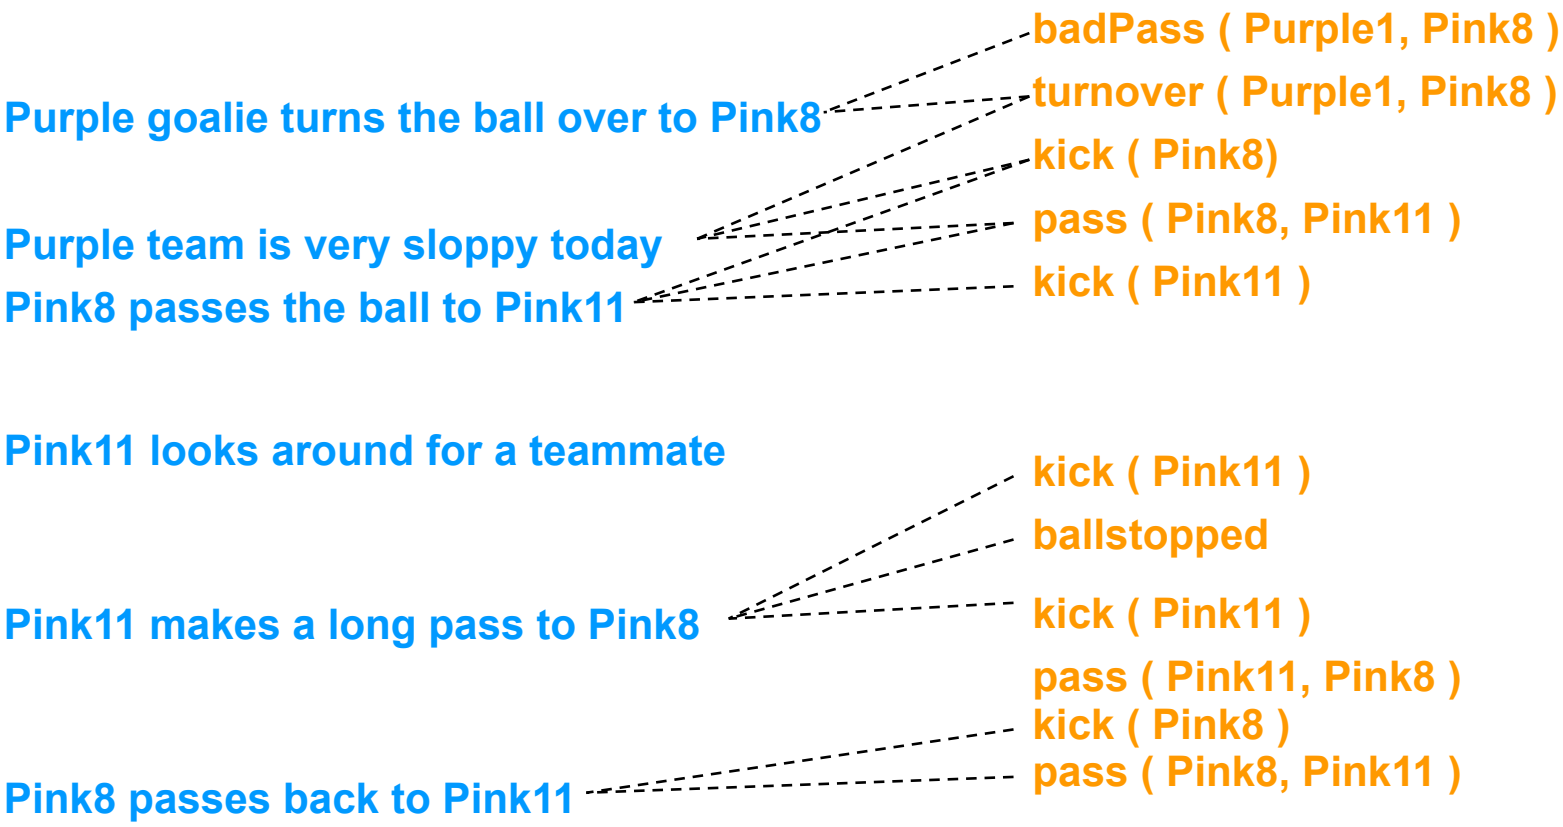

KRISPER

1. Assume every possible meaning for a sentence is correct

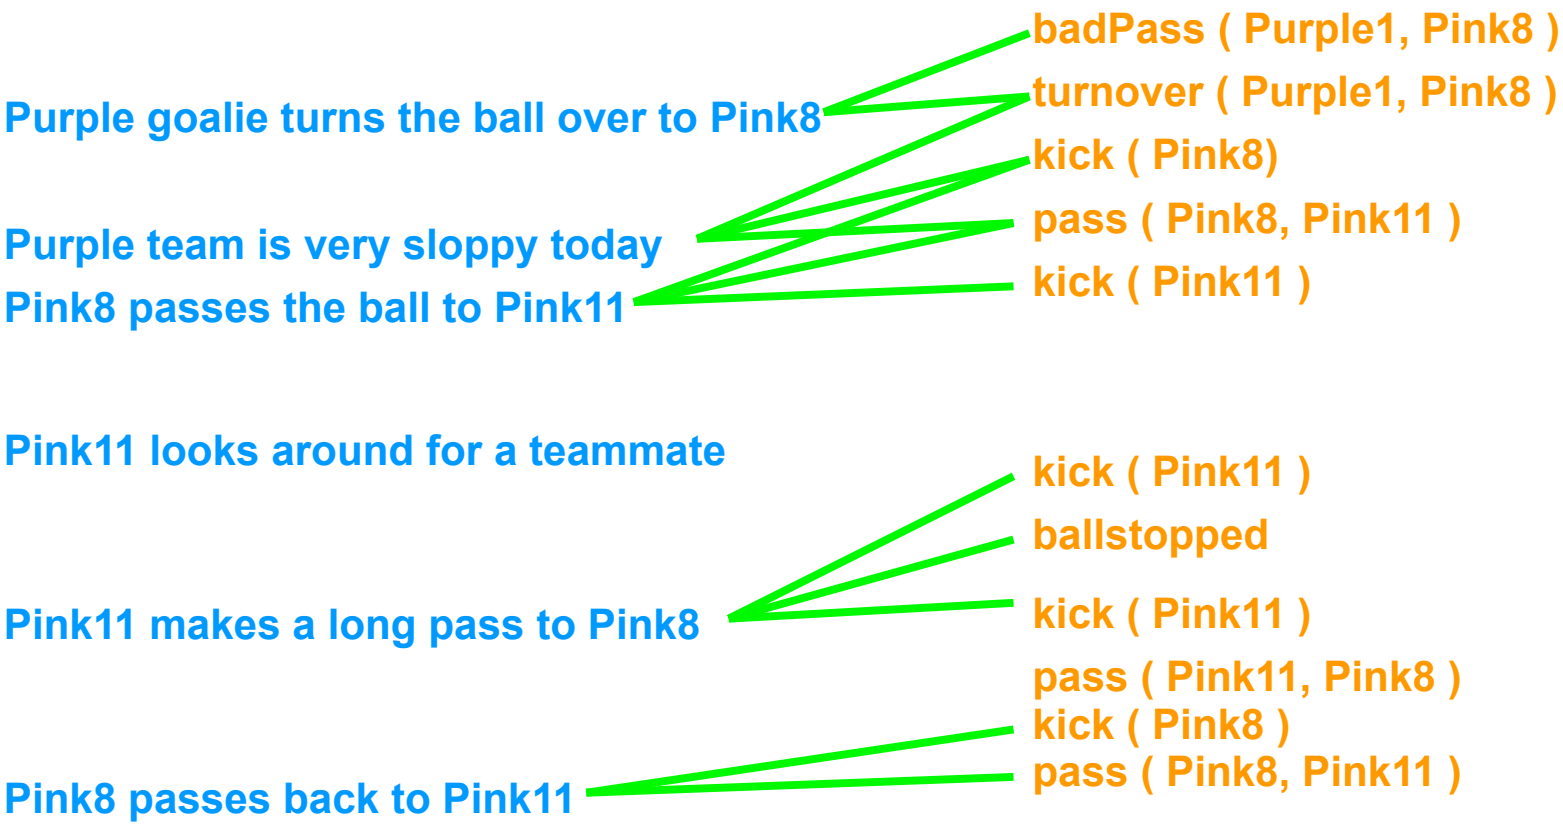

KRISPER

