

Discussion #3

Student-led discussions

Discussion leader

Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog

Satwik Kottur¹ and José M.F. Moura¹ and Stefan Lee^{2,3} and Dhruv Batra^{3,4}
¹Carnegie Mellon University, ²Virginia Tech, ³Georgia Tech, ⁴Facebook AI Research

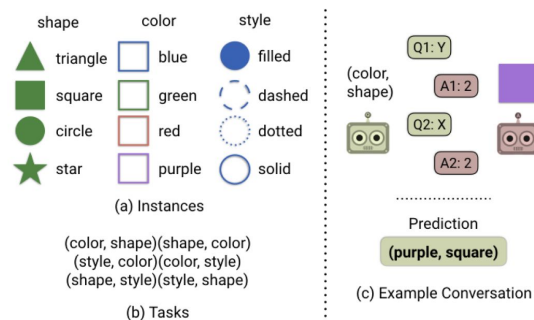


Figure 1: (a) *Task & Talk*: The testbed for our study is cooperative 2-player game, *Task & Talk*, grounded in a synthetic world of objects with 4 *shapes* \times 4 *colors* \times 4 *styles*. (b) Q-BOT is assigned a *task* – to inquire about the state of an ordered pair of attributes. (c) An example gameplay between the two agents - Q-BOT asks questions depending on the task which are answered by A-BOT conditioned on the hidden instance visible to only itself. At the end, Q-BOT makes a prediction of pair of attributes (*purple, square*).

Summary

- What are the conditions that lead to the emergence of human-interpretable or compositional grounded language?
 - Invented languages not compositional, interpretable, or natural
 - There are conditions that must be present before compositional grounded language emerges
- Some definitions:
 - Grounding: mapping words to physical concepts
 - Compositionality: combining knowledge of simpler concepts to describe richer concepts
- Task and talk game:
 - Information asymmetry: A-bot sees the object, Q-bot knows the task, bots exchange utterances to help Q-bot find the answer to the task
 - Cooperative reward: bots rewarded identically based on accuracy of Q-bot's prediction

Experiment Settings

- Overcomplete vocabularies: When vocabulary of A-bot is greater than number of instances
 - Codebook that maps symbols to object instances, no dialog is necessary
 - Fails for test instances with novel combinations of attributes

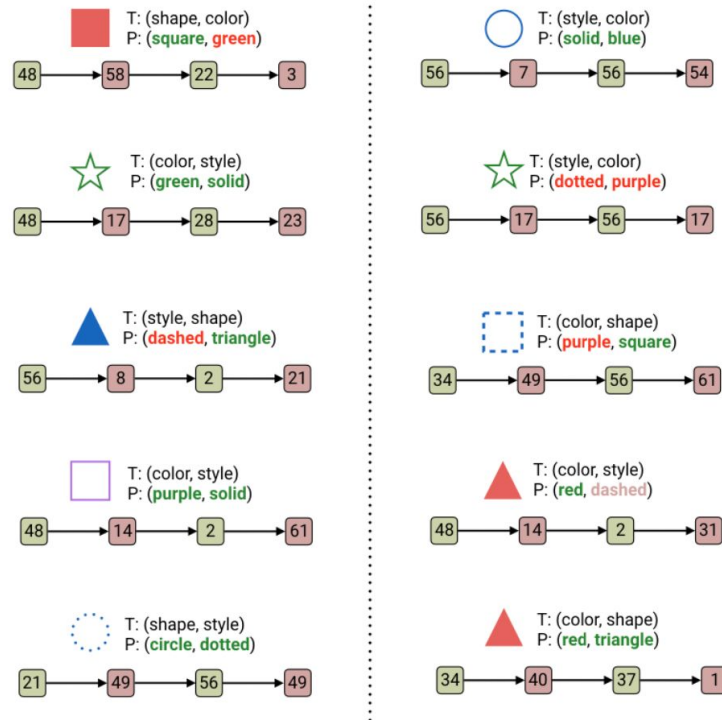


Figure 3: Overcomplete vocabularies setting ($|V_Q| = |V_A| = 64$, Sec. 4.1). Owing to a large vocabulary, we denote the tokens using numbers, as opposed to English alphabet characters shown in other figures. A-BOT mostly ignores what Q-BOT asks and instead conveys the instance using pairs of symbols across rounds unique to an instance, leading to a highly non-human intuitive and non-compositional language.

