Learning Perceptually Grounded Word Meanings from Unaligned Parallel Data

Stefanie Tellex

Pratiksha Thaker, Joshua Joseph, Thomas Kollar, Steven Dickerson, Matthew Walter, Robin Deits, Ashis Gopal Banerjee, Dimitar Simeonov, Seth Teller, Nicholas Roy

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

CSAIL

How do we tell them what to do?

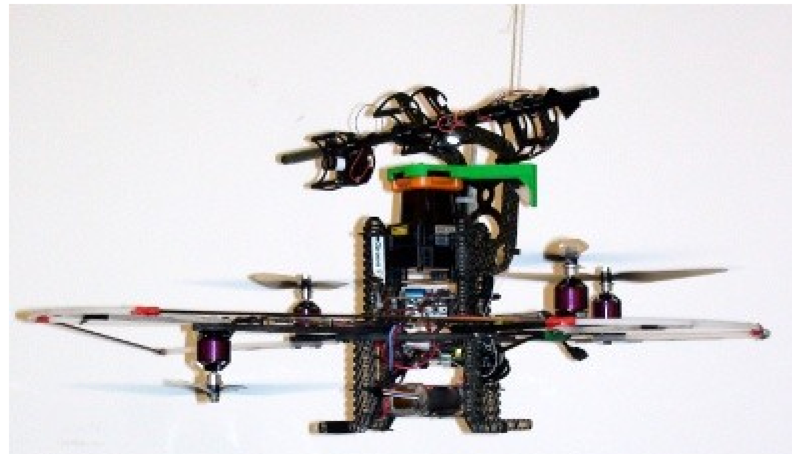

Existing Interfaces

Courtesy of iRobot, Inc.

Agile (MIT)

Courtesy of AeroVironment, Inc.


Return to Camp Charlie and take the eastern path to the landing zone.