2. Resulting NL-MR pairs are weighted and given to semantic parser learner

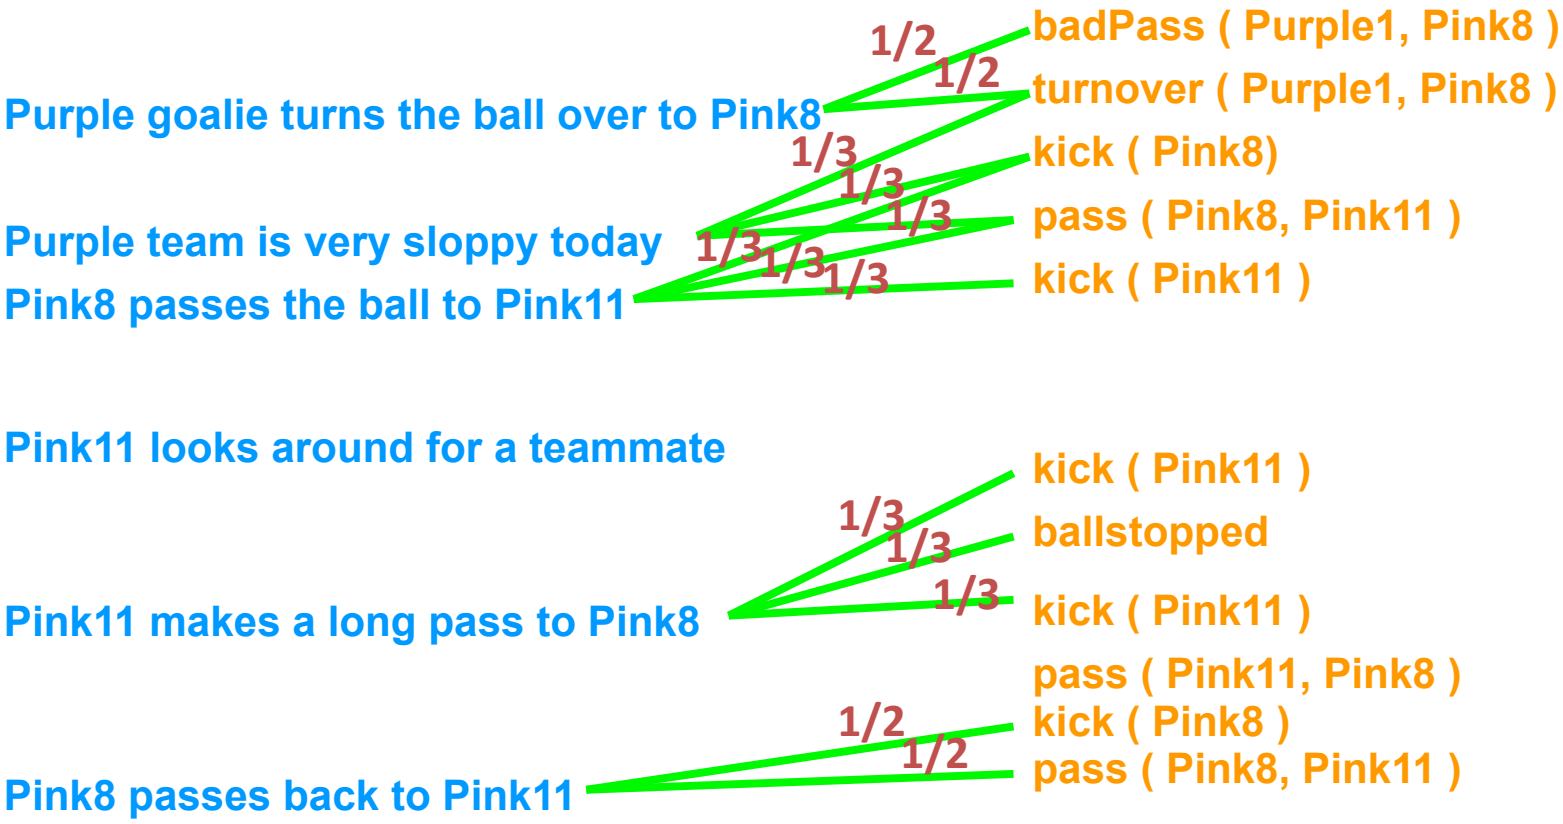

KRISPER

3. Estimate the confidence of