- Attribute-value vocabulary: Size of vocabulary depends on number of attribute values
 - Tasks were encoded in order-agnostic manner, second round of utterance of Q-bot is unnecessary
 - Able to communicate the task, fails on unseen attribute value combinations (not compositional)

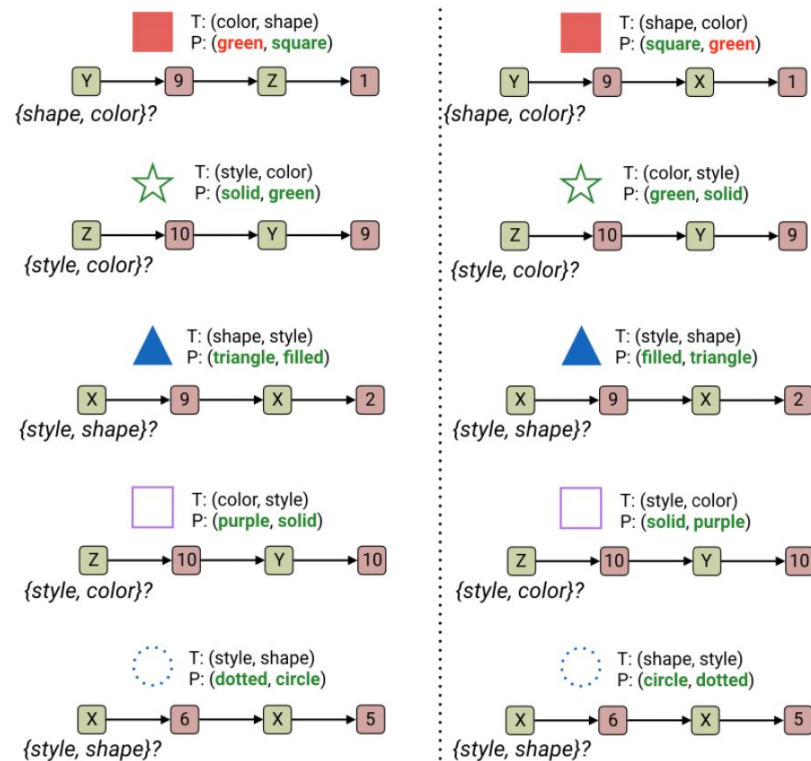


Figure 4: Attribute and Value vocabulary setting ($|V_Q| = 3$, $|V_A| = 12$, Sec. 4.2). We show symmetric tasks for each instance on either side to illustrate the similarities in the language between the agents. As seen here, Q-BOT maps symmetric tasks in an order-agnostic fashion, and uses only the first token to convey task information to A-BOT.

- Memoryless A-bot: limit vocabulary to reduce ‘synonyms’ learnt, remove memory (no longer able to distinguish between rounds)
 - Results in learned language that grounds individual symbols into attributes and their states

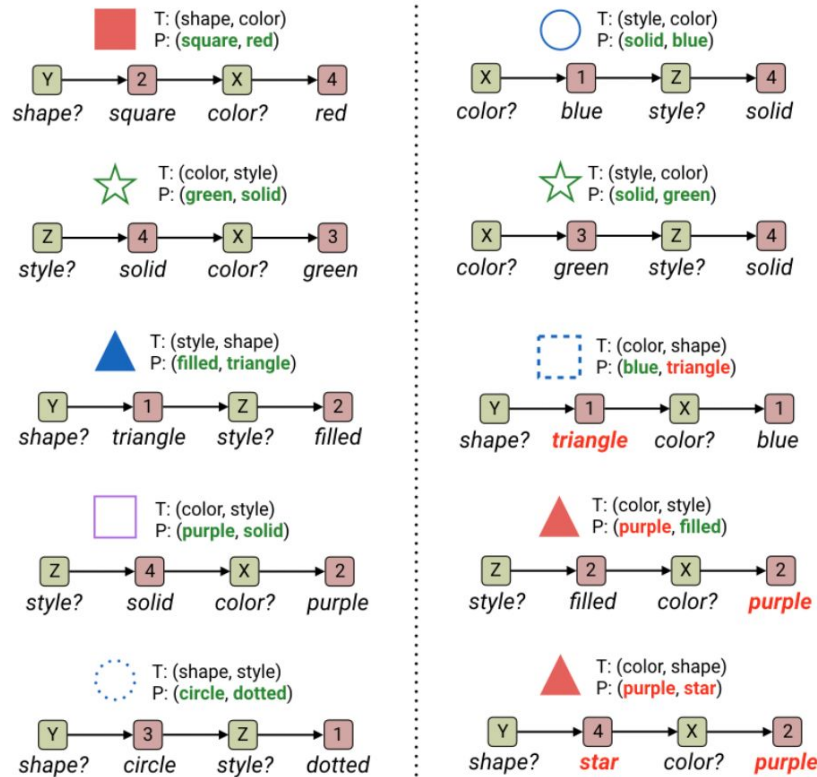


Figure 5: Example dialogs for memoryless A-BOT, minimal vocabulary setting ($|V_Q| = 3, |V_A| = 4$, Sec. 4.3). Learnt language is consistent and grounded, denoted below each token. Incorrect predictions on unseen instances (right, bottom) are also shown. Notice that the incorrectly predicted attribute is still from the right category (a color attribute for *color*).