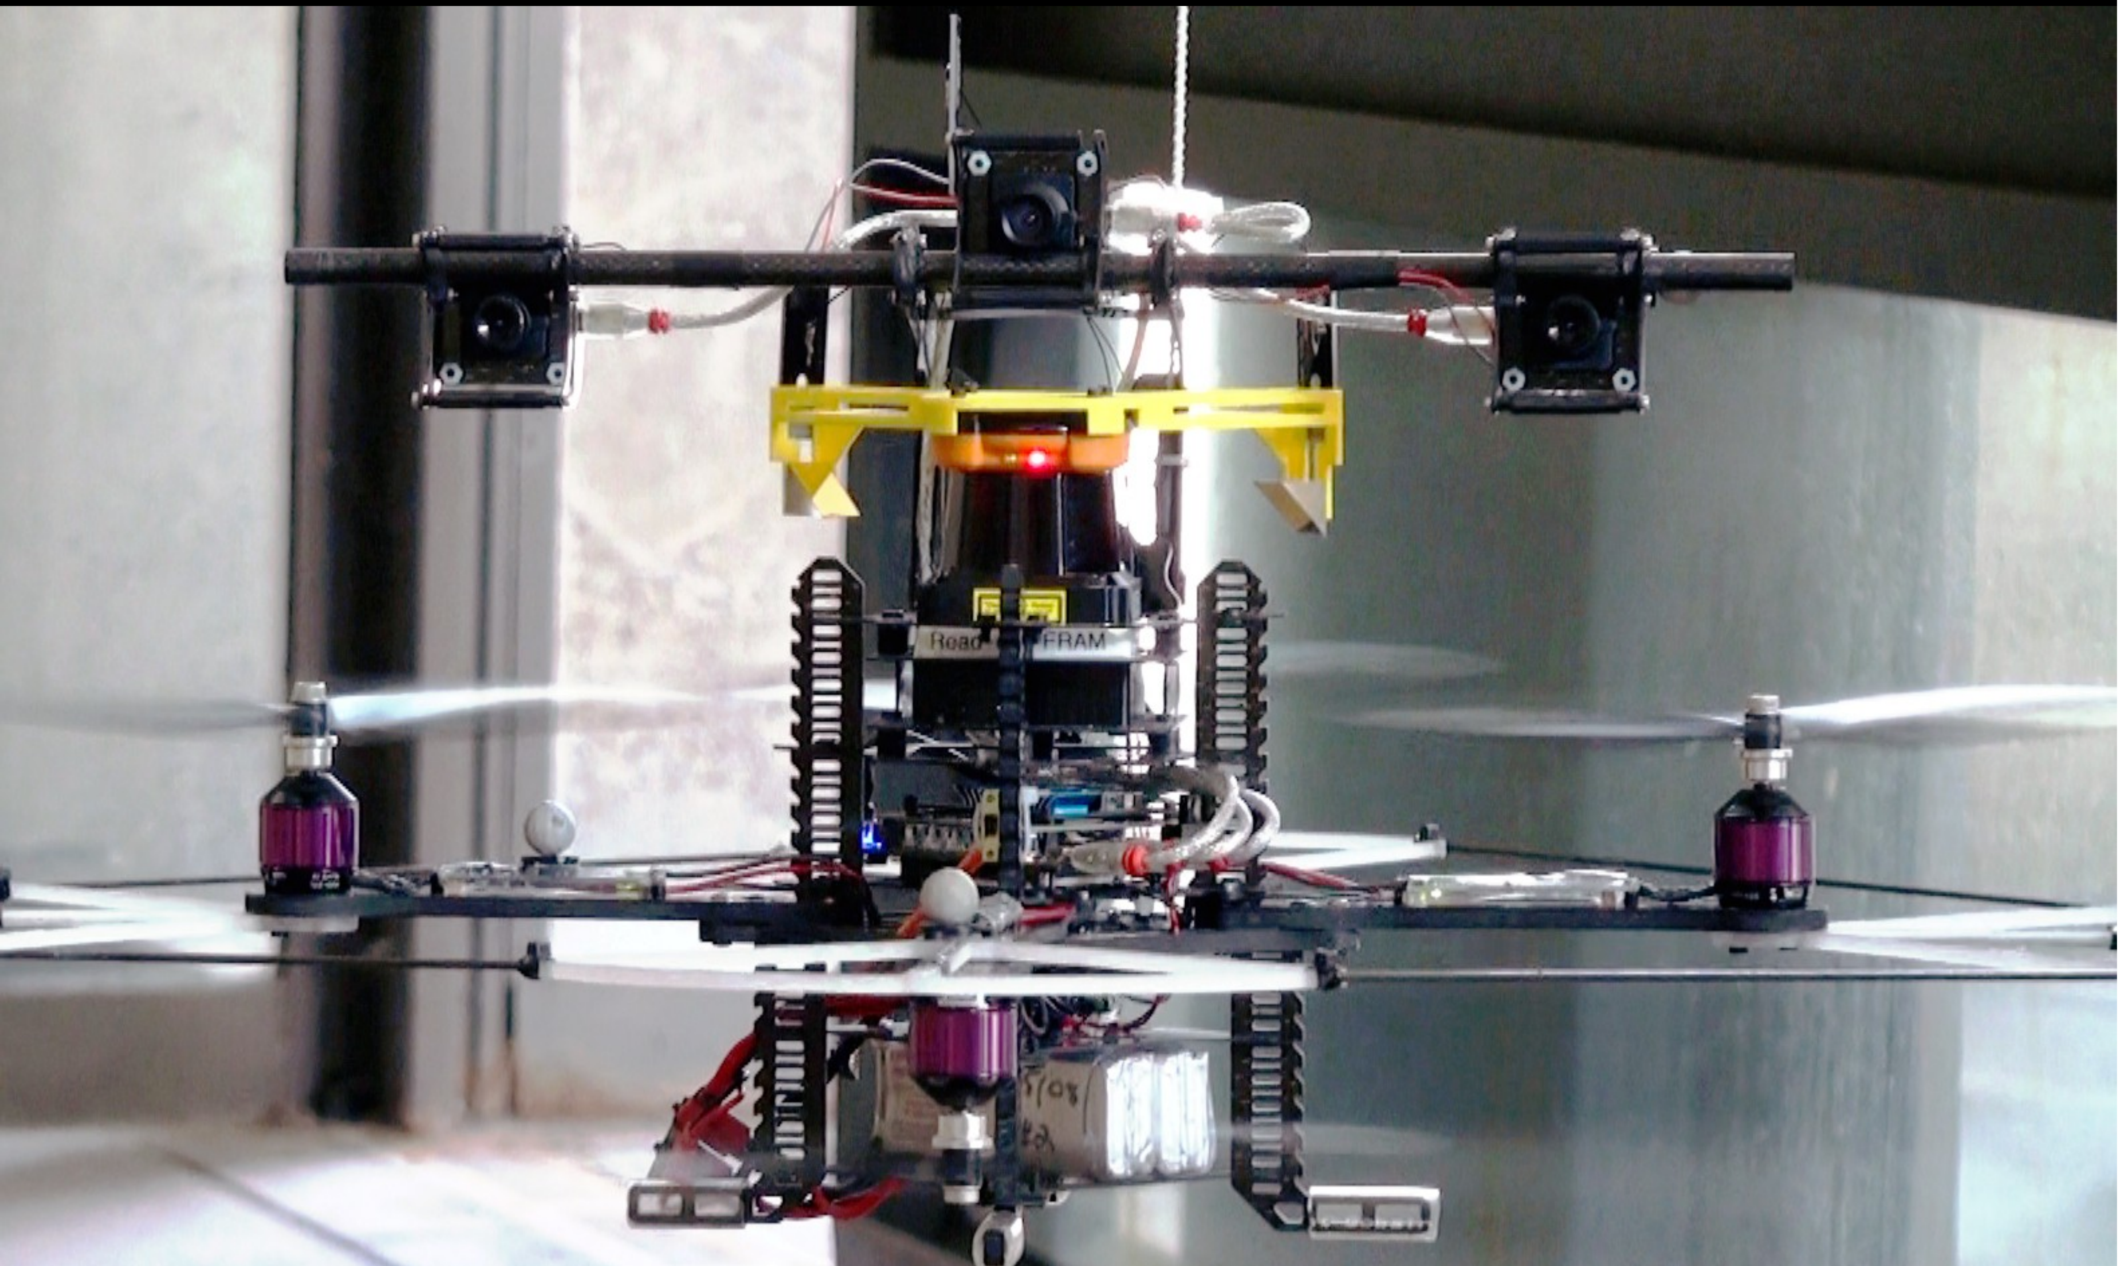


Put the pallet on the truck.
MIT
MIT
CSAIL
BAE SYSTEMS
TOYOTA
TOYOTA
MIT

Symbol Grounding Problem

“The pallet of boxes on the left.”


Put the pallet on the truck.
MIT
CSAIL
BAE SYSTEMS
TOYOTA
MIT

another tyre pallet on the trolley.

Place the pallet of tires on the left side of the trailer.

lift the tire pallet to the truck

Arrange tire pallet to the truck.

Please lift the set of six tires up and set them on the trailer, to the right of the set of tires already on it.

Load the skid right next to the other skid of tires on the trailer.

Place a second pallet of tires on the trailer.

Place the pallet of tires that is on the forklift next to the pallet of tires that is already loaded on the trailer.

Lift tire pallet. Move to unoccupied location on truck. Lower tire pallet. Reverse to starting location. Lower forks. End.

Put the tire pallet on the trailer.

Lift the tire pallet in the air, then proceed to deposit it to the right of the
tire pallet already on the table right in front of you.

Place the pallet of tires on the right side of the truck.

Lift the tire pallet and proceed forward to set it on the platform directly ahead, to the right of the tire pallet already there.

another tyre pallet on the trolley.

Place the pallet of tires on the left side of the trailer.

lift the tire pallet to the truck

Arrange tire pallet to the truck.

Please lift the set of six tires up and set them on the trailer, to the right of the set of tires already on it.

Load the skid right next to the other skid of tires on the trailer.

Place a second pallet of tires on the trailer.

Place the pallet of tires that is on the forklift next to the pallet of tires that is already loaded on the trailer.

Lift tire pallet. Move to unoccupied location on truck. Lower tire pallet. Reverse to starting location. Lower forks. End.

Put the tire pallet on the trailer.

Lift the tire pallet in the air, then proceed to deposit it to the right of the
tire pallet already on the table right in front of you.

Place the pallet of tires on the right side of the truck.

Lift the tire pallet and proceed forward to set it on the platform directly ahead, to the right of the tire pallet already there.

another tyre pallet on the trolley.

Place the pallet of tires on the left side of the trailer.

lift the tire pallet to the truck

Arrange tire pallet to the truck.

Please lift the set of six tires up and set them on the trailer, to the right of the set of tires already on it.

Load the skid right next to the other skid of tires on the trailer.

Place a second pallet of tires on the trailer.

Place the pallet of tires that is on the forklift next to the pallet of tires that is already loaded on the trailer.

Lift tire pallet. Move to unoccupied location on truck. Lower tire pallet. Reverse to starting location. Lower forks. End.

Put the tire pallet on the trailer.

Lift the tire pallet in the air, then proceed to deposit it to the right of the
tire pallet already on the table right in front of you.

Place the pallet of tires on the right side of the truck.

Lift the tire pallet and proceed forward to set it on the platform directly ahead, to the right of the tire pallet already there.

another tyre pallet on the trolley.

Place the pallet of tires on the left side of the trailer.

lift the tire pallet to the truck

Arrange tire pallet to the truck.

Please lift the set of six tires up and set them on the trailer, to the right of the set of tires already on it.

Load the skid right next to the other skid of tires on the trailer.

Place a second pallet of tires on the trailer.

Place the pallet of tires that is on the forklift next to the pallet of tires that is already loaded on the trailer.

Lift tire pallet. Move to unoccupied location on truck. Lower tire pallet. Reverse to starting location. Lower forks. End.

Put the tire pallet on the trailer.

Lift the tire pallet in the air, then proceed to deposit it to the right of the
tire pallet already on the table right in front of you.

Place the pallet of tires on the right side of the truck.

Lift the tire pallet and proceed forward to set it on the platform directly ahead, to the right of the tire pallet already there.

How can we make robots robustly understand natural language commands?

How do we learn word meanings?

Załaduj towar na ciężarówkę.

Zdejmij paletę z ciężarówki
i postaw między koparkami.
Przewieź towar z ciężarówki
na drugą stronę parkingu.

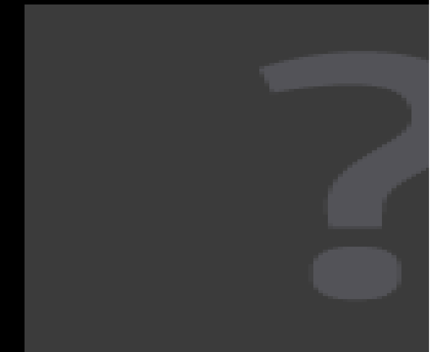

Postaw skrzynię na ziemi.
Opróżnij wózek widłowy.
Odstaw towar.

Podnieś paletę z ziemi i
postaw na ciężarówce.
Załaduj towar na
ciężarówkę.

Take the pallet off the truck and put it between the bulldozers.
Move the cargo from the truck to the other side of the lot.

Put the pallet on the ground.
Unload the forklift.
Put down the load.

Pick up the pallet off the ground and take it to the truck.
Put the pallet on the truck.

Załaduj towar na ciężarówkę.

Załaduj towar na ciężarówkę.

Put the pallet on the truck

Załaduj towar na ciężarówkę.

Załaduj towar na ciężarówkę.

Related Work

- Symbolic approaches.
  - Winograd (1970), Hsiao, Mavridis and Roy (2003), MacMahon et al. (2006), Siskind (2001), Allen et al. (2007)
  - Carefully model the structure of language.
  - Challenge: robustly understanding novel utterances from untrained users.

- Probabilistic Grounding Models.
  - Mavridis and Roy (2006). Shimizu and Haas (2009), Matuszek, Fox, and Koscher (2010). Kollar, Tellex, Roy, and Roy (2010).
  - Probabilistic models that map between language and aspects of the external world.
  - Challenge: modeling the hierarchical, compositional structure of language.

- Learning Language through Action
  - Branavan et al. (2010), Liang et al. (2010), Vogel and Jurafsky (2010).
  - Learn from language by acting in the world and observing an external reward signal.
  - Challenge: learning mappings between words and perceptual features.

- Generalized Grounding Graphs (this talk)
  - Robustly understand natural language commands from untrained users.
  - Model the compositional and hierarchical structure of language.
  - Learning word meanings in terms of perceptual features.

Related Work

- Generalized Grounding Graphs (Tellex et al. 2011)
- Reinforcement Learning for Mapping Instructions to Actions (Branavan et al. 2010)

“Go to the truck”

What should the robot do?

“Go to the truck”

What should the robot do?

“Go to the truck”

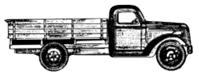

Approach:

1. Define an objective function.
2. Search for the action that maximizes the function.
3. Execute that action.

$$\underset{\gamma_1}{argmax}\ f\left(\text{Go to the truck}, \gamma_1\right)$$

“Go to the truck”

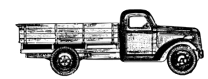

Approach:

1. Define an objective function.
2. Search for the action that maximizes the function.
3. Execute that action.

$$\underset{\gamma_1}{argmax} \; p\left(\gamma_1 | \text{Go to the truck}\right)$$

More Complex Language

$$\underset{\gamma_1}{argmax}\ p(\gamma_1|\text{Go to the truck})$$

$$\underset{\gamma_1}{argmax}\ p(\gamma_1|\text{Go to the pallet})\ ?$$

More Complex Language

$$\underset{\gamma_1, \gamma_2}{argmax} \; p(\gamma_1, \gamma_2 | \text{Go to the truck})$$

$$\underset{\gamma_1, \gamma_2}{argmax} \; p(\gamma_1, \gamma_2 | \text{Go to the pallet})$$

another tyre pallet on the trolley.