each NL-MR pair using the resulting trained semantic parser

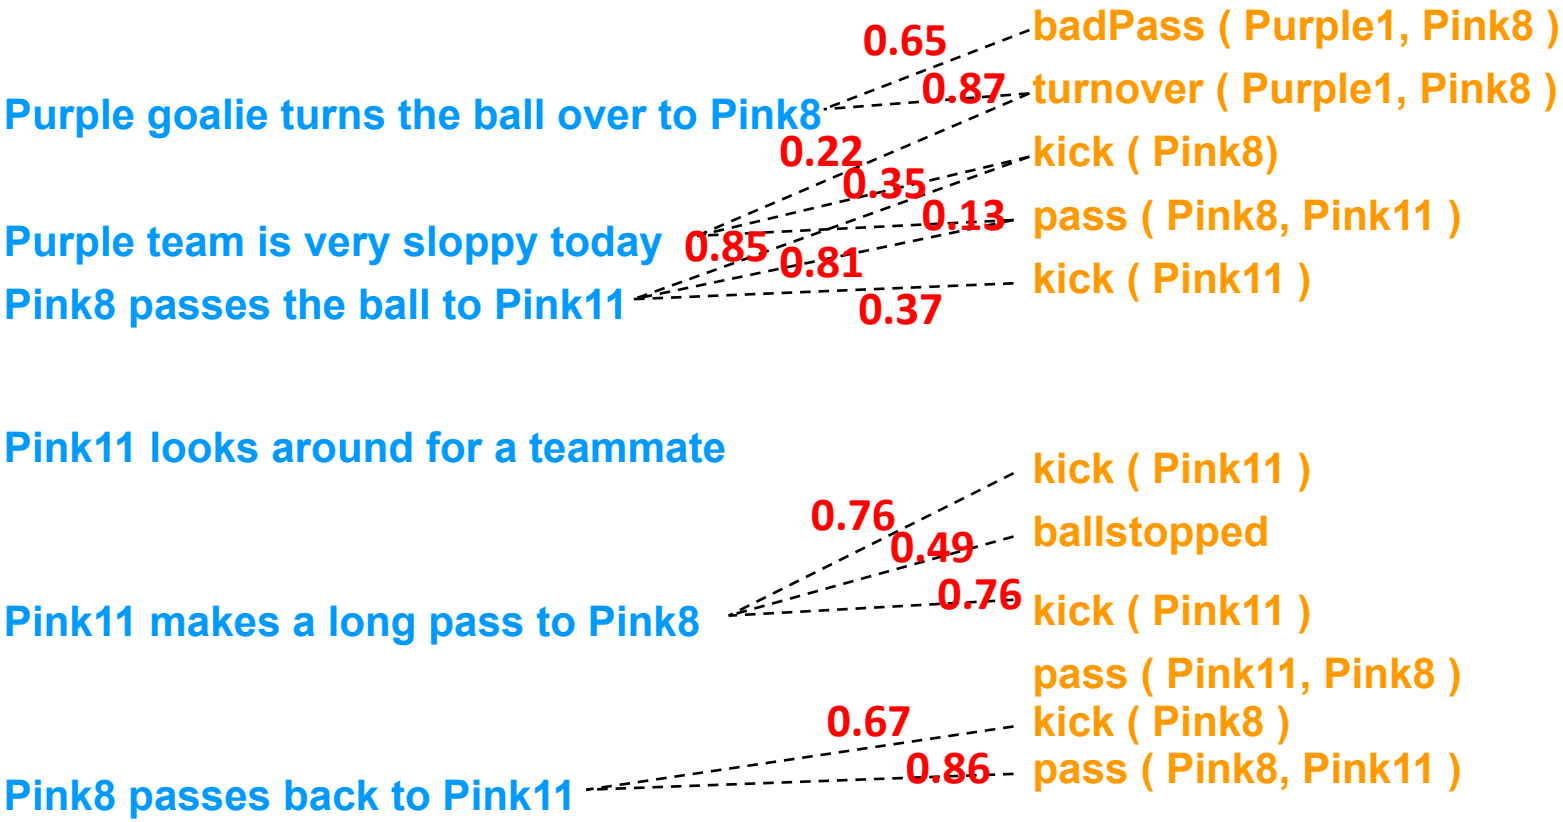

KRISPER

4. Use *maximum weighted matching* on a bipartite