Emergent Communication

with Entangled Input

Andre Cianflone

McGill School of Computer Science

Mila

Dialogue agents

3,625 views | Aug 27, 2018, 08:19am

Chatbots Are Killing Customer Service. Here's Why.

Christopher Elliott Contributor

I'm a consumer advocate. I write about customer service.

amazon alexa prize

Google Duplex: An AI System for Accomplishing Real-World Tasks Over the Phone

Tuesday, May 8, 2018

Training dialogue agents

- Traditional approach: **supervised learning**, eg: Ubuntu Dialogues Corpus (Lowe et al, 2015)
- Downside:
    - Not grounded
    - Not necessarily compositional
    - Not goal oriented

Outline

- Emergent Communication (EC)
- EC with symbolic input
- EC with pixel input
- EC with generative models

Alternative: AI learns its own language

- Learn end-to-end in **virtual environment**
- To solve a task, agents **must learn to communicate**
- **Reward** only for solving task, not language
- Language emerges

Emergent Communication

- "Language derives meaning from its use" (Wittgenstein, 1953)
- Let agents **evolve a language** in order to solve a task. Desired properties:
    - **Grounded**: Meaning from use, objects, experience
    - **Compositional**: meaning of the whole composed from the meaning of smaller units
- For example:
    - Referential games (Lazaridou et al, 2018)
    - Negotiate resource allocation (Cao et al, 2017)
    - Self-driving cars (Resnick et al, 2017)

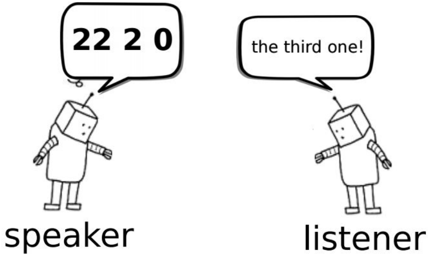

Motivations

- Need common language for **AI agents to coordinate**
- Agents need language capacity to **communicate with humans**
- Informs us on the evolution of language in humans, necessary conditions for biological language, insights into cognitive science
- Agents produce structured messages when they uncover the true **factors of variation**
- Maybe: **Transfer learning**, emergent language can help learn natural language

Emergent Communication with Symbolic Input

Referential Game: Task & Talk

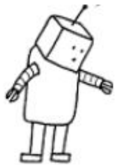

Agent-A

given object: 

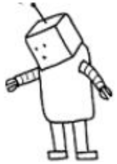

Agent-Q

guess attribues:
(colour, shape)?

Agent-A and
Agent-Q
communicate
in 2 rounds

Agent-Q guesses
the two attributes

Reward is shared,
train with
REINFORCE

Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog, Kottur et al, 2017

Unnatural Emergence

- Task is **solved perfectly**, however...
- With large vocabulary, every possible combination is encoded
- **Some utterances are ignored**

Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog, Kottur et al, 2017

Necessary Conditions

- Constraints force the language to be more "**natural**"
- Must **limit vocabulary** size
- A-Bot must **forget past utterance** →now individual symbols are grounded to attributes
- Language is **compositional**

Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog, Kottur et al, 2017

Communication Result

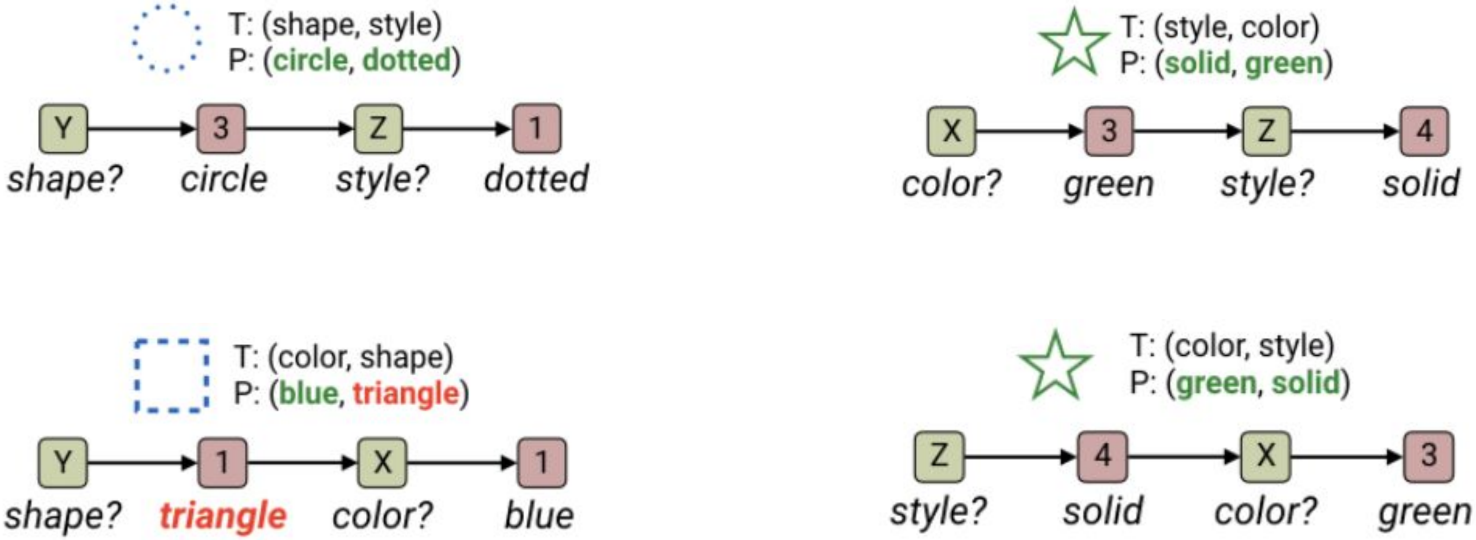

Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog, Kottur et al, 2017

Entangled Input

Entangled Input

- Typical setup has disentangled input →
- But the real world is entangled...

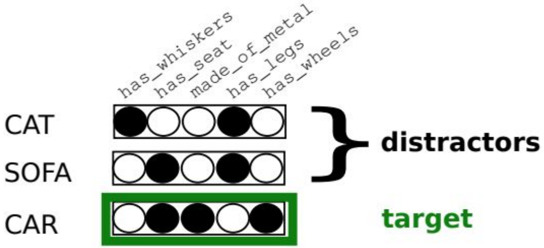

Referential Game with Images

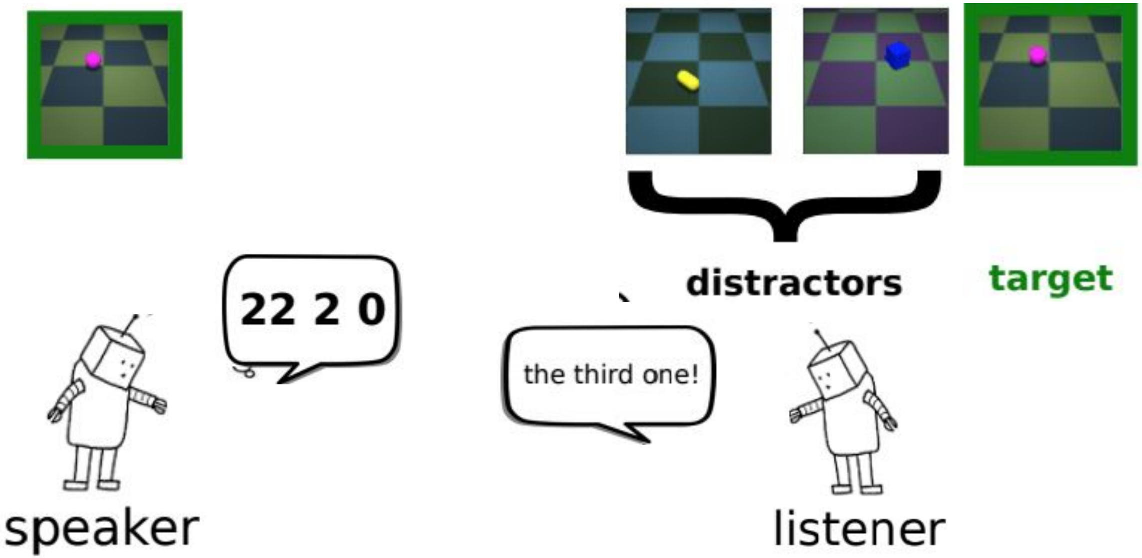

Results

game	distractors	balanced	viewpoints	lexicon size	random	train	test	topographic ρ
A	20	No	No	1068	5.0	93.7	93.6	0.13
B	2	No	No	13	50.0	93.2	93.4	0.006
C	2	No	Yes	8	50.0	86.0	85.7	0.07
D	2	Yes	Yes	5	50.0	90.4	89.9	0.06

Table 3: Communicative success of agents playing different games. Columns **random**, **train** and **test** report percentage accuracies. Column **topographic** ρ reports the topographic similarity between the symbolic representation of scenes and the generated messages ($p < 0.01$, permutation test).

