Vocabulary Alignment for Collaborative Agents: a Study with Real-World Multilingual How-to Instructions

IIIA Institut d'Investigació en Intel·ligència Artificial

Paula Chocron, Paolo Pareti

University of Southampton

Vocabulary Alignment

Setting: collaboration in open multi-agent systems.

Problem: agents use different languages or dialects. Linguistic resources might not be available, might be expensive to use or not properly contextualised.

Approach: interaction-based alignment techniques.

Collaborative Interaction

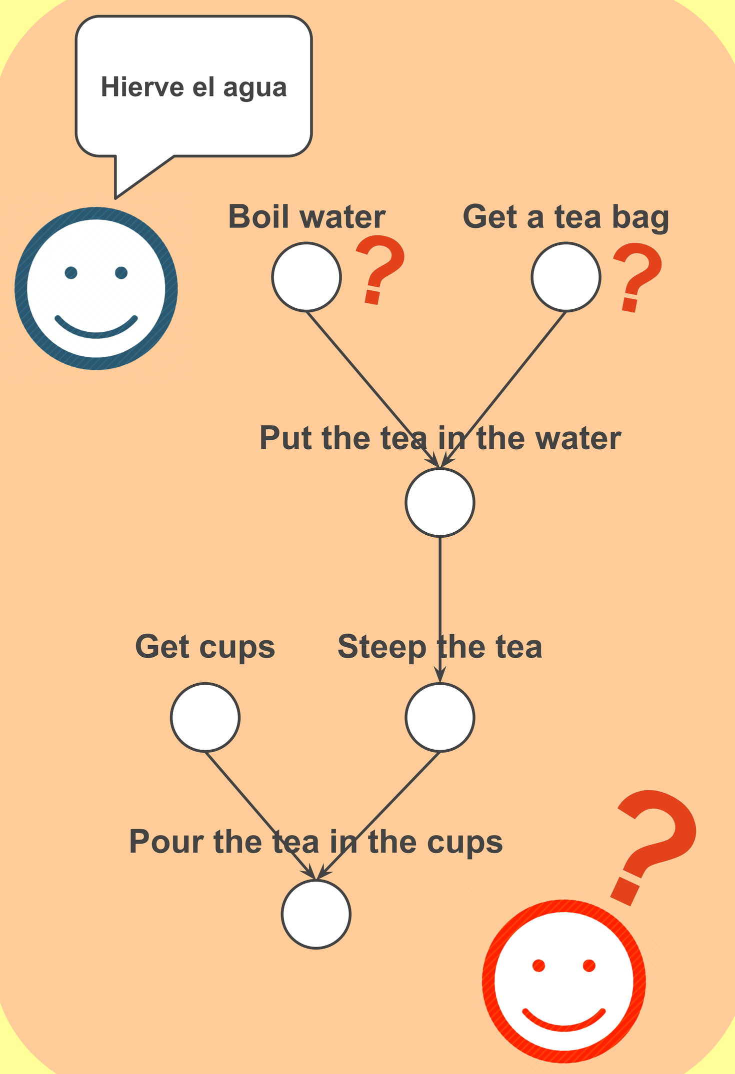

Agents follow the same protocols, performing actions in turns.

They communicate which action they took with a message.

They only know how to perform a subset of the required actions, but their capabilities are complementary.

Interpretation and Learning

They interpret messages using a confidence distribution over the mappings of the words in the foreign vocabulary.

It uses dynamically learned word frequencies.

$$\alpha^a(t^b) \in \underset{t^a \in Poss^a}{\operatorname{argmax}} (\delta(t^b, t^a))$$

$$\delta(t^b, t^a) = \max_{c \in Perm_m(t^b)} \delta_p(c, t^a) - \rho * abs(n - m)$$

$$\delta_p(c, t^a) = \frac{\sum_{0 \leq i \leq m} \omega(c[i], t^a[i])}{Freq(c[i]) + Freq(t^a[i])}$$

After each interaction, an agent updates its confidence distribution over word mappings.

It increases the confidence of the mappings between the foreign words of the received message, and the local words in the label of the next expected actions.

for all $t^a \in Poss^a$ and $v^a \in t^a$:

$$\omega(v^a, v^b) = \omega(v^a, v^b) + \max_{c \in Perm_{v^b, v^a}} \delta_p(c, t^a)$$

