Interactive installation for collaborative creation of a language:

Active negotiation of new linguistic conventions

William Schueller, Pierre-Yves Oudeyer

Kreyon Conference,
Rome 08/09/2017

Social Convention?

- Arbitrary choice (among equivalent options)
- Interest in agreeing with the group

Road side for driving

Greeting protocol

Electric plugs

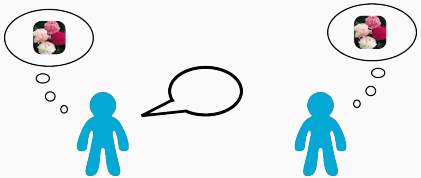

Language

Social Convention?

- Arbitrary choice (among equivalent options)
- Interest in agreeing with the group

- Potential conflicts
- **Dynamics?**

Road side for driving

Greeting protocol

Electric plugs

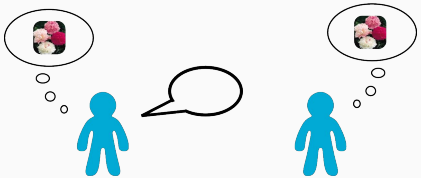

Language

Social Convention?

- Arbitrary choice (among equivalent options)
- Interest in agreeing with the group

- Potential conflicts
- **Dynamics?**

Road side for driving

Greeting protocol

Electric plugs

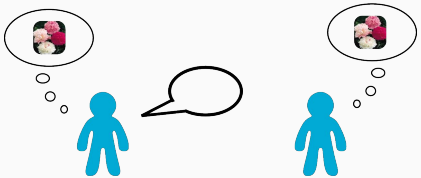

Language

Social Convention?

- Arbitrary choice (among equivalent options)
- Interest in agreeing with the group

- Potential conflicts
- **Dynamics?**

1 limite meridionale del tipo *ortiga* «ortica» (*-c-* > *-g-*)
2 \\ \\ \\ \\ *sal* «sale» (caduta di *-e*)
3 \\ \\ \\ \\ *cavei* «capelli» (*-p-* > *-v-*)
4 \\ \\ \\ \\ *spala* «spalla» (*-ll-* > *-l-*)
5 \\ \\ \\ \\ *slèr, slar* «sellaio» (caduta della vocale protoni)
6 \\ \\ \\ \\ *pâ* «pane»
7 \\ \\ \\ \\ *incô, inci* «oggi»
8 limite settentrionale del tipo *ferraru* «fabbro»
9 \\ \\ \\ \\ *frate* «fratello»
10 \\ \\ \\ \\ *femmina* «donna»
11 \\ \\ \\ \\ *figliomo* «mio figlio»
12 \\ \\ \\ \\ *tene le spalle larghe* «ha le spalle larghe»
13 \\ \\ \\ \\ *cosa* «coscia»
14 \\ \\ \\ \\ *lu cimice* «la cimice»
15 \\ \\ \\ \\ *fagu* «faggio»
16 \\ \\ \\ \\ *mondone* «montone» (*-nt-* > *-nd-*)
17 \\ \\ \\ \\ *dienti* «denti» (metafonesi)
18 \\ \\ \\ \\ *acitu* «aceto» (metafonesi)

I numeri sono quelli dei punti d'inchiesta dell'*AIS*.

Competition between words for same meaning: femmina/donna, pa/pane, acitu/aceto, …

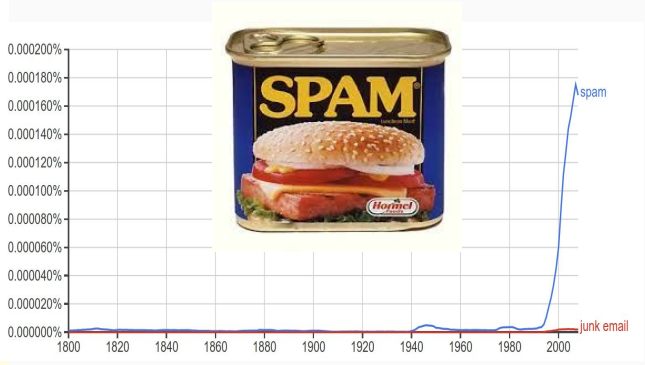

New word-meaning associations

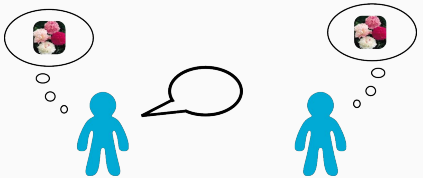

Language

The Naming Game

- Multi-agent model
- No central control or information (2 individuals per interaction)

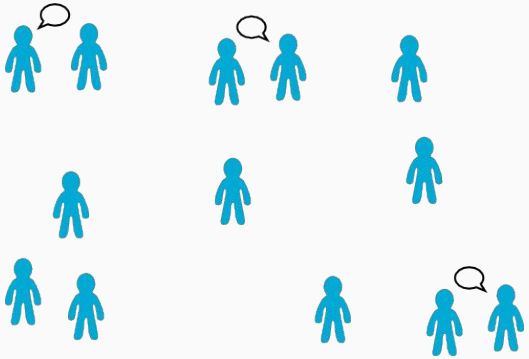

Origins: the Talking Heads Experiment (Luc Steels)

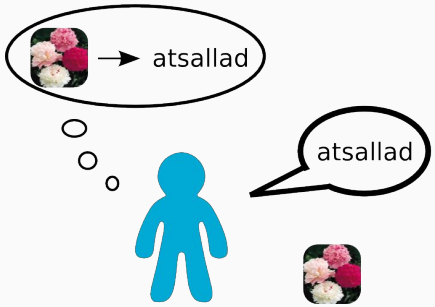

Speaker

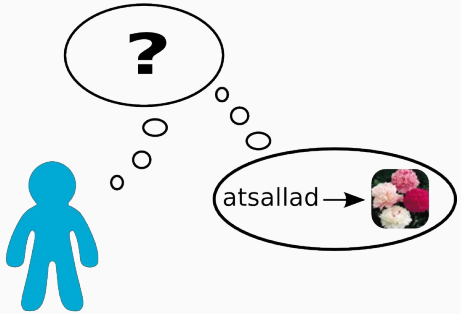

Hearer

The Naming Game

- Multi-agent model
- No central control or information (2 individuals per interaction)

Innovation: inventing new meaning-word associations

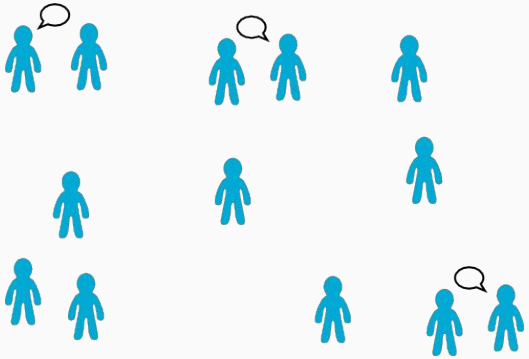

Origins: the Talking Heads Experiment (Luc Steels)

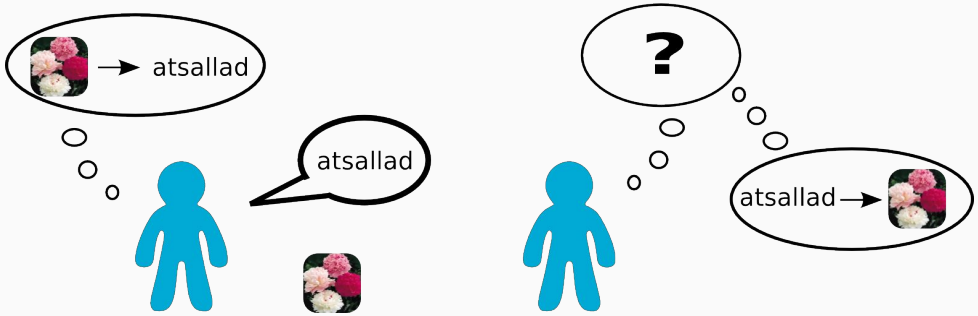

Speaker

Hearer

The Naming Game

- Multi-agent model
- No central control or information (2 individuals per interaction)
- Simple rules of update after interaction
- Dealing with synonymy