Emergence of Linguistic Communication from Referential Games with Symbolic and Pixel Input. Lazaridou et al, 2018

Consistency of Messages

- **Topographic similarity:** semantically similar languages have similar messages

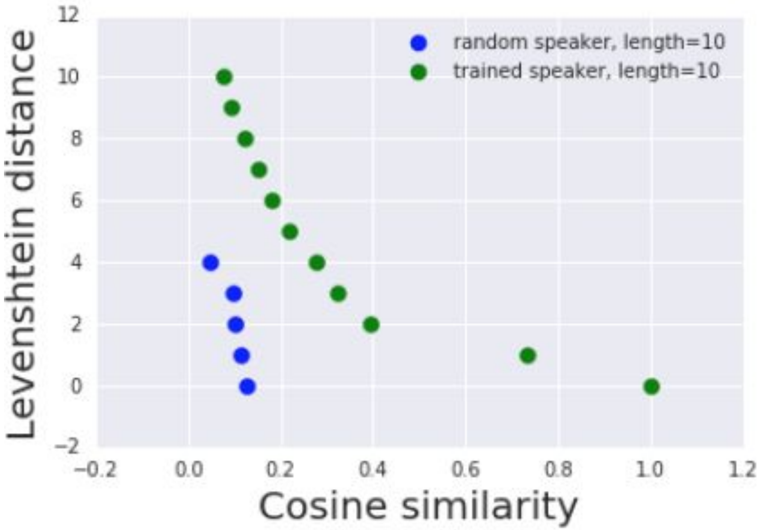

Emergence of Linguistic Communication from Referential Games with Symbolic and Pixel Input. Lazaridou et al, 2018

Generative Models

EC with Generative Models

- We want to see if we can **evolve language in a generative model** setup, conditional GANs in our case
- **Disentangle factors of variation**: high level concepts are words, completely unsupervised setup
- Natural way to **manipulate generative models**: write "words"

EC with Conditional GAN (EC-GAN)