Place the pallet of tires on the left side of the trailer.

lift the tire pallet to the truck

Arrange tire pallet to the truck.

Please lift the set of six tires up and set them on the trailer, to the right of the set of tires already on it.

Load the skid right next to the other skid of tires on the trailer.

Place a second pallet of tires on the trailer.

Place the pallet of tires that is on the forklift next to the pallet of tires that is already loaded on the trailer.

Lift tire pallet. Move to unoccupied location on truck. Lower tire pallet. Reverse to starting location. Lower forks. End.

Put the tire pallet on the trailer.

Lift the tire pallet in the air, then proceed to deposit it to the right of the
tire pallet already on the table right in front of you.

Place the pallet of tires on the right side of the truck.

Lift the tire pallet and proceed forward to set it on the platform directly ahead, to the right of the tire pallet already there.

More Complex Language

$$\underset{\gamma_{1},\gamma_2}{argmax}\ p(\gamma_{1,}\gamma_2|\text{Go to the truck})$$

$$\underset{\gamma_{1},\gamma_2}{argmax}\ p(\gamma_{1,}\gamma_2|\text{Go to the pallet})$$

$$\underset{\gamma_1 \ldots \gamma_N}{argmax}\ p(\gamma_1 \ldots \gamma_N|language)$$

γ_k are *groundings*, or objects, places, paths, and events in the external world. Each γ_k corresponds to a constituent phrase in the language input.

“the truck”

$$p(\phi|\gamma,\lambda^f)$$

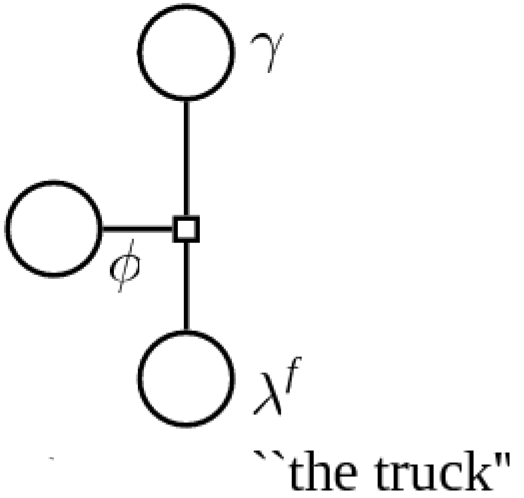

“the truck”

$$p(\phi|\gamma,\lambda^{f})=0.9$$

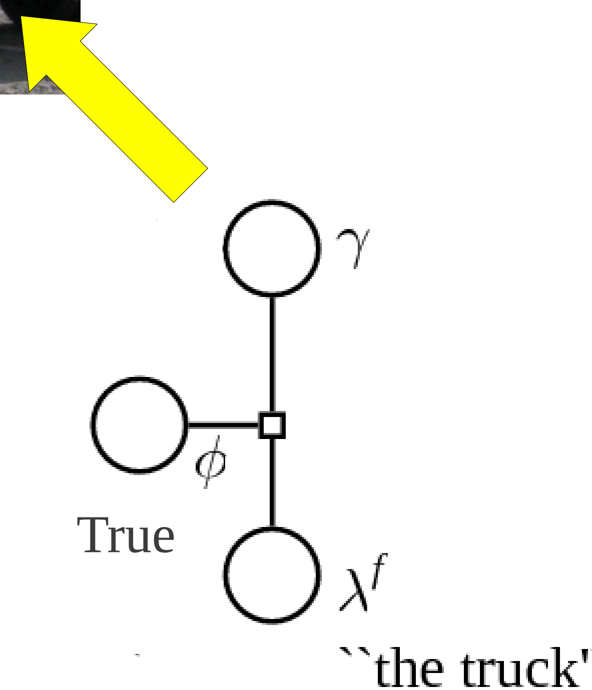

“the truck”

$$p(\phi|\gamma,\lambda^f)=0.1$$

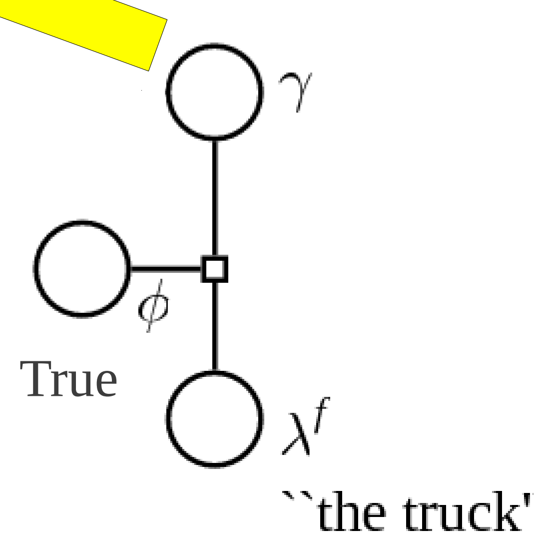

“on the truck”

“on the truck”

$$p(\phi_5|\gamma_{5,}\lambda_5^f)$$

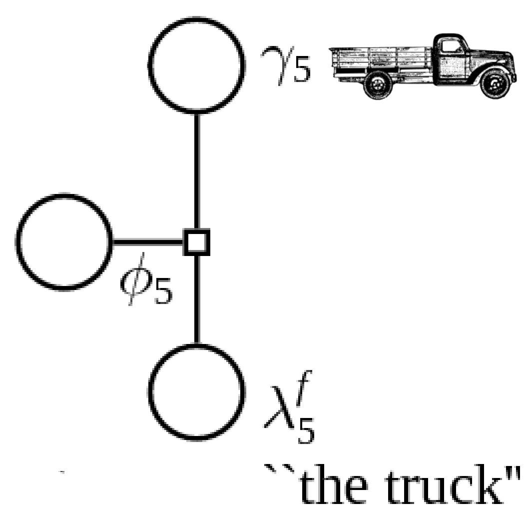

“on the truck”

$$p(\phi_3|\gamma_4,\gamma_5,\lambda_3^r)\times p(\phi_5|\gamma_5,\lambda_5^f)$$

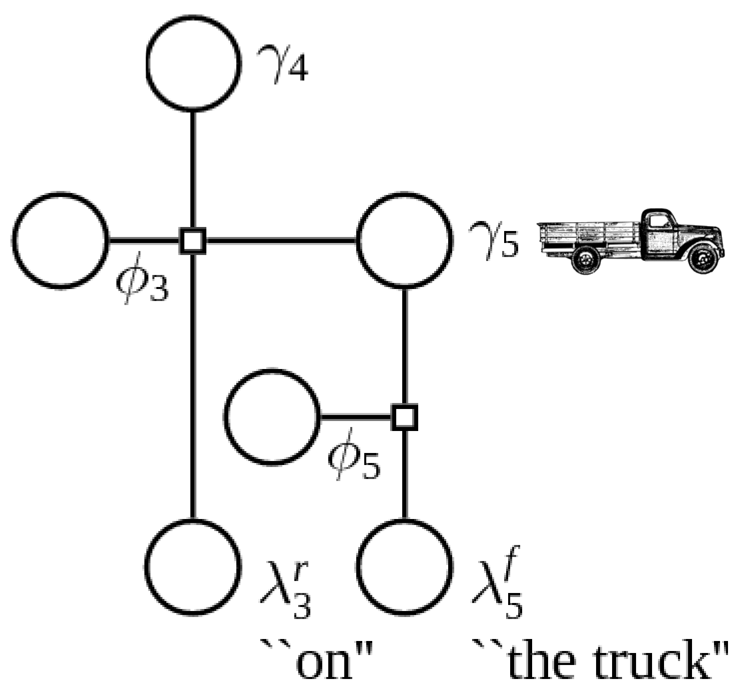

“on the truck”

$$p(\phi_3|\gamma_4, \gamma_5, \lambda_3^r) \times p(\phi_5|\gamma_5, \lambda_5^f)$$