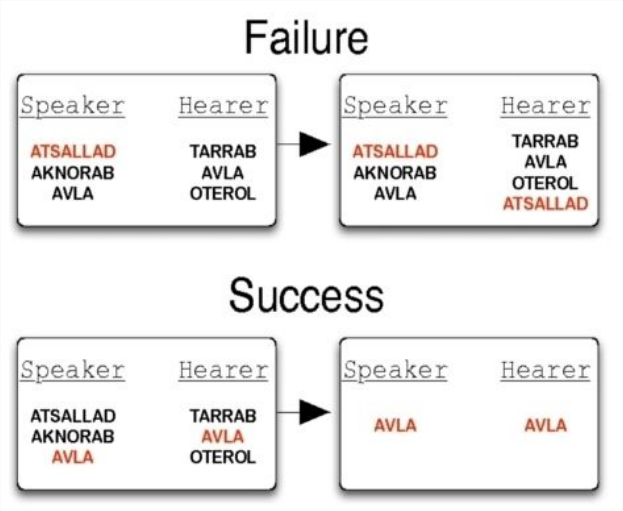

Andrea Baronchelli and Vittorio Loreto

The Naming Game

- Multi-agent model
- No central control or information (2 individuals per interaction)
- Simple rules of update after interaction
- Dealing with synonymy
- **Explosion of complexity !**
  Before reaching consensus

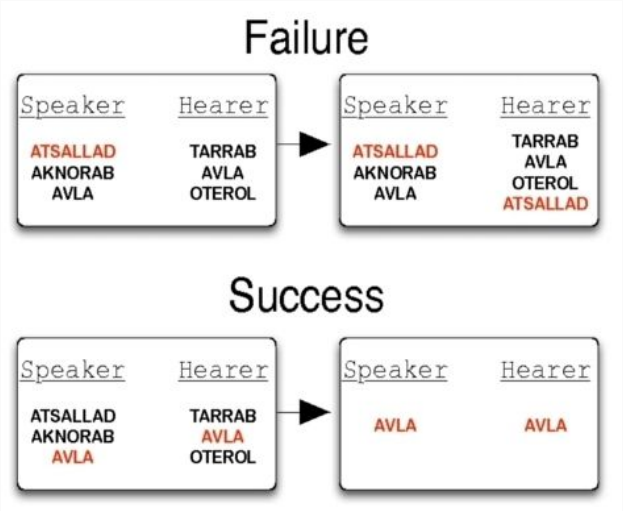

Andrea Baronchelli and Vittorio Loreto

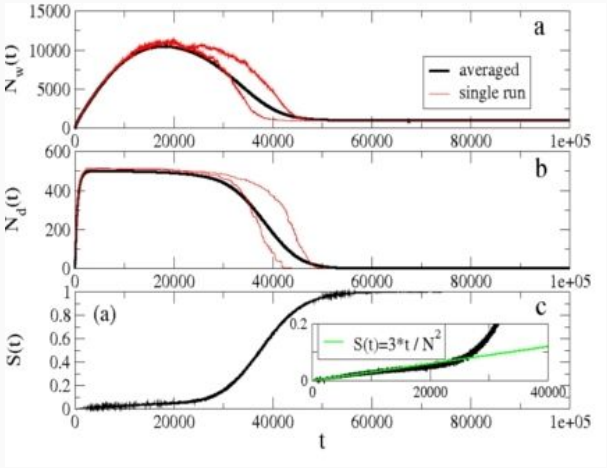

Average memory per agent
Here, at peak, ~10 words per meaning

Number of different words for same meaning

Reaching consensus,
% of successful interactions

Control of complexity?

Young children are really efficient learners, how do they do? How do they influence their learning process?

Control of complexity?

- **Flow theory:** choose challenges neither too complex, neither too simple.
- Intrinsic motivation, curiosity: controlling complexity growth
- Developmental paths (learn to move arms before walking, helps to keep balance; babble before talking, etc)

Young children are really efficient learners, how do they do? How do they influence their learning process?

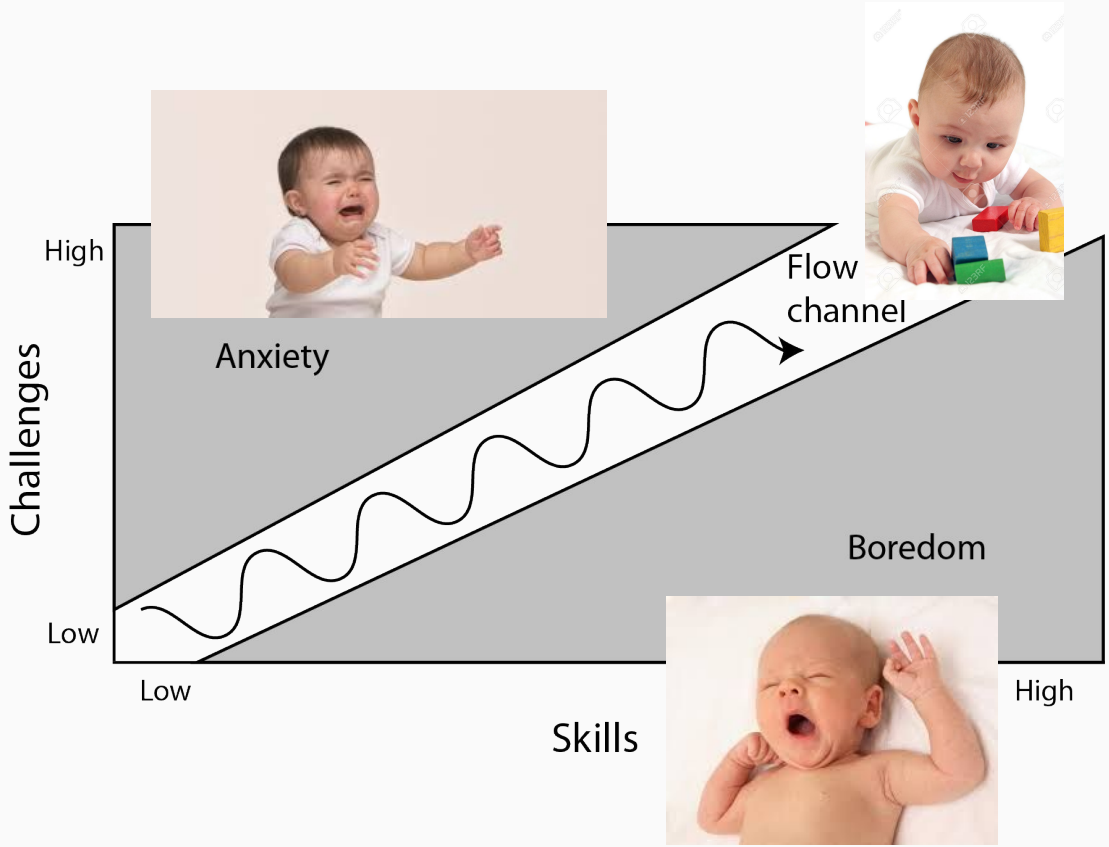

Control of complexity?

- **Flow theory:** choose challenges neither too complex, neither too simple.
- Intrinsic motivation, curiosity: controlling complexity growth
- Developmental paths (learn to move arms before walking, helps to keep balance)
- **Learning progress maximization**
- Motor skills learning, language learning (babbling, syllables, simple words, sentences), …
- What about “common social learning”? (Agreeing on new conventions)

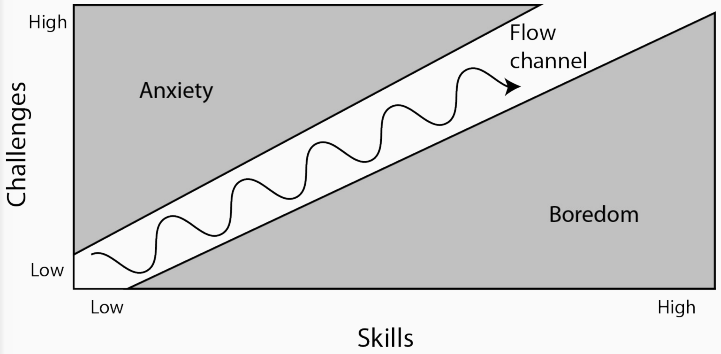

Pierre-Yves Oudeyer

Robots learning basic motor and language skills : like children?