graph to find the best NL-MR pairs [Munkres, 1957]

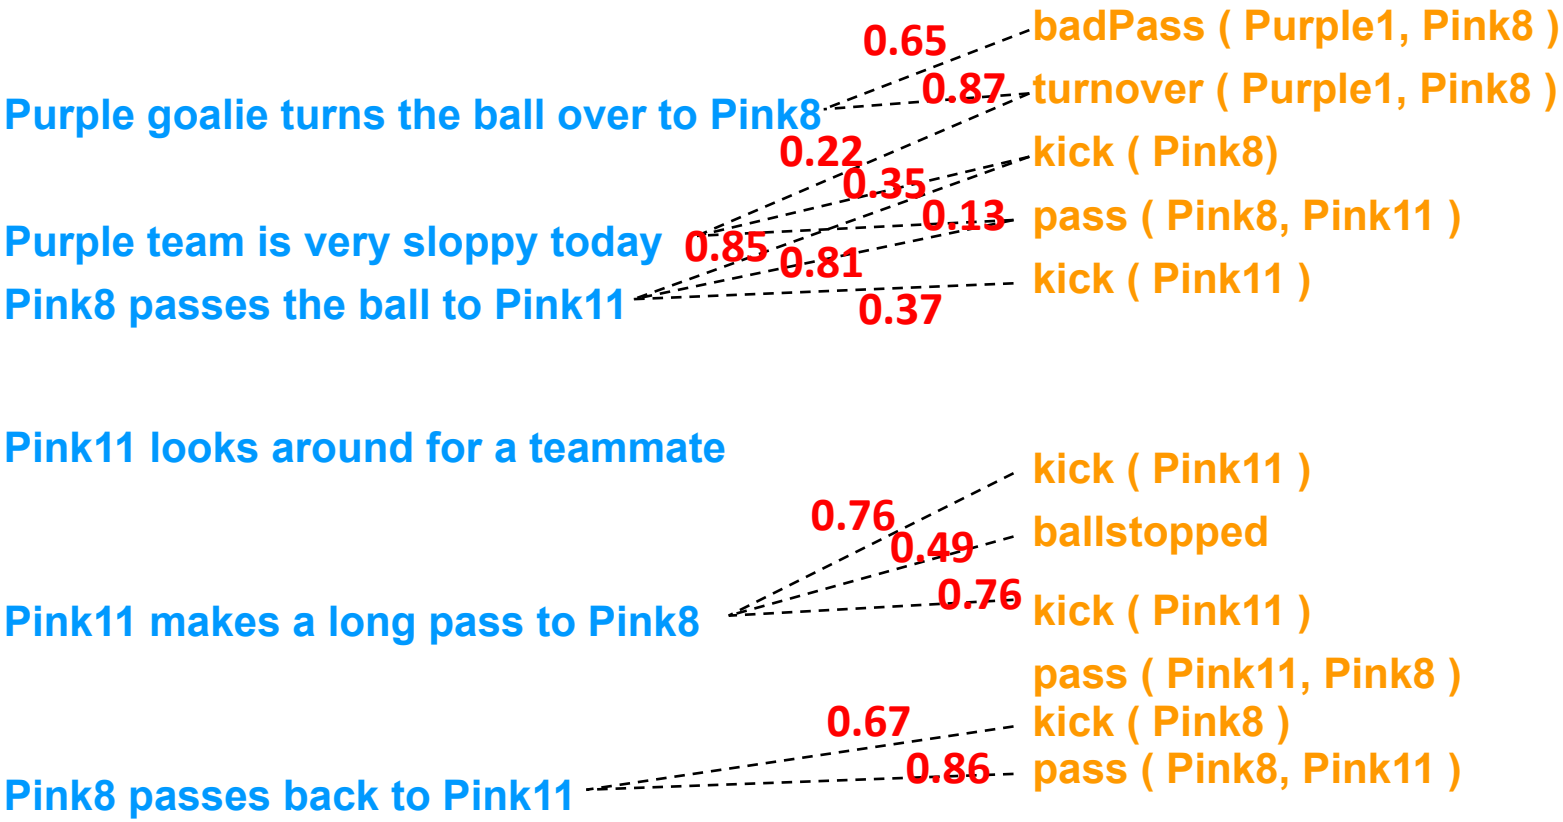

KRISPER

4. Use *maximum weighted matching* on a bipartite graph to find the best NL-MR pairs [Munkres, 1957]

KRISPER

5. Give the best pairs to the semantic parser learner in the next iteration, and repeat until convergence