Setting	Vocab.		Memory		Seen (%)		Unseen (%)		Characteristics
	V_Q	V_A	A	Q	Both	One	Both	One	
Overcomplete (§4.1)	64	64	✓	✓	100	100	25.6	79.5	<ul style="list-style-type: none"> • Non-compositional language • Q-BOT insignificant • Inconsistent A-BOT grounding across rounds • Poor generalization to unseen instances
Attr-Value (§4.2)	3	12	✓	✓	100	100	38.5	88.4	<ul style="list-style-type: none"> • Non-compositional language • Q-BOT uses one round to convey task • Inconsistent A-BOT grounding across rounds • Poor generalization to unseen instances
NoMem-Min (§4.3)	3	4	✗	✓	100	100	74.4	94.9	<ul style="list-style-type: none"> • Compositional language • Q-BOT uses both rounds to convey task • Consistent A-BOT grounding across rounds • Good generalization to unseen instances

Table 2: Overview of settings we explore to analyze the language learnt by two agents in a cooperative game, Task & Talk. Last two columns measure generalization in terms of prediction accuracy of **both** or at least **one** of the attribute pair, on a held-out test set containing unseen instances.

Main Points from Discussion Posts

- What we liked:
 - Clearly stated contributions and differentiation from predecessor's works
 - Problem formulation: Modeling resembled closely with human dialogue
 - Evolution of experiments: from overcomplete vocabulary to iteratively adding constraints
- What we hoped to see:
 - More extensive related works (cultural evolution paper was not cited)
 - Ablative studies on what changes if action and state space become larger
 - User studies to collect data on language understanding
 - What happens if the agents have prior knowledge, how would the conversation evolve?
 - How do results generalize to more realistic tasks? Would compositionality evolve when the state space is more complex?
- General comments:
 - Goals of communicating the attributes of real-world objects may not be aligned with other goals of language such as compositionality
 - Agreement with the authors on importance of incorporating human-guided actions and state restrictions to make agents perform desirable behaviour

ITERATED LEARNING FOR EMERGENT SYSTEMATICITY IN VQA

Ankit Vani*

Mila, Université de Montréal

Max Schwarzer

Mila, Université de Montréal

Yuchen Lu

Mila, Université de Montréal

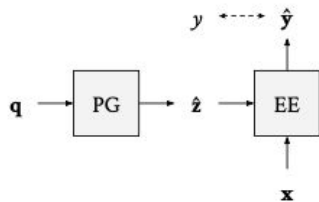
Eeshan Dhekane

Mila, Université de Montréal

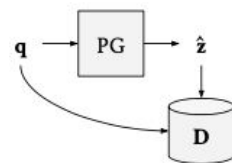
Aaron Courville

Mila, Université de Montréal, CIFAR Fellow

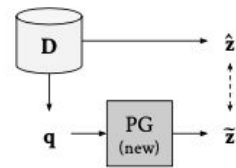
Summary



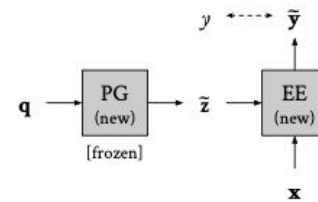
(a) Interacting phase.



(b) Transmitting phase.



(c) Program generator learning phase.



(d) Execution engine learning phase.

- A “use case” of cultural transmission in a specific ML problem space - visual question answering (VQA)
- Conceptually modeling a VQA task as a **collaborative game of two agents**
 - **Program generator (PG)** and **execution engine (EE)** in neural module networks (NMNs) as communicating agents in a cooperative VQA game
 - Programs (i.e. messages drawn from an emergent language) are passed between agents
- Tested on two VQA datasets, SHAPES and CLEVR
- Achieving similar performance as Johnson et al. (2017b) and Vedantam et al. (2019), but using a very small number of ground-truth programs for supervision

Main Points from Discussion Posts

Neat idea and overall clear presentation of the contributions

Found connection between this paper and the previous two - from theory to a specific application scenario 😊

The idea of “taking inspiration from human behavior to inform ML design” - potentially more explainability

The idea of viewing program generator and execution engine as two collaborative agents

Reinitializing the execution engine (but why?)