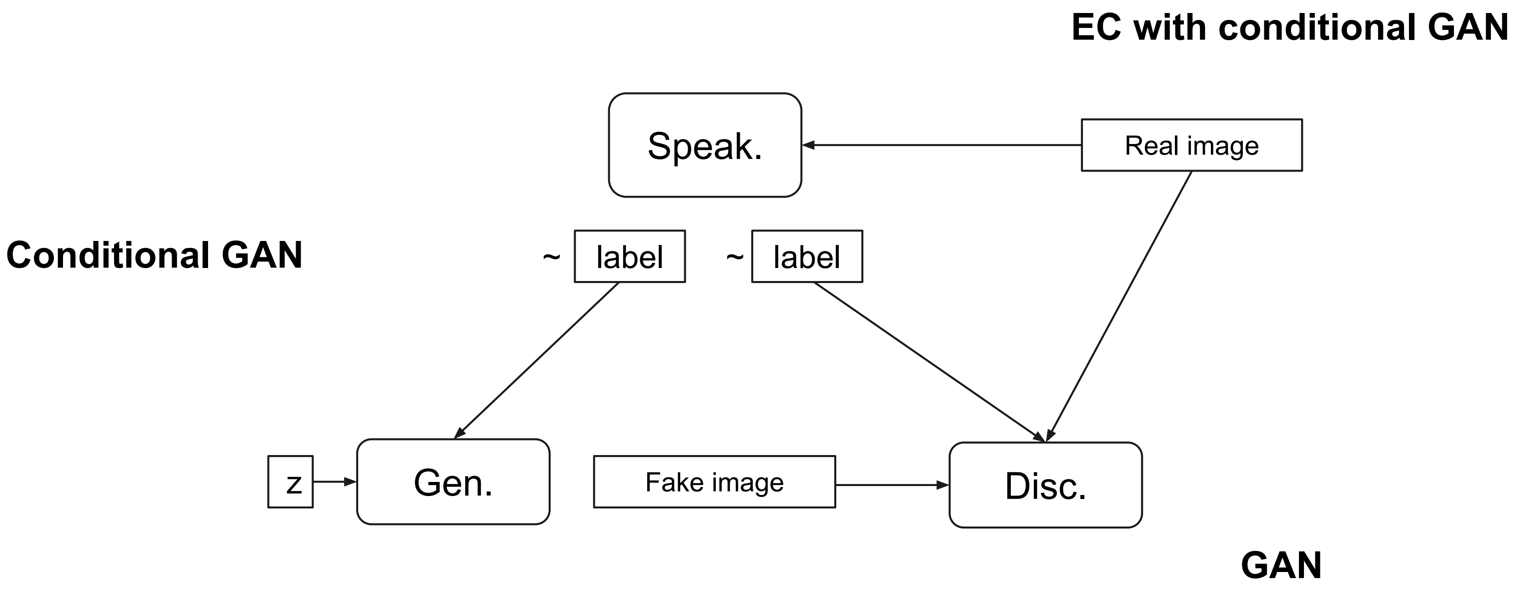

EC-GAN

- **Cooperative/competitive** task with 3 deep ConvNets
- Let a description of images evolve to help discriminator and generator compete
- Message is **discrete**, no gradient flow to speaker from GAN
- Completely **unsupervised**: No real labels seen by any agent

EC-GAN Objective

Generator $J^{(G)} = -\frac{1}{2}\mathbb{E}_z \log D(G(z)|\tilde{y})$ Try to fool G

Discriminator $J^{(D)}\left(\theta^{(D)}, \theta^{(G)}\right) = -\frac{1}{2}\mathbb{E}_{x\sim p_{data}} \log D(x|\tilde{y}) - \frac{1}{2}\mathbb{E}_z \log(1 - D(G(z)|\tilde{y}))$

Speaker $\theta \leftarrow \theta + \alpha \ln \pi(A_t|S_t, \theta)(G + \mathcal{H})$

$$R = JSD(D(x|\tilde{y})||D(x|\hat{y})) + JSD(D(G(z))|\tilde{y})||D(G(z))|\hat{y}), \quad \hat{y} \sim U$$

Policy contribution to D compared to random policy

Policy contribution to G compared to random policy

Initial Results

Initially, seems like some correlation?

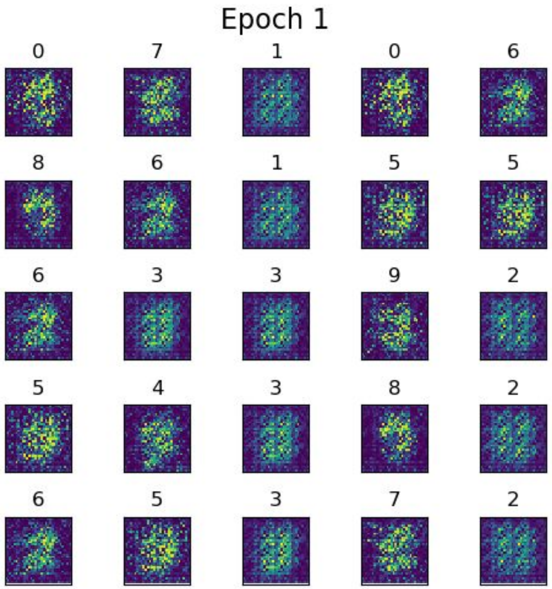

Stability Issues

GANs are unstable

Gan + RL = more instability

If policy converges too quickly, the discriminator can learn quite quickly at the generator's expense

Eventually, Generator error forces drastic policy change. Policy becomes uniform and the conditional GAN converges to a regular GAN, ignoring labels!