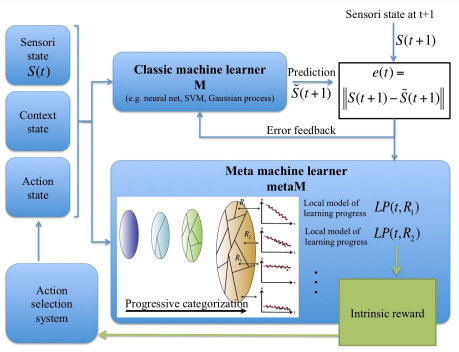

Interest models

Control of complexity?

What about “common social learning”? (Agreeing on new conventions)

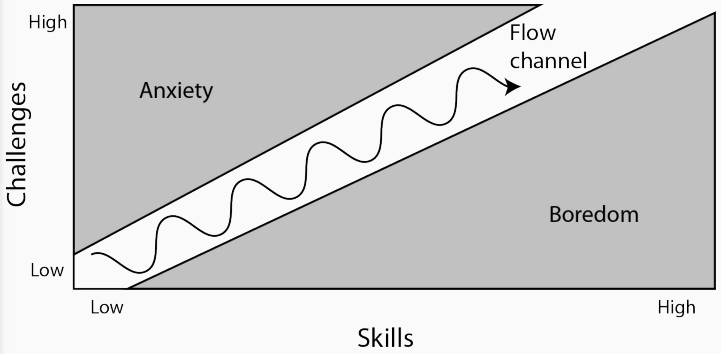

Pierre-Yves Oudeyer

Robots learning basic motor and language skills : like children?

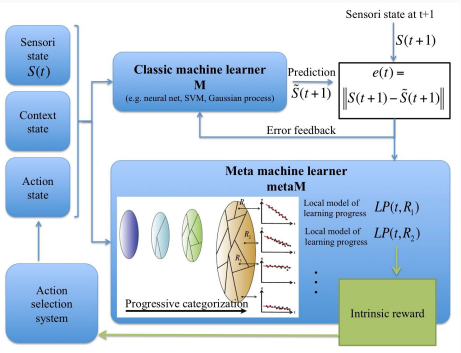

Interest models

Active control of complexity in the Naming Game?

- Several choices!

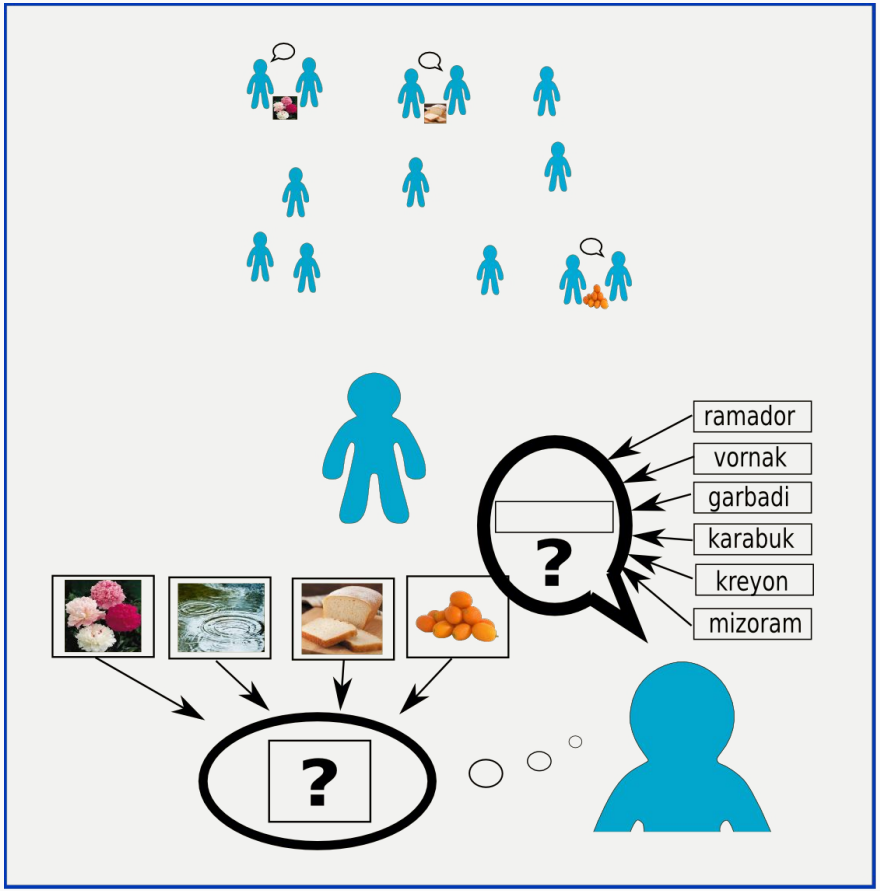

Active control of complexity in the Naming Game?

- Several choices!
  - Topic choice (before interaction)
  - Acceptance policy (after interaction): trust or not the other
  - ...

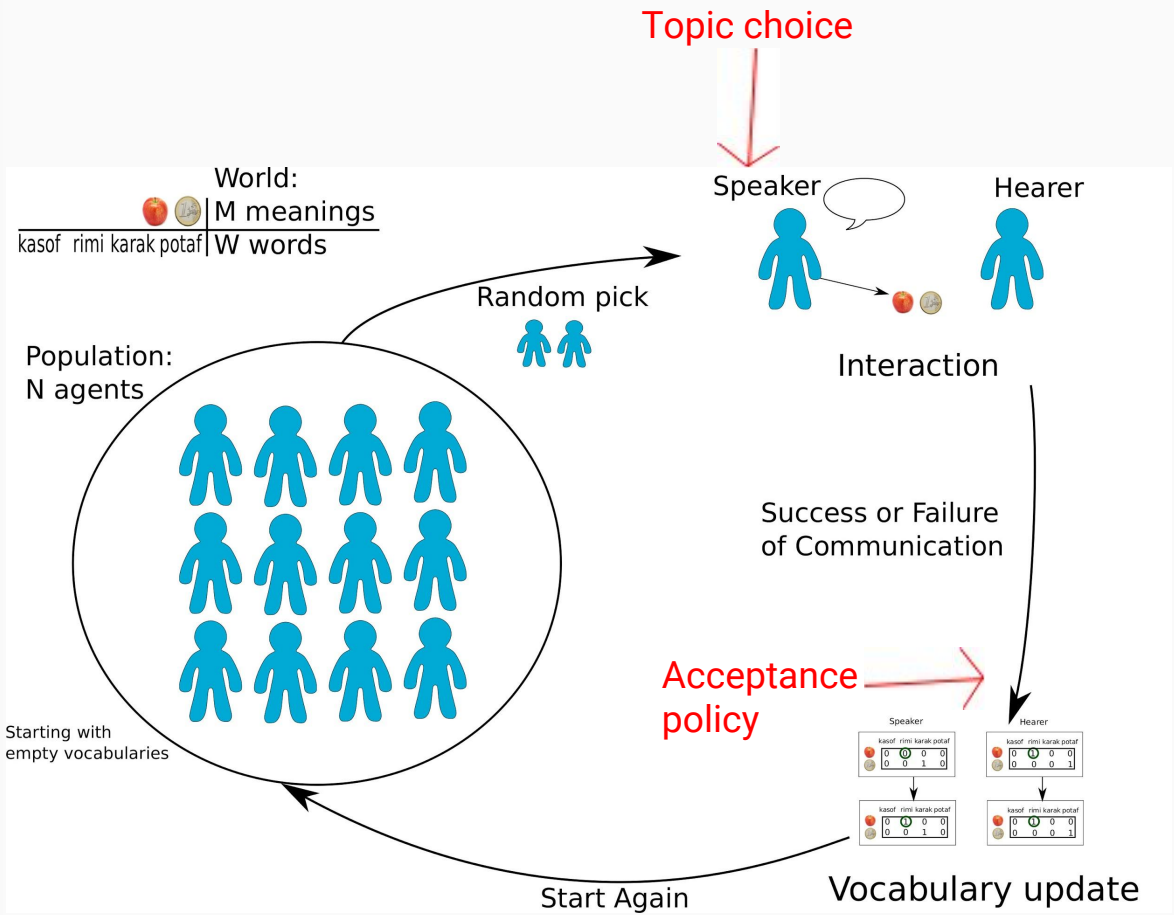

Active Topic Choice

- Stick to few meanings, later only explore others

CONTROLLING RATE OF INNOVATION (exploitation/exploration)