Seeing a trend in applying IL in algorithms; useful with data-hungry data sets

Would be interesting to see applications in other areas

Criticisms

Harder dataset than beyond SHAPES

Lack of comparison with other SOTA VQA methods, although partly understandable

Experiment design could be improved - few tasks and rationale of choice

Discuss its limitations and whether it's actually able to capture the deeper composition

Discuss IL in more complex scenarios

Archaeologist

Scientific Peer Reviewer (Skeptic)

Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog

Pros

- Covers important topic of interpretability and compositional issue in invented language.
- Studies generalization to novel environments
- Experiments are well designed
- Diagrams explain the methodology and ideas clearly.

Ratings

Scientific Quality: 4

Originality: 4

Relevance: 5

Clarity and Presentation: 3

Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog

Cons

- **Insufficient discussion on interpretability:** Major focus of discussion is on compositionality
- **Clarity on evaluation criteria could help future work**
 - Most of the discussion focuses on empirical understanding of results.
 - May be some user study would be helpful
- **No comparison with existing methods:** they do cite the base paper.
 - It would be nice to see existing area of work such as traditional rule-based system etc.

Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog

Cons

- No ablation study
- Limited discussion on future work, challenges and limitations
- No discussion on social and ethical impact of this work
- Limited explanations on words from the specific domain make it hard for people from other research areas to understand this.

Huge Wish

- **Missing study or citation on how dialogue emerges in human-human interaction.**

Iterated Learning For Emergent Systematicity in VQA

Main Contributions

- Uses iterated learning as a mechanism to promote **systematic generalization** by inducing compositional structure in Neural Module Networks (NMNs)
- Systematic Generalization is evaluated on **SHAPES-SyGeT** and **CLEVR/CLOSURE** using a small amount of ground-truth layout supervision

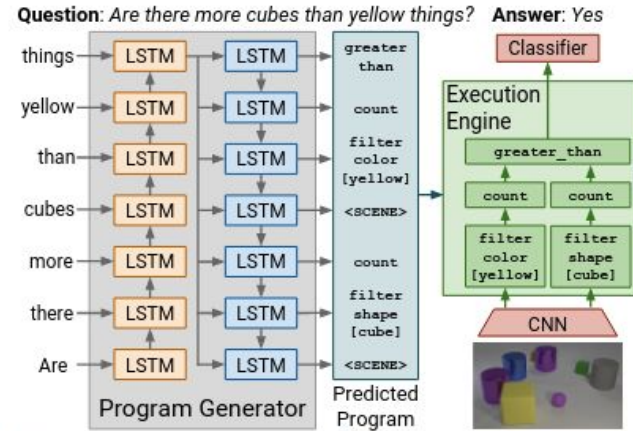


Figure 2. System overview. The **program generator** is a sequence-to-sequence model which inputs the question as a sequence of words and outputs a program as a sequence of functions, where the sequence is interpreted as a prefix traversal of the program's abstract syntax tree. The **execution engine** executes the program on the image by assembling a neural module network [2] mirroring the structure of the predicted program.

Inferring and Executing Programs for Visual Reasoning

Iterated Learning For Emergent Systematicity in VQA

Strengths

- Related work presented in ML and CogSci lay a strong foundation for proposed approach
- First paper to extend iterated learning beyond emulating what was done* in humans to **emerge compositional behavior for VQA (programs)**
 - Interesting design choice to model the program generator and execution engine as communicating agents
 - Requires less supervision than prior work on NMNs (form of supervised self-play)
- Ablation studies on SyGeT is mostly thorough
- Datasets chosen allows for **controlled experimentation**

*(Ren et al. (2020), Guo et al. (2020), Cogswell et al. (2020))

Iterated Learning For Emergent Systematicity in VQA

Weakness

- Performance on large-scale VQA datasets (as observed with GQA) is unclear - paper shares insight only on datasets which result in good performance even without imposing compositional structure (SHAPES and CLEVR)
- Gloss over decisions made wrt *learning bottleneck*
 - How do we decide the modes of reinitialization and # of generations?
 - For SyGeT - 5000 steps result in near perfect accuracy; **why choose PG phase length of 2000?**
- Limitations/**Failure modes** on SyGeT and CLOSURE templates could have been probed (qualitative results)
- Evaluation of