Data Acquisition

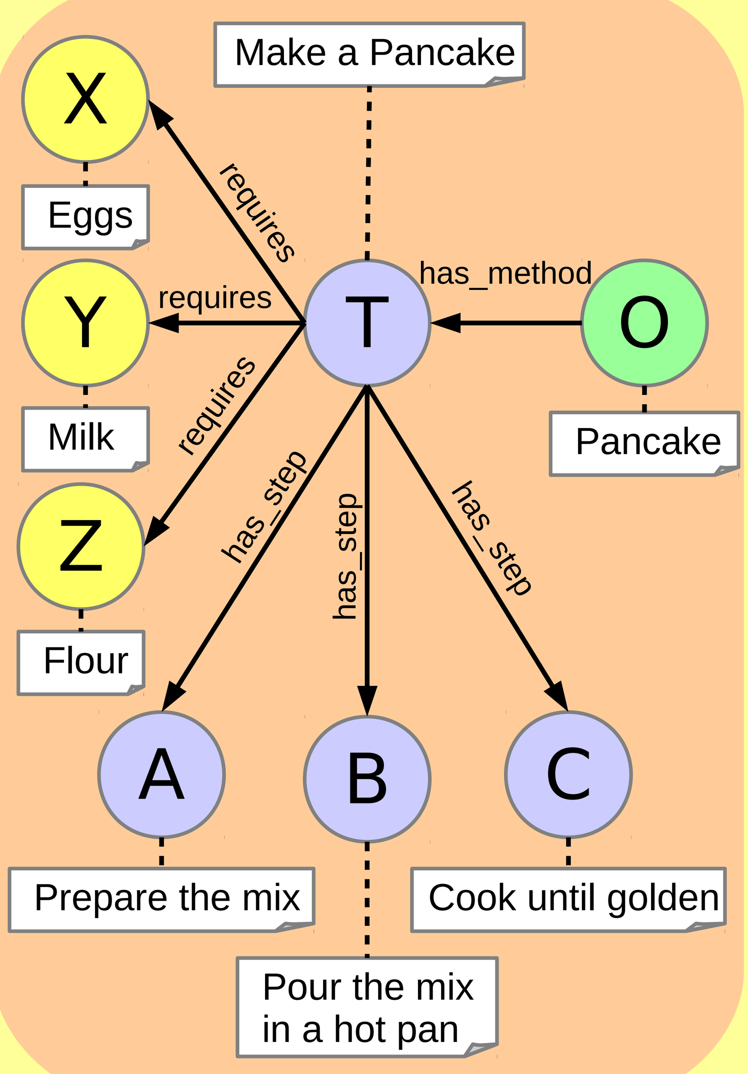

Real word data is taken from **The Human Know-How Dataset**.

http:// w3id.org/knowhow/dataset

- it contains over 200K interlinked instructions
- versions in 16 languages
- extracted from the following websites:

wikiHow The world's how to manual.

Snapguide

Experimental Evaluation

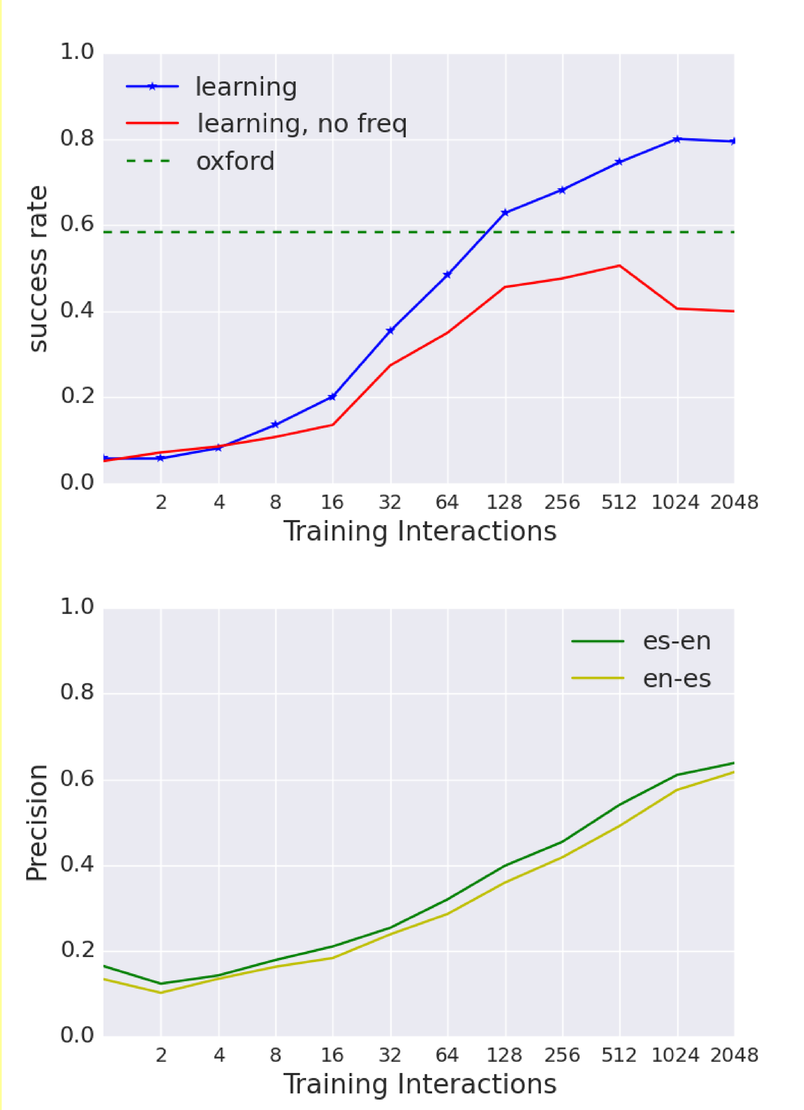

We converted cooking-related instructions into protocols, obtaining 327 structurally compatible bilingual protocols (English and Spanish)

Agents update their alignments using randomly chosen protocols for *n* iterations before their success rate is measured.

Findings

- Our simple approach reaches a higher success rate than the baseline after ~100 interactions.
- Considering word frequency significantly increases the success rate.
- The correctness of the word-alignments improves with training, although agents do not have a notion of semantically correct alignments.
- Our approach discovers contextually-correct translations between words that are not listed in the dictionary.
- Adapting existing approaches to real-world data requires rethinking core assumptions.