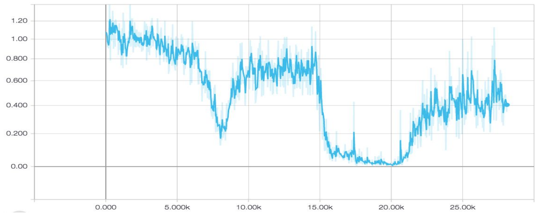

Discriminator loss

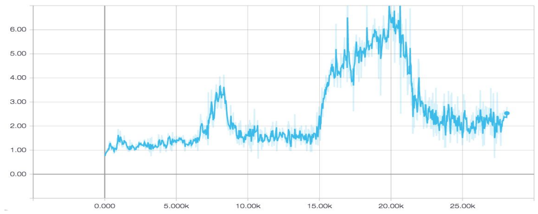

Generator loss

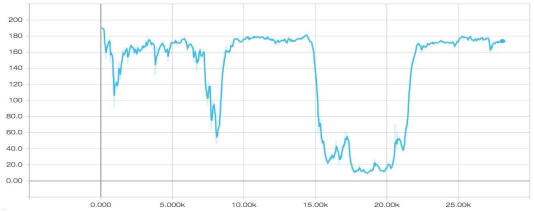

Policy entropy

Results

The EC-GAN learns **grounded symbols**, unsupervised

Tested on various vocabulary sizes: 10, 20, 30 discrete symbols

Generator images are not as good as GAN conditioned on true labels

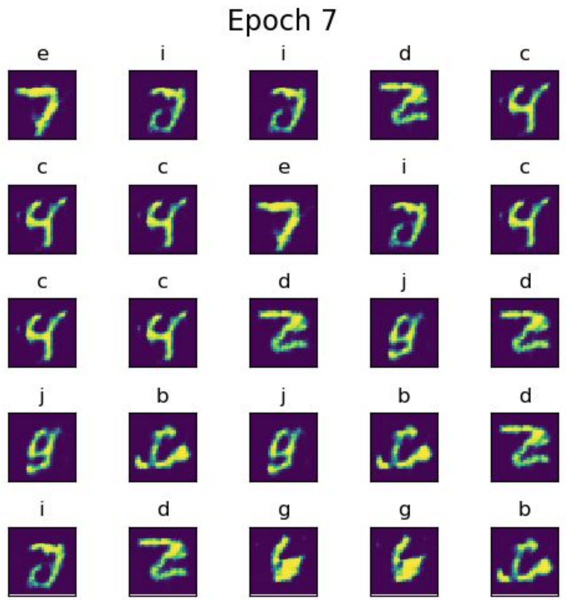

Conclusion

What helps:

- Fine tune hyper parameters, entropy regularization, return definition

- Bound the loss

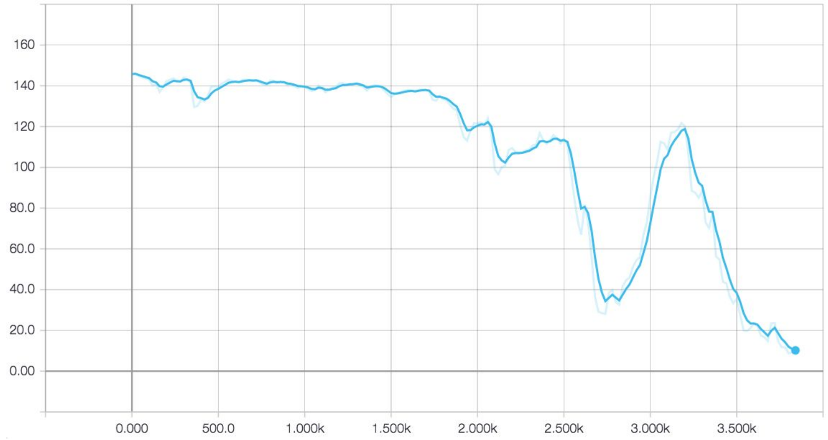

Train D/G until under loss upper bound

Future

- Simple AE baseline
- Different dataset
- Episodic GAN
- Language probing: Train classifier conditioned on language
- Variable length message. Does message complexity correlate with dataset complexity?
- Two-way messages

Thank you!

andre.cianflone@mail.mcgill.ca

References

Cao, Kris, Angeliki Lazaridou, Marc Lanctot, Joel Z. Leibo, Karl Tuyls, and Stephen Clark. "Emergent Communication through Negotiation." NIPS 2017 Workshop on Emergent Communication.

Kottur, Satwik, José MF Moura, Stefan Lee, and Dhruv Batra. "Natural Language Does Not Emerge'Naturally'in Multi-Agent Dialog." EMNLP (2017).

Lazaridou, Angeliki, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. "Emergence of linguistic communication from referential games with symbolic and pixel input." ICLR (2018).

References

Lowe, Ryan, Nissan Pow, Iulian Serban, and Joelle Pineau. "The Ubuntu Dialogue Corpus: A Large Dataset for Research in Unstructured Multi-Turn Dialogue Systems". SIGDIAL (2015)

Resnick, Cinjon; Kulikov, Ilya; Cho, Kyunghyun; Weston, Jason. "Vehicle Communication Strategies for Simulated Highway Driving", NIPS 2017 Workshop on Emergent Communication.