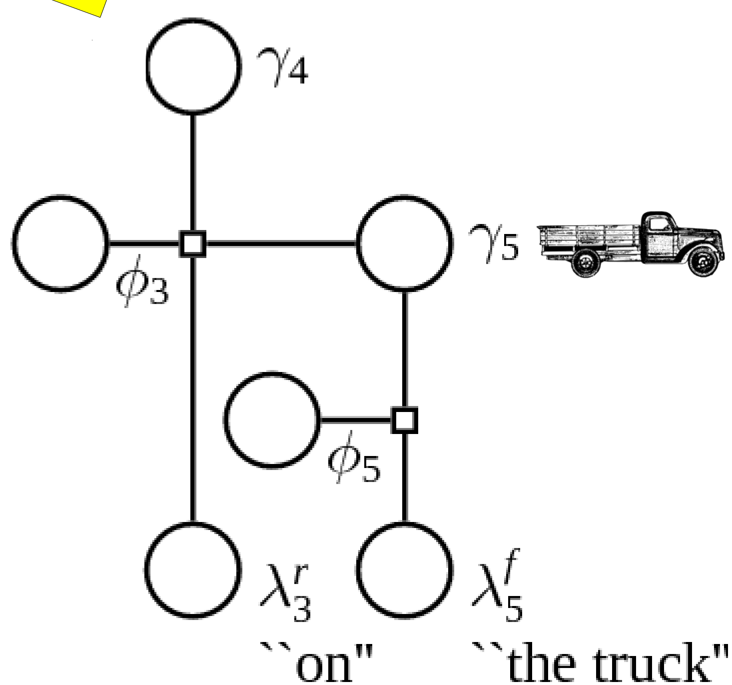

“on the truck”

$$p(\phi_3|\gamma_4,\gamma_5,\lambda_3^r)\times p(\phi_5|\gamma_5,\lambda_5^f)$$

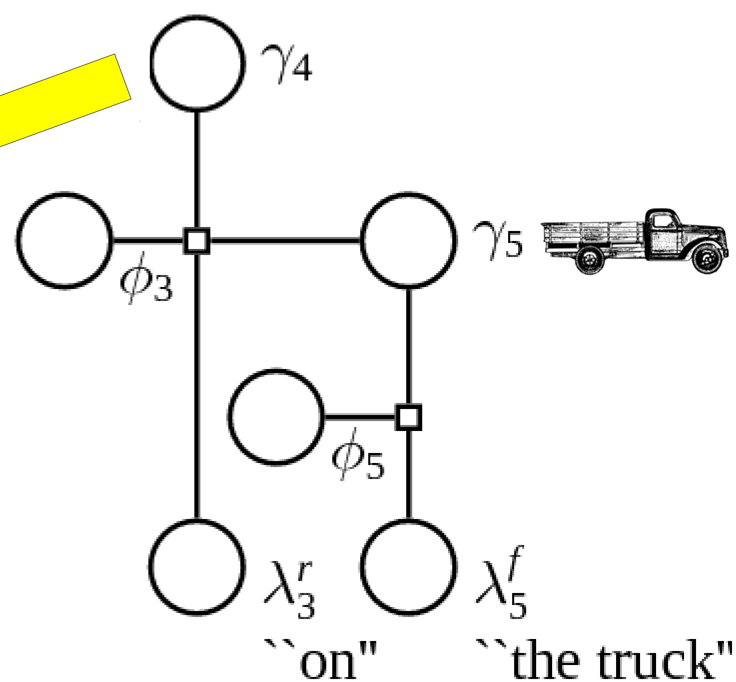

“Go to the pallet on the truck.”

$$p(\phi_1|\gamma_1,\gamma_2,\lambda_1^r)\times p(\phi_2|\gamma_2,\gamma_4,\lambda_2^r)\times p(\phi_4|\gamma_4,\lambda_4^f)\times p(\phi_3|\gamma_4,\gamma_5,\lambda_3^r)\times p(\phi_5|\gamma_5,\lambda_5^f)$$

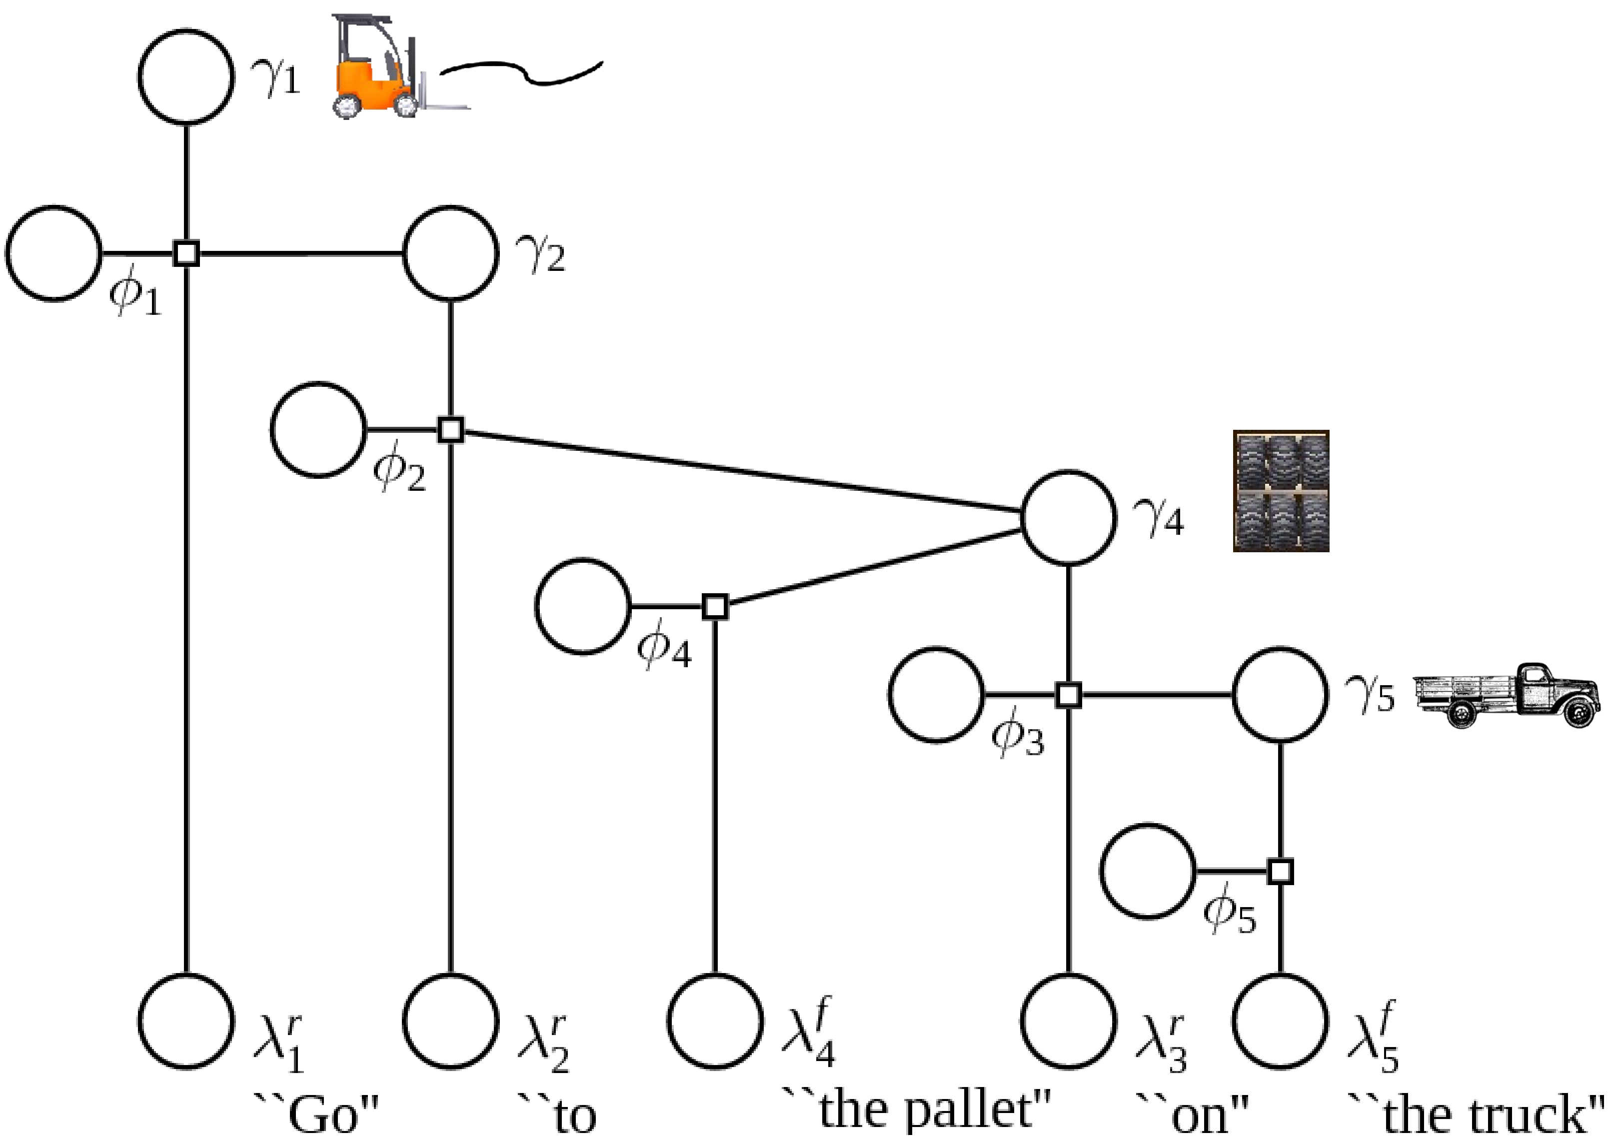

“Put the pallet on the truck.”