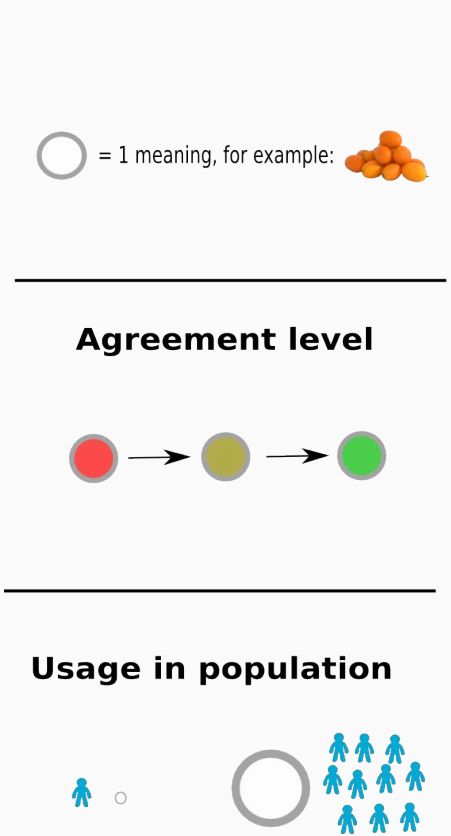

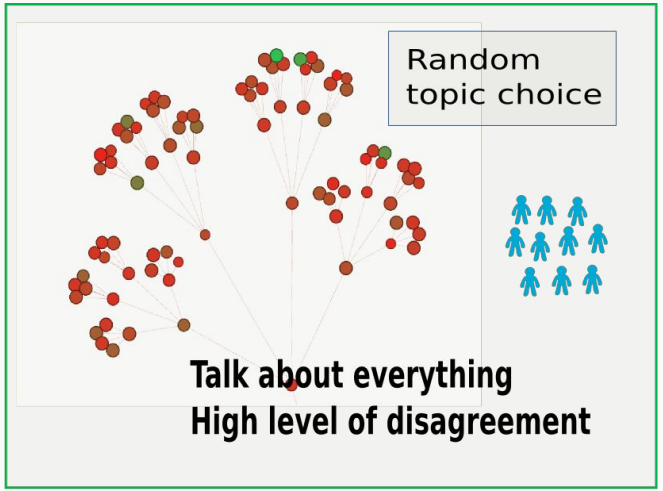

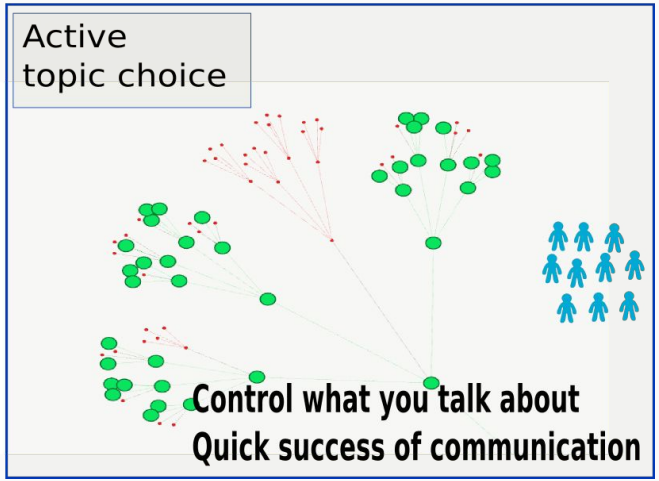

https://www.dropbox.com/s/o71jjsv7q8pk30w/ngal_struct_1.mp4?dl=0

Active Topic Choice

- Stick to few meanings, later only explore others
- Several algorithms based on measures of confidence

Random choice

Active choice

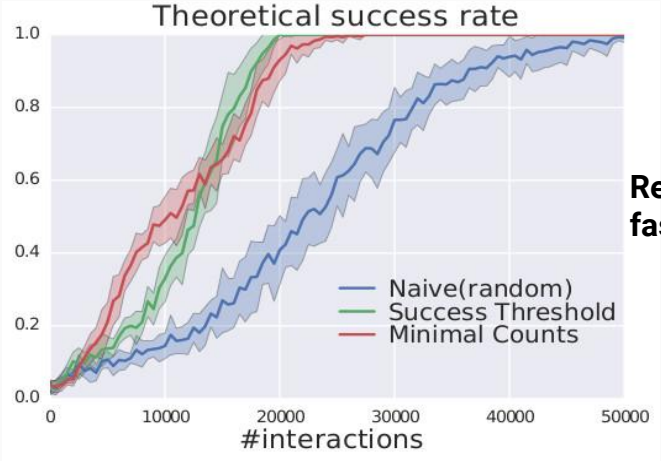

Reach consensus faster

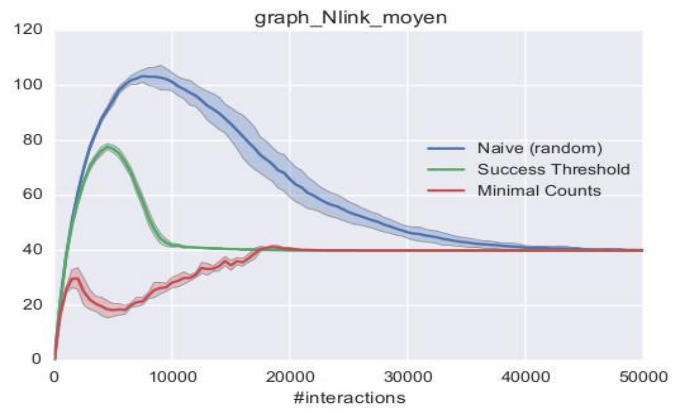

Less memory usage

Active control of complexity growth !

Acceptance policy: “Do I trust what he just said?”

- Stick to few meanings, later only explore others
- Several algorithms based on measures of confidence

Local complexity (memory) remains low!

Normal

Active

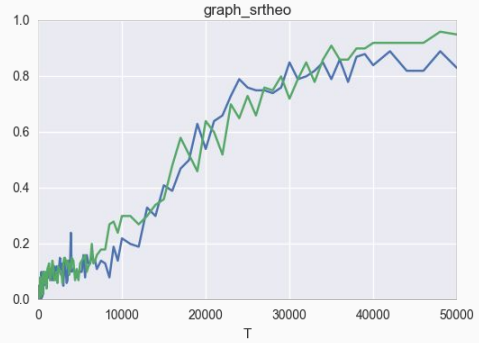

Success rate

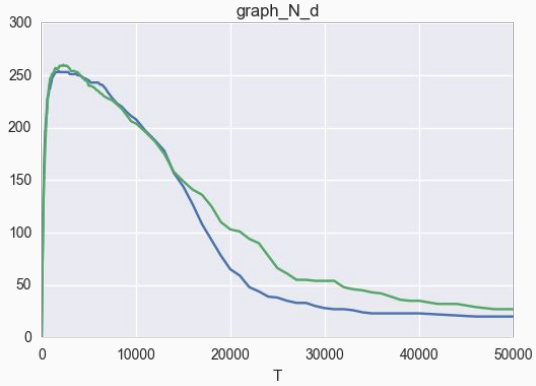

Global complexity

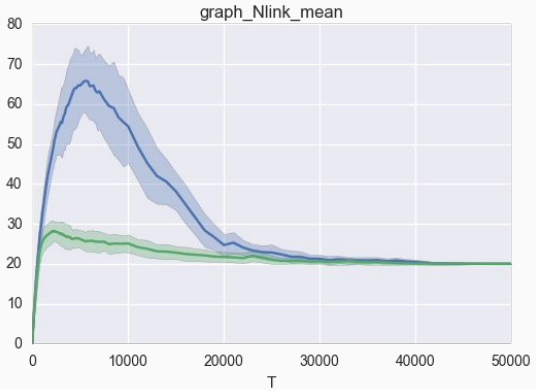

Local complexity

Human behavior?