$$p(\phi_1|\gamma_{1,}\lambda_1^r)\times p(\phi_2|\gamma_{2,}\lambda_2^f)\times p(\phi_3|\gamma_{3,}\gamma_{4,}\lambda_3^r)\times p(\phi_4|\gamma_{4,}\lambda_4^f)$$

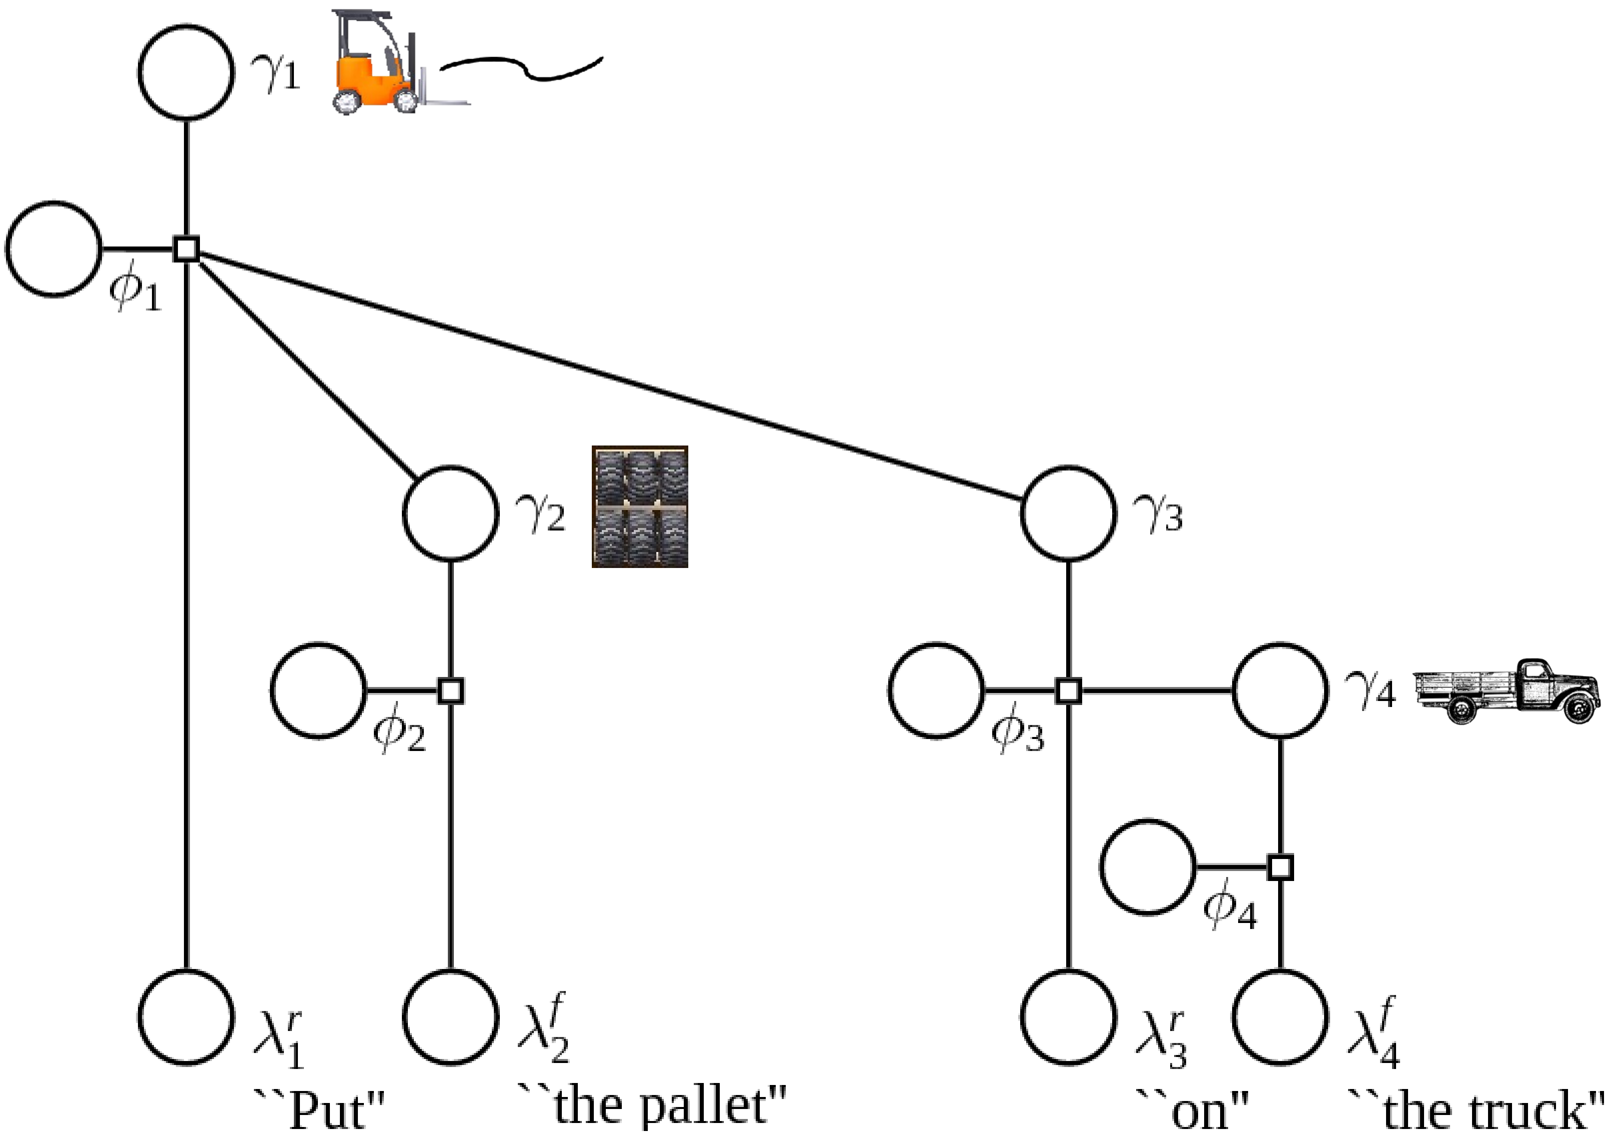

“Go to the pallet on the truck.”

“Put the pallet on the truck.”