Previous Naming Game experiments:

Issues:

- Recruiting users
- Getting/keeping them motivated
- If interacting together, should not abandon
- Haing relevant measures and results

Human behavior?

Previous Naming Game experiments:

- A. Baronchelli and D. Centola: single meaning, no active topic choice possible, focused on social network properties

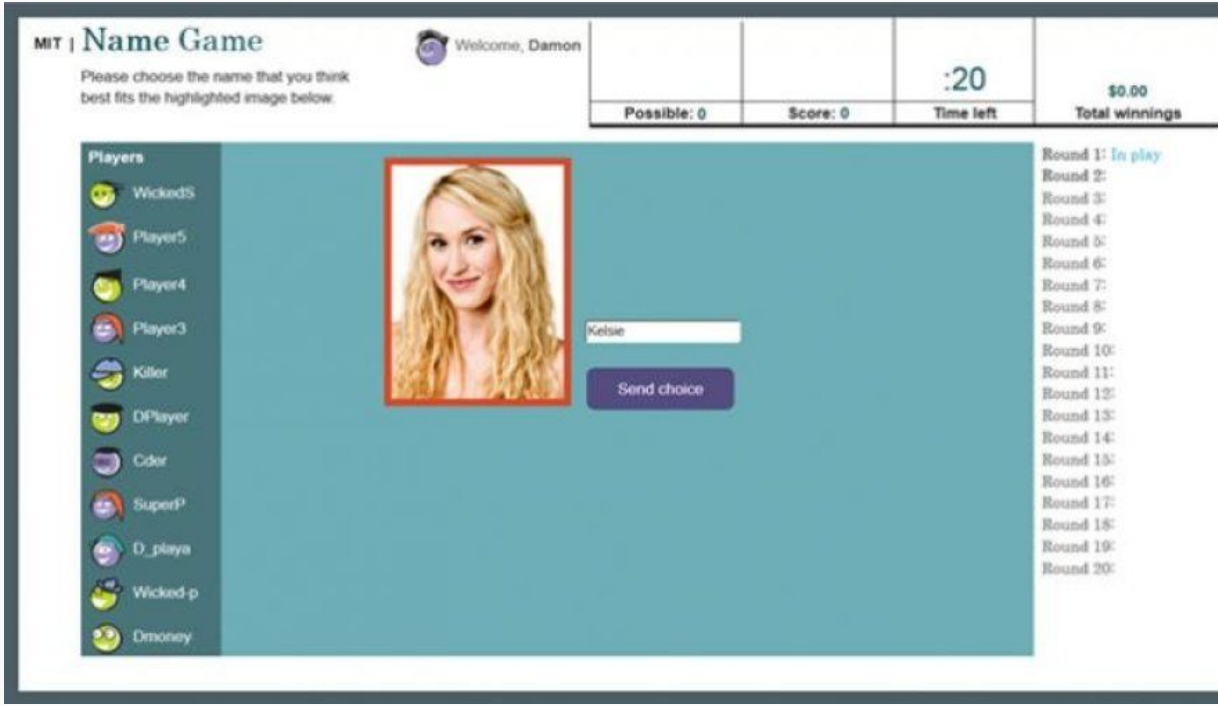

Human behavior?

Previous Naming Game experiments:

- A. Baronchelli and D. Centola: single meaning, no active topic choice possible, focused on social network properties

Important issue:
Gather participants, at the same time and keep them interested!

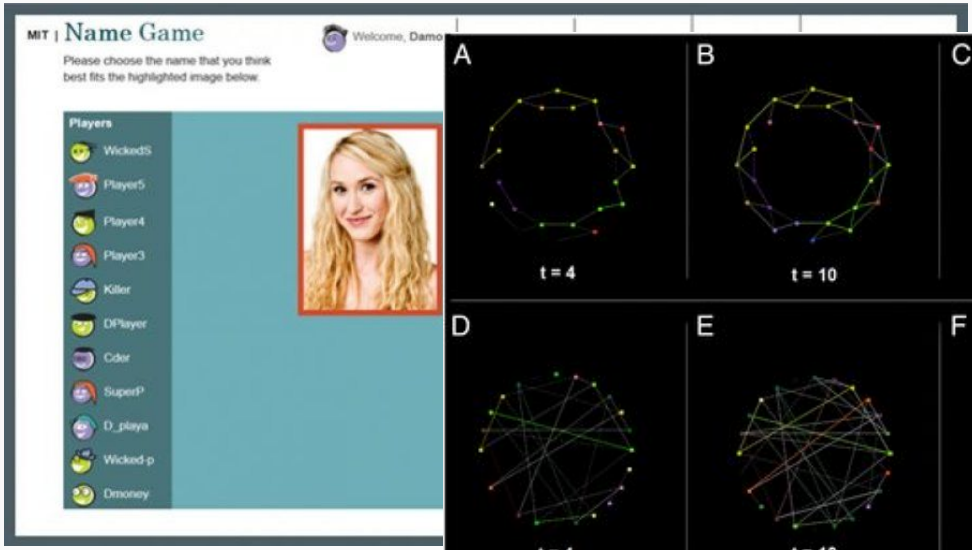

Network of potential neighbors: Consensus reached only when everyone interacts with everyone

Human behavior?

Previous Naming Game experiments:

- Talking Heads (Luc Steels) : too global, consensus not a goal per se, more about the artificial agents

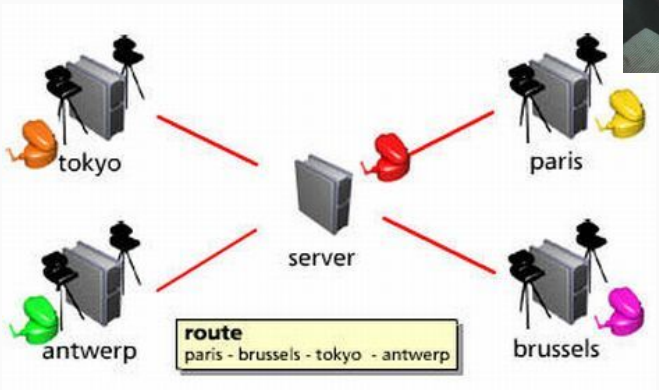

People can interact with the agents and force their own words in the system. But:

1. **Actual words already exist (square, red, circle, …)**
2. **The experiment focused on the dynamics of the simulated agents' conventions**

Human behavior?

Previous Naming Game experiments:

- Ergo-robots: artistic display of the model

Free participants, but data is not relevant!

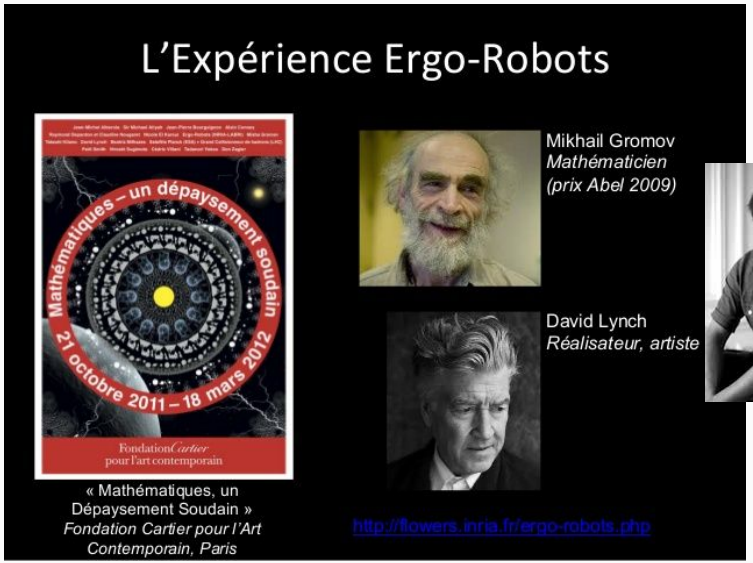

Pierre-Yves Oudeyer
A good supervisor

New experiment platform : an actual game !

- Focused on meaning choice and meaning exploration
- For the moment only one user at a t
  + simulated agents
- Maximizing a score (probability of successful communication at the e

- Measures:
  - Bias innovation vs. reusing known meanings
  - Rational choice?
  - Persistence of own inventions (IKEA effect?)

http://kreyon.net/naming-game

Or just outside, with the other experiments !

New experiment platform : an actual game !

Maybe by tomorrow, bonus types of experiments:

- Successfully communicating about meanings unravels new ones (motivation other than the score)
- Collective experiment, reusing agents who interacted with other people.

http://kreyon.net/naming-game

Or just outside !

Thanks

Elisabetta Falivene

Théo Segonds

Usual strategies

Success Threshold

Minimal Counts

Hearer s choice

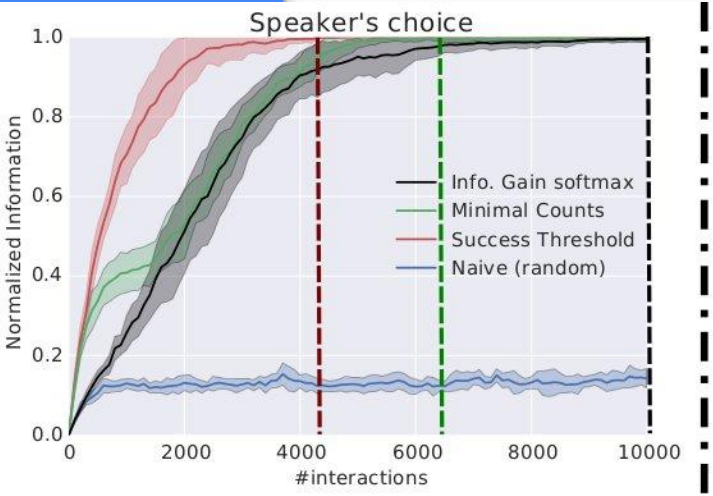

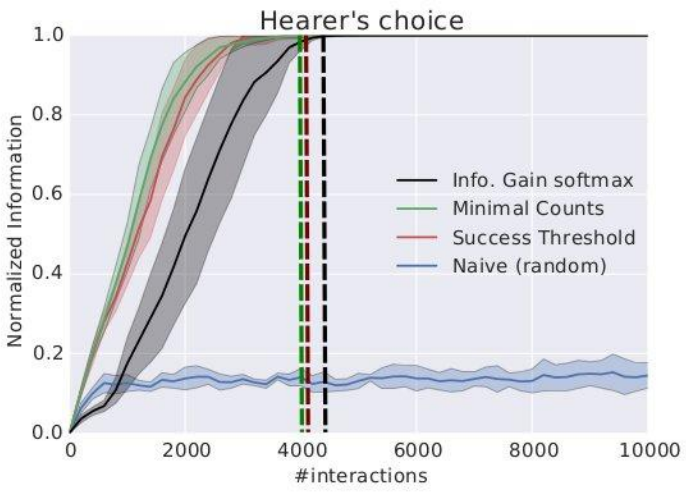

- Naive strategy converges slowly (after 1.000.000 interactions – not depicted here).
- **Hearer's choice policy is more efficient** for all active learning strategies.
- Last 5% of information are acquired slower when the speaker is choosing.

Vertical lines show full convergence time for each strategy. (M=W=N=20, averaged over 8 trials)

Hearer s choice

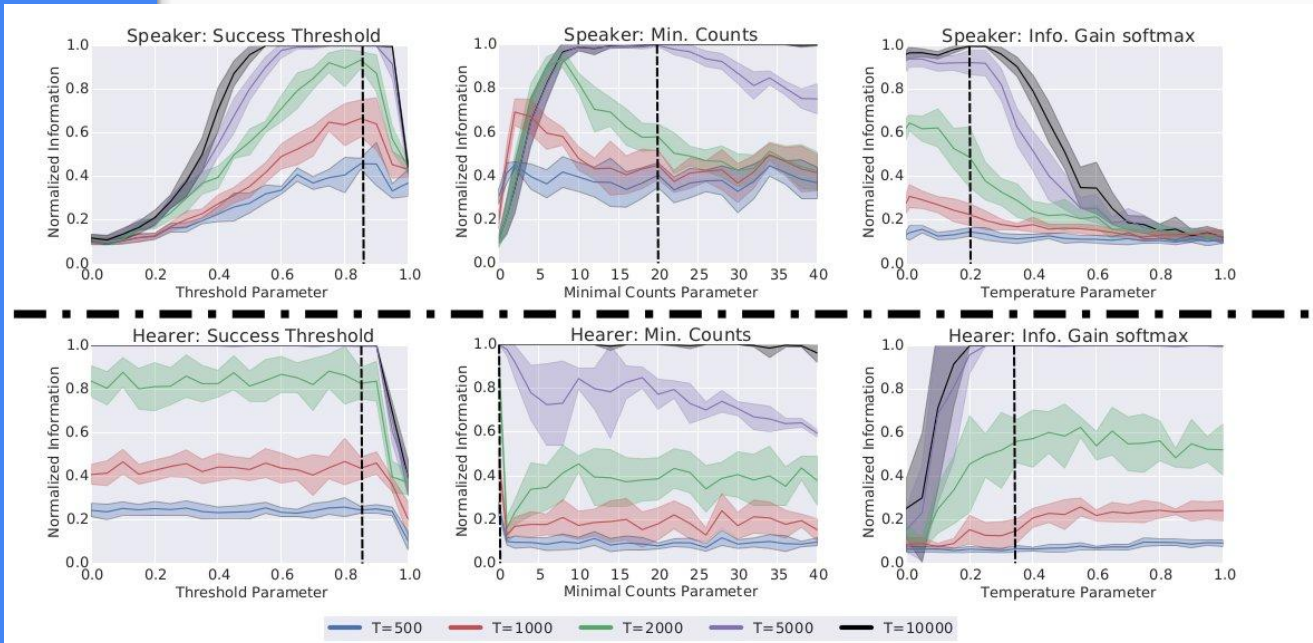

Convergence speed dependance on strategy parameters, for active 3 strategies and 2 active interaction scenarios. In all cases **hearer's choice scenario parameters are more robust** to change in value. Snapshots are taken for concurrent strategies spanning a relevant parameter interval, at different time steps (500, 1000, 2000, 5000, 10.000 interactions). Vertical lines show parameter values chosen for the comparisons in results section. (M=W=N=20, averaged over 8 trials)

Online platforms

- Amazon Mechanical Turk
- CrowdCurio

Python library

Important features:

- Modularity
- Open source
- Reproducible research
- Scalable on cluster
- Easy to read, modify and evolve (Supposed to be)