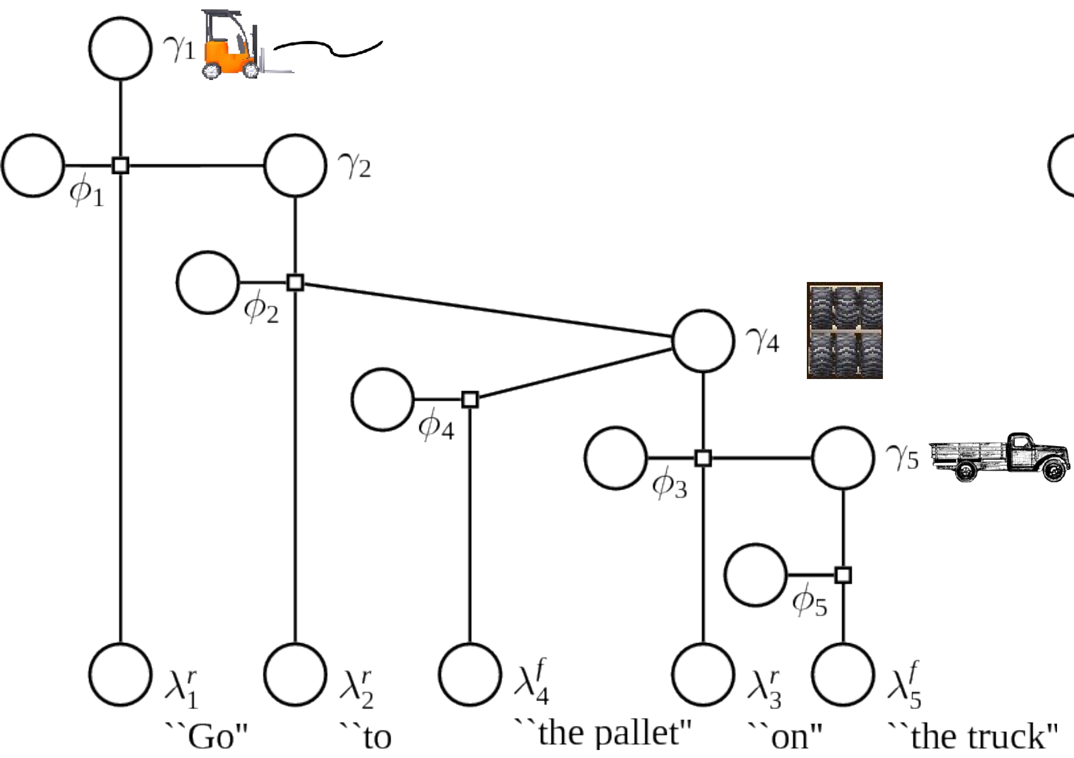

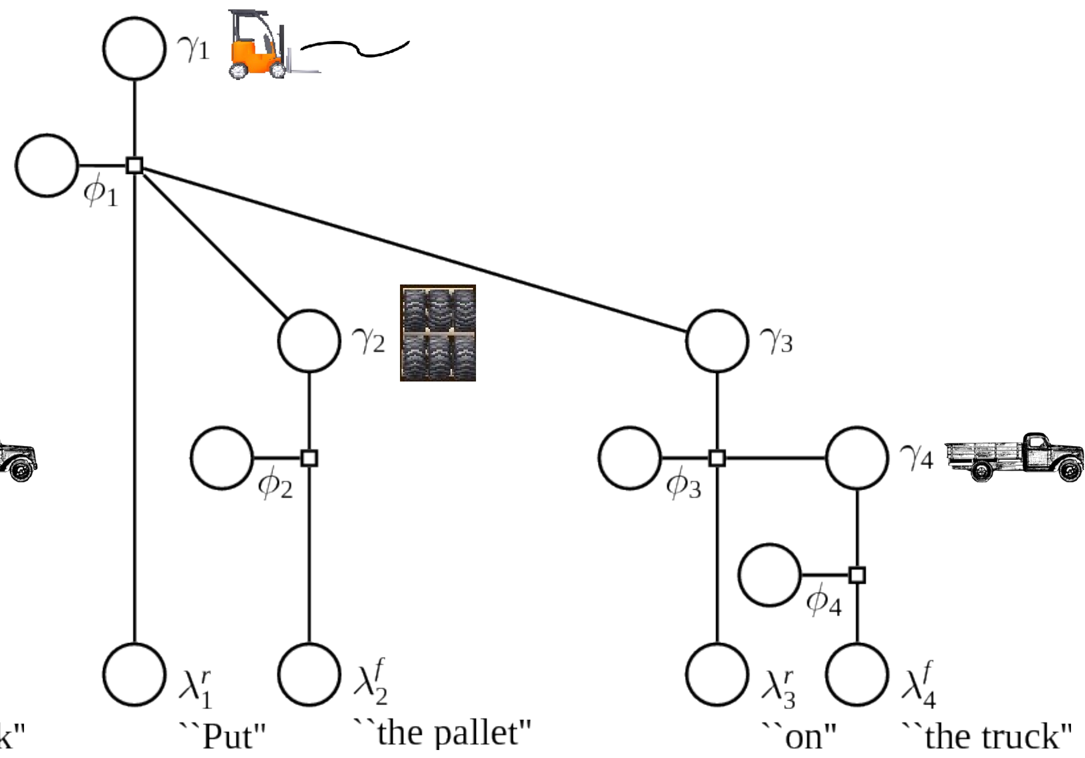

$$p(\phi_1|\gamma_1,\gamma_2,\lambda_1^r)\times p(\phi_2|\gamma_2,\gamma_4,\lambda_2^r)\times p(\phi_4|\gamma_4,\lambda_4^f)\times p(\phi_3|\gamma_4,\gamma_5,\lambda_3^r)\times p(\phi_5|\gamma_5,\lambda_5^f)$$

$$p(\phi_1|\gamma_1,\lambda_1^r)\times p(\phi_2|\gamma_2,\lambda_2^f)\times p(\phi_3|\gamma_3,\gamma_4,\lambda_3^r)\times p(\phi_4|\gamma_4,\lambda_4^f)$$

Generalized Grounding Graphs

$$p(\Phi|\gamma_1 \ldots \gamma_N, language) = \prod_i p(\phi_i|\lambda_i, \gamma_1 \ldots \gamma_N)$$

$$= \prod_i \frac{1}{Z} \exp\left(\sum \mu_k s_k(\phi_i, \lambda_i, \gamma_1 \ldots \gamma_N)\right)$$

- Factor for each linguistic constituent.
- Each factor is a log-linear model.
- Training requires fully aligned data, with positive and negative examples.

μ_k - Feature weight.
s_k - Binary-valued feature function.

Grounding Decision Problem

- Action: Assign a value to all variables in the grounding graph.

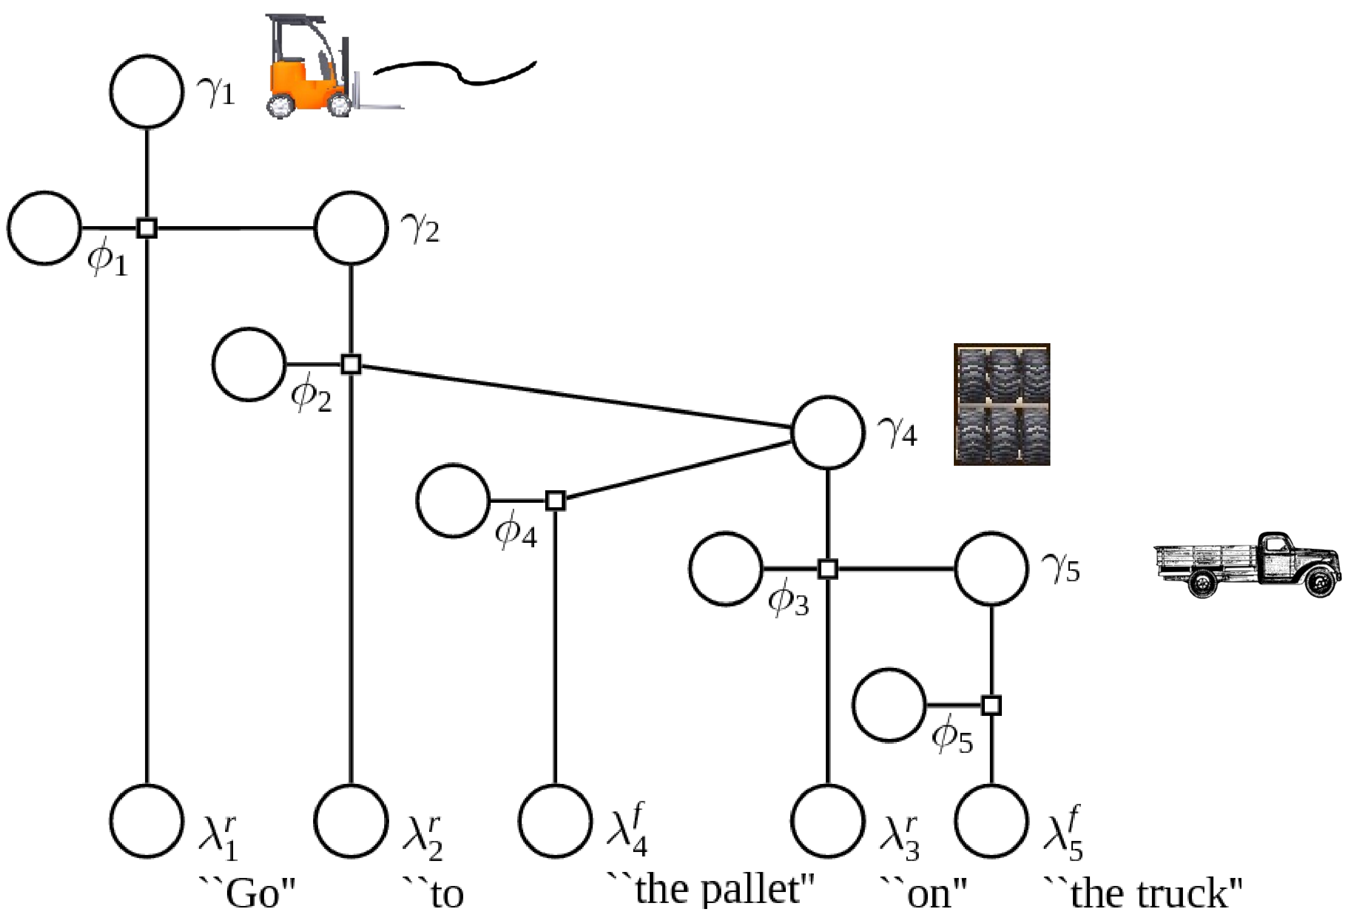

Grounding Decision Problem

- Action: Assign a value to all variables in the grounding graph.