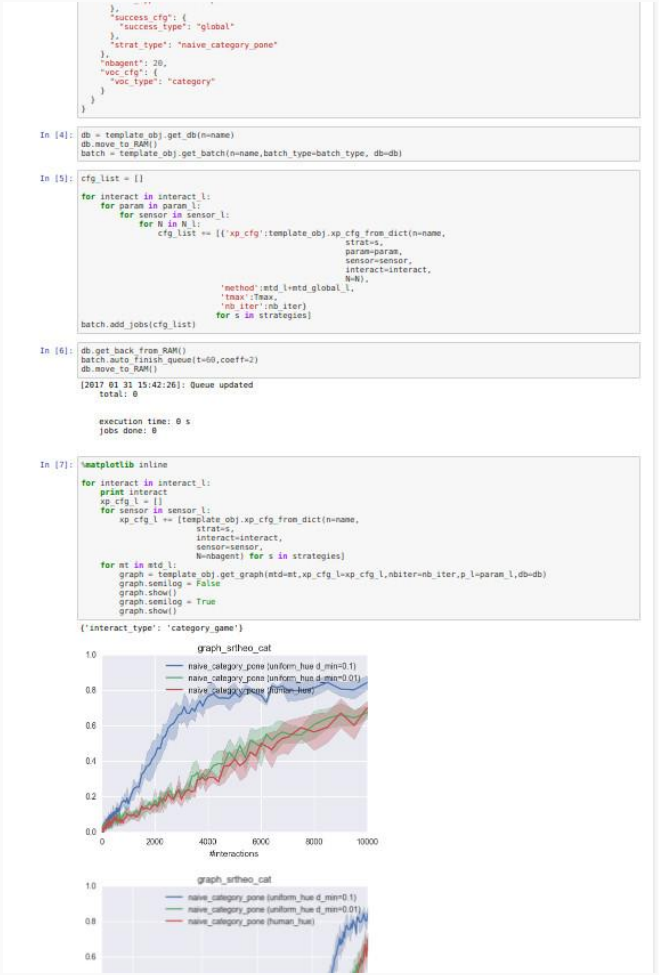

http://github.com/flowersteam/naminggamesal

Human behavior?

Previous Naming Game experiments:

- Iterated Learning (K. Smith and S. Kirby):
  every user is a “generation”, no need to sync users

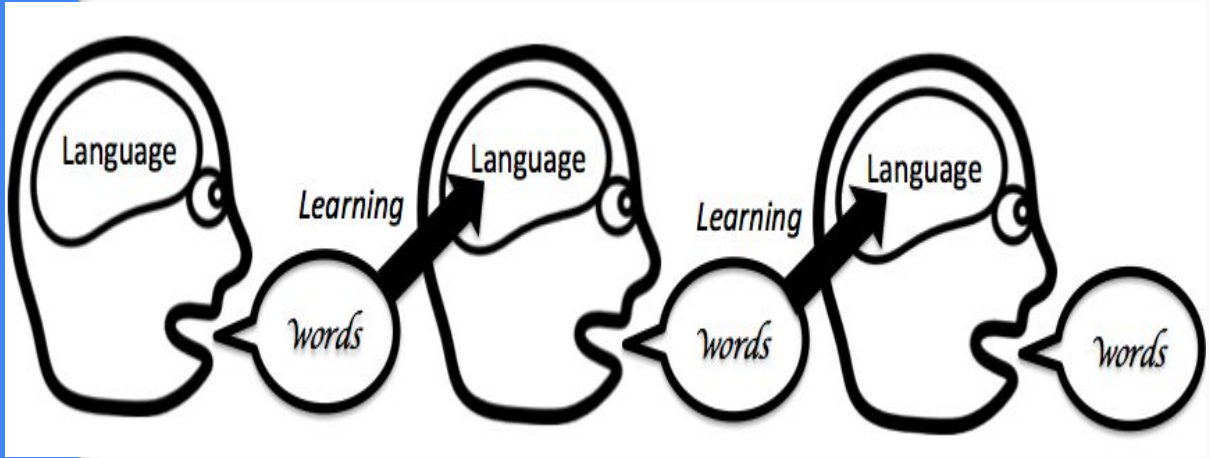

Definition of the local measure of probability of success

What is used by the agent?

- Vocabulary
- Memory of past interactions

Idea: build an approximation of the ‘average vocabulary’ in the population using this memory, and compute probability of success

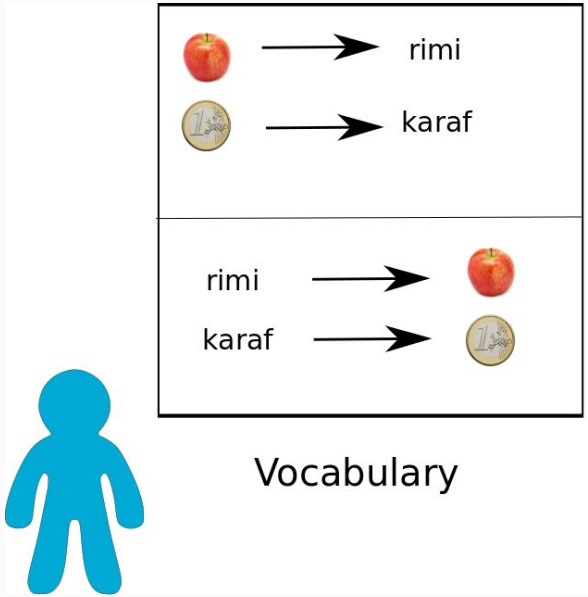

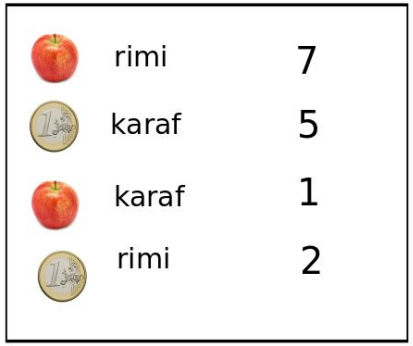

Definition of the local measure of probability of success

Normalization: only retain maximum t interaction per meaning or word

(ignore past interactions older than t -> sliding window per meaning and per word)

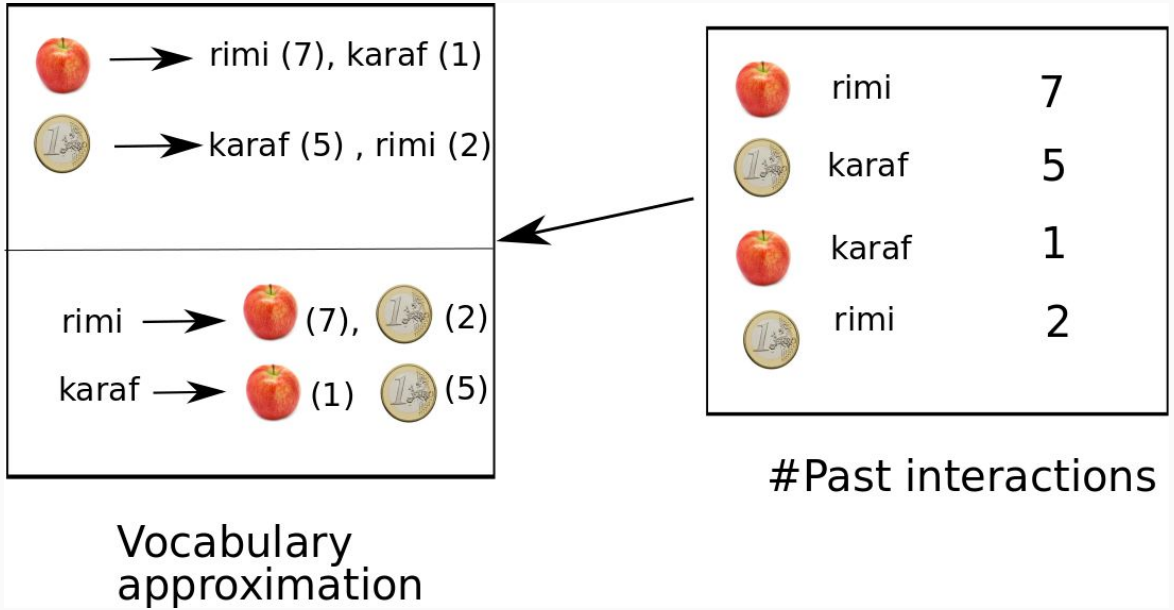

Definition of the local measure of probability of success

Normalization: only retain maximum t =10 interactions per meaning or word

(ignore past interactions older than t -> sliding window per meaning and per word)