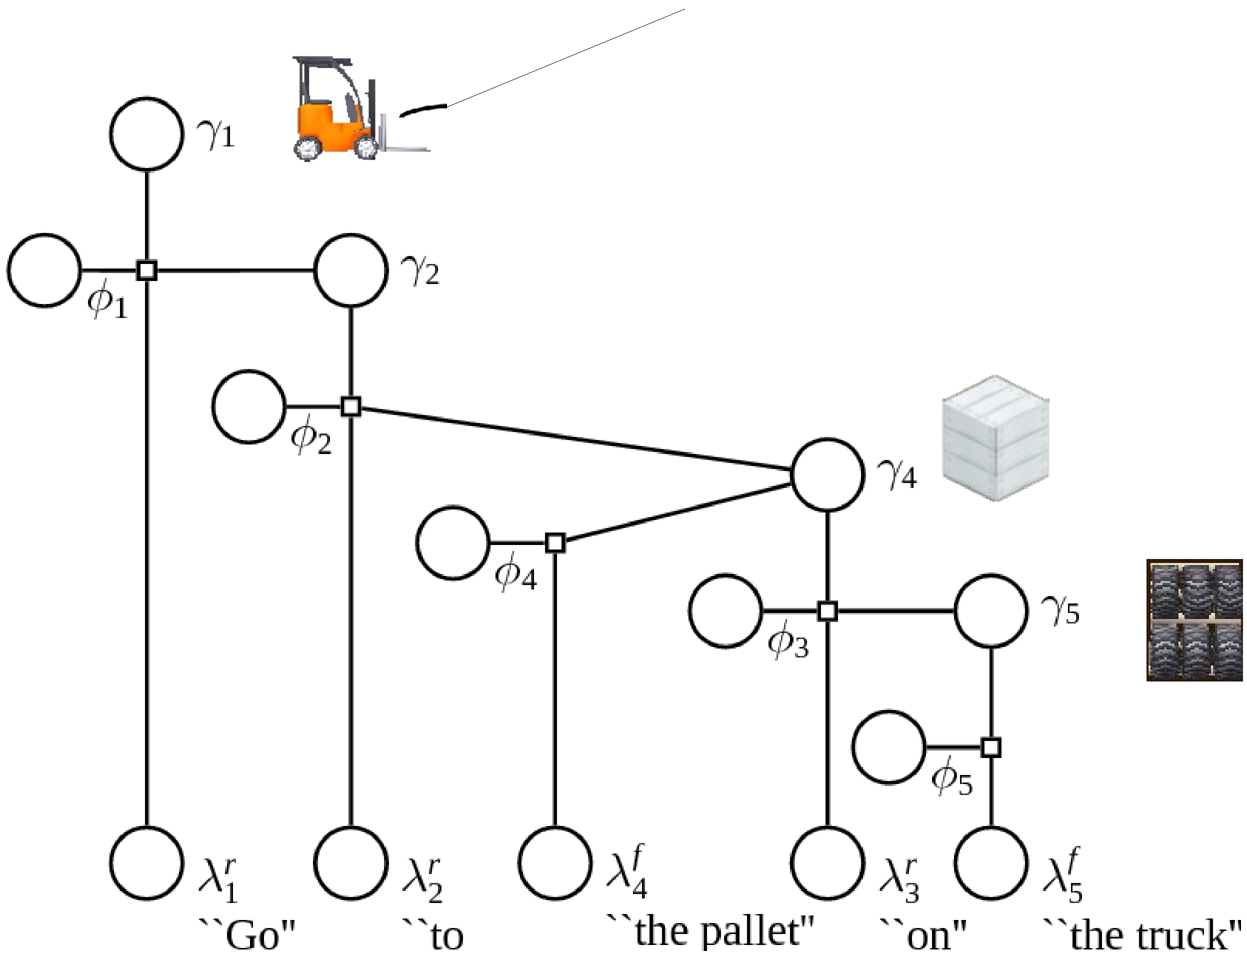

Physical State/Action Space

- State:
  - Location of the robot.
  - Location of all objects.
- Primitive actions:
  - Pick up an object.
  - Put down an object.
  - Drive to locations in a topological map.
- Each state defines a space of values for each grounding variable.

Grounding Decision Problem

Action 1: γ_{go}=[Drive to location 1], γ_{pallet}=[pallet 1]

Action 2: γ_{go}=[Drive to location 1], γ_{pallet}=[pallet 2]

Action 3: γ_{go}=[Drive to location 2], γ_{pallet}=[pallet 1]

Action 4: γ_{go}=[Drive to location 2], γ_{pallet}=[pallet 2]

Action 5: γ_{go}=[Drive to location 1], γ_{pallet}=[pallet 1]

Action 6: γ_{go}=[Drive to location 1], γ_{pallet}=[pallet 2]

Action 7: γ_{go}=[Drive to location 2], γ_{pallet}=[pallet 1]

Action 8: γ_{go}=[Drive to location 2], γ_{pallet}=[pallet 2]

...

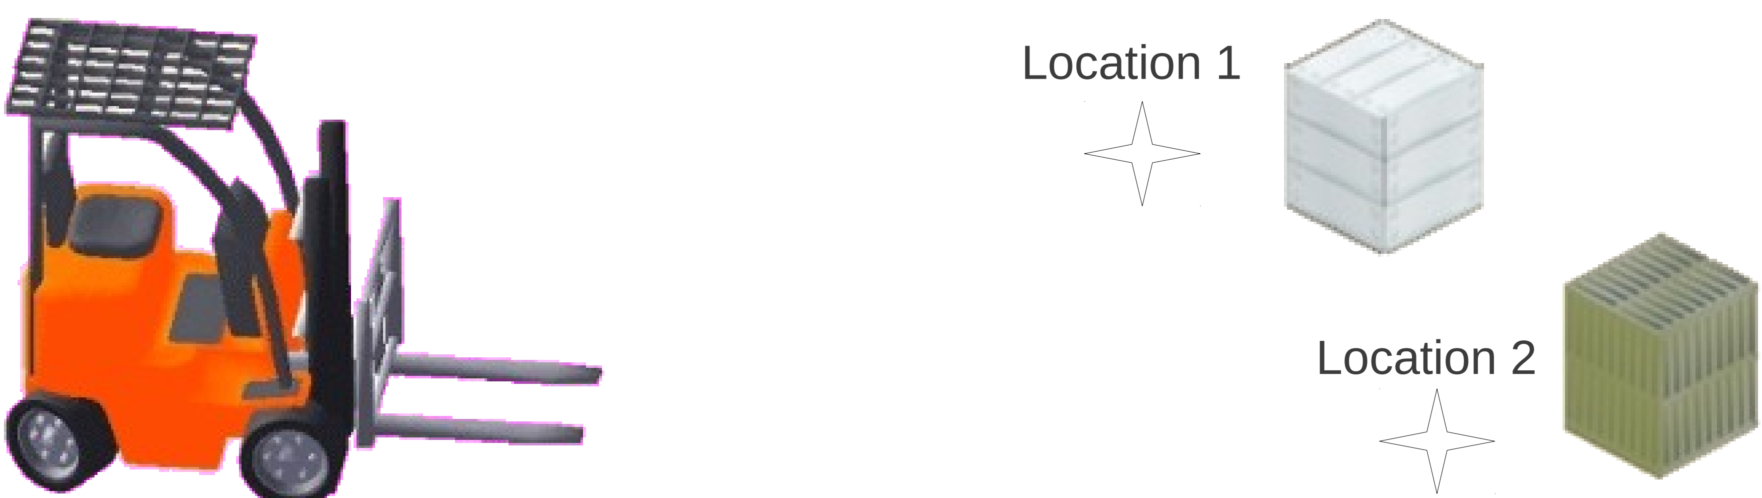

Reward

- Base reward:
  - +1 for choosing the right top-level action (regardless of other groundings).
  - -1 otherwise.
- Base reward is scaled so the sum of positive reward equals the sum of negative reward.

Policy Gradient

- Assume a parametric form for the policy that factors based on G[3].
- Optimize model parameters via gradient ascent.
- Top-level reward signal is distributed among factors for each linguistic constituent.

$$argmax_{\Theta} E_{p(\Gamma|\Phi,\Lambda,M,\Theta)} r(\Gamma, g_a)$$

Evaluation

- Corpus of “pick up” and “to” examples issued in simple environments.