Consider 2 possible options: success when speaker, success when hearer, then average over meanings

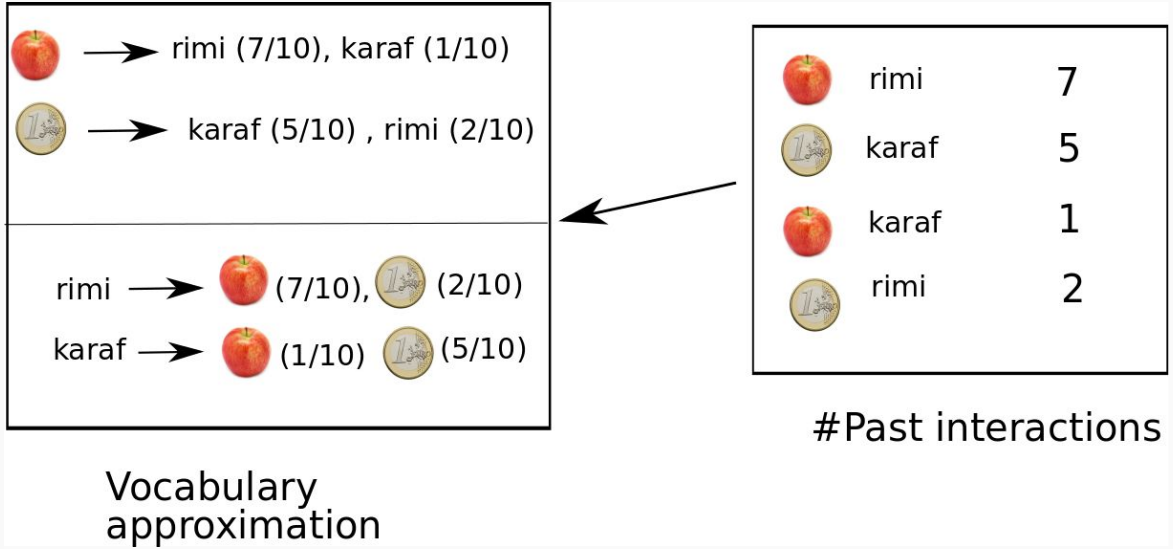

SR(,speaker) = (→ rimi)$_{Voc}$ x (rimi →)$_{Mem}$ + (→ karaf)$_{Voc}$ x (karaf →)$_{Mem}$

= 1 x 7/10 + 0 x 1/10

= 7/10

SR( ) = 1/2 * SR(,speaker) + 1/2 * SR (, hearer)

= 0.49

SR_total = 1/M * SR( ) + 1/M * SR ( )

= 0.37

Definition of the local measure of probability of success

Importance of time scale:

- If too small, too much confidence over new conventions (can be own), and forgetting lots of info
- If too high, evaluation of probability of success can be too small

But always, when over the time period all interactions agree with voc, SR = 1

Example with the previous example:

t = 9 -> SR = 0.457

t = 11 -> SR = 0.306

t= 100 -> SR = 0.0037

t < 9: impossible because 7+2=9, would have had to remove oldest interactions from memory

Here, numbers do not sum to 10, some information is still missing from the agent's point of view!

Multi-Armed Bandit

- Algorithms for exploit/explore
- Updating beliefs about reward, not using actual computation of the reward (very costly in our case)

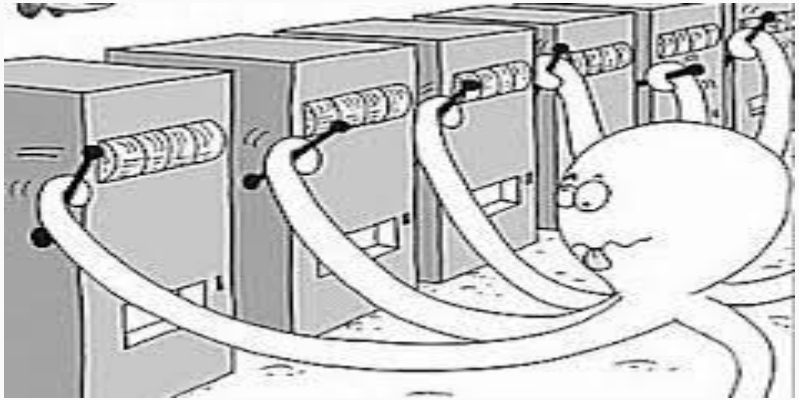

Implemented:

Thompson Sampling, one “slot machine” per meaning + 1 for exploration. New machine (=newly explored meaning) takes current state of exploration machine

Abstract: Collective exploration through communication

Naming Game on a balanced tree (4 children per node), opening up adjacent possible only if condition of ‘good communication’, for example:

- At least 1 success
- At least k success
- Being certain about this meaning (estimated proba of success = 1)

N=100, M=W=10^4

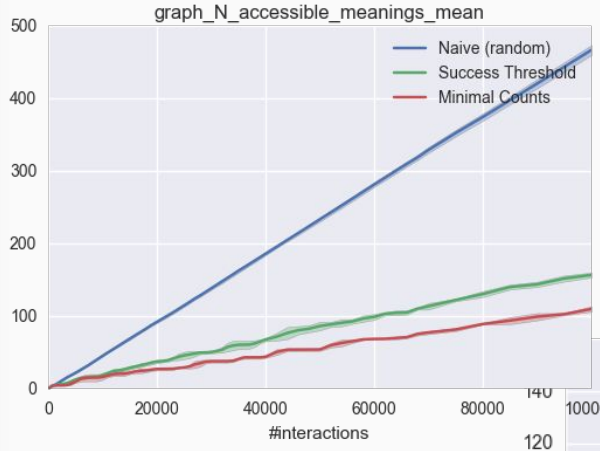

1 success

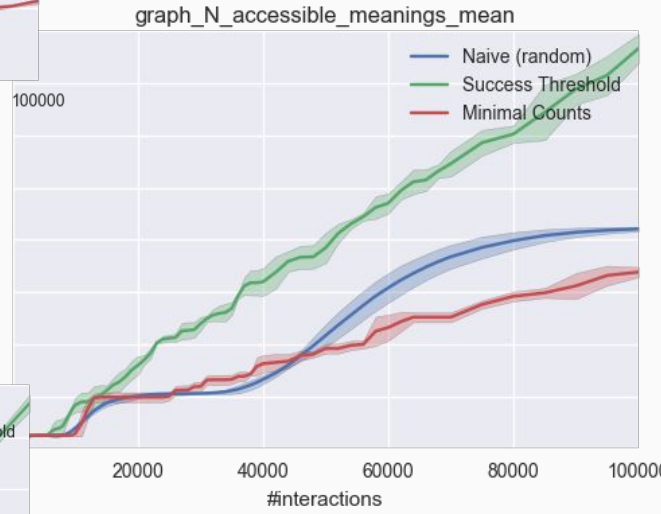

10 success

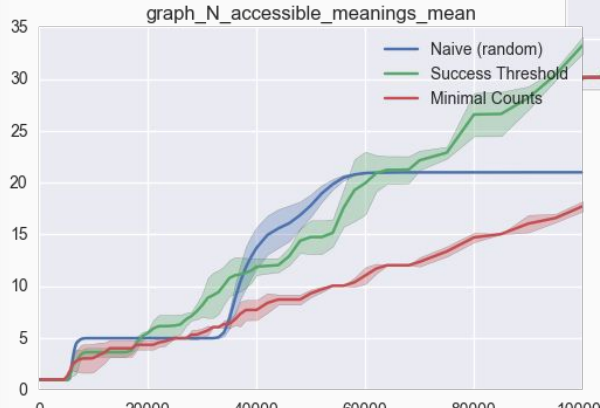

Proba of success (m) = 1