Results

Training

	% Correct, Actions	% Correct, Concrete Noun Phrases
Random	11%	45%
Our Approach	80%	71%

Testing

	% Correct, Actions	% Correct, Concrete Noun Phrases
Random	17%	32%
Our Approach	67%	73%
Fully Supervised	73%	82%

Heat Map

$$p\left(\gamma_1 | \text{Go to the truck}\right)$$

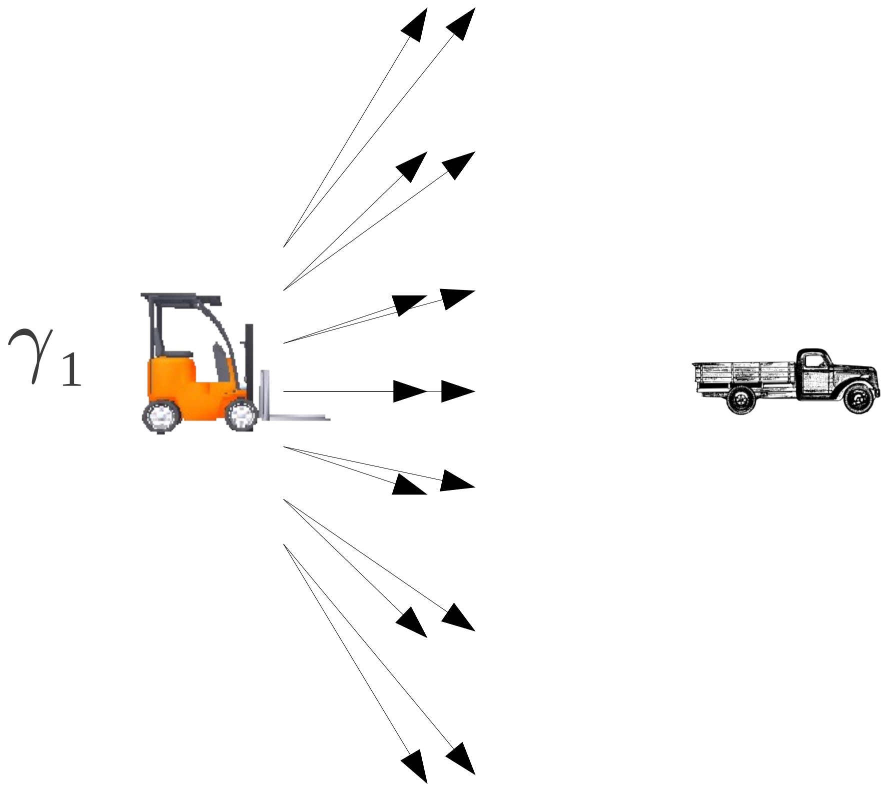

$$p\left(\gamma_1 | \text{Go to the truck}\right)$$

Trained from aligned data

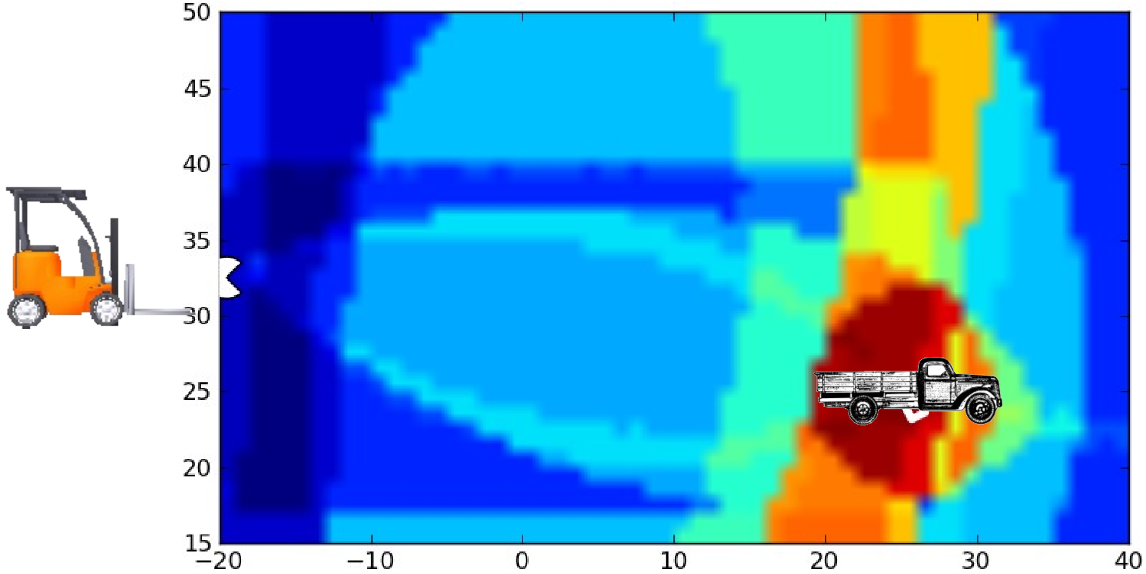

$$p\left(\gamma_1 | \text{Go to the truck}\right)$$

Trained from aligned data

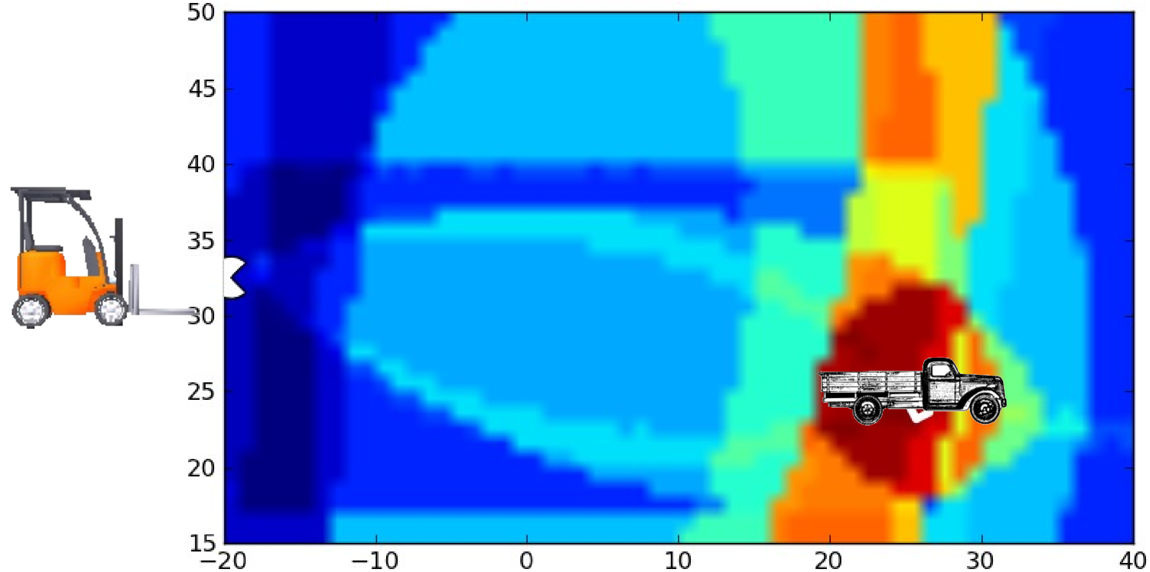

Trained from unaligned data

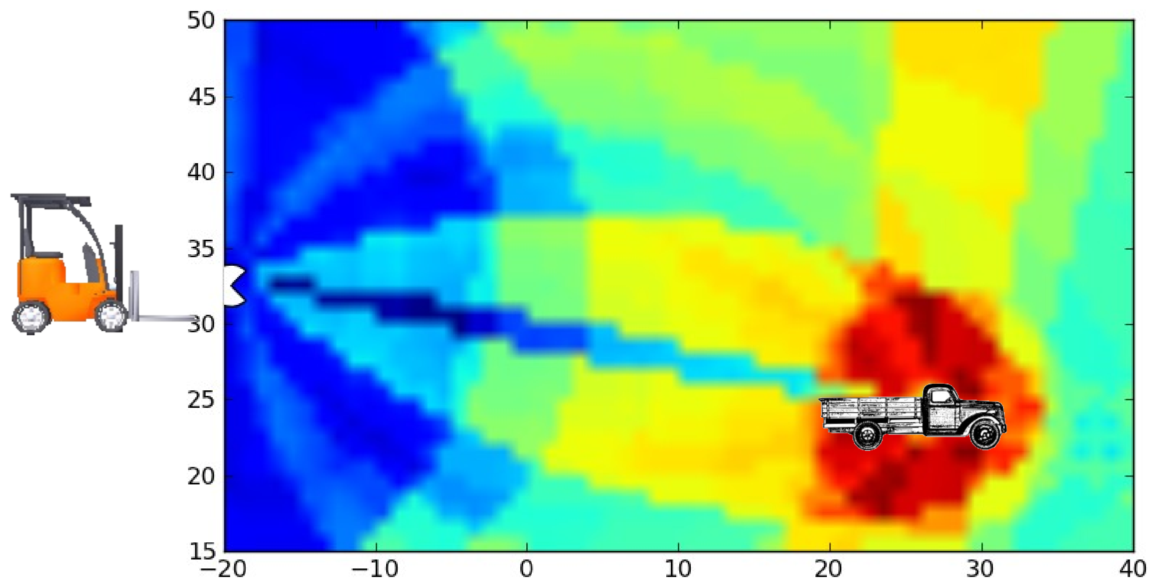

Future Work

- Scaling to very large corpora.
- Generalizing learned word meanings.
  - Multiple platforms. (PR2)
- Information-theoretic human-robot dialog.
- Speech.

Conclusion

- Policy gradient algorithm for learning grounded word meanings from unaligned parallel corpora.
- Parametric form for the policy that factors according to the structure of language.
- Demonstration of learning grounded word meanings from an unaligned parallel corpus of “pick up” and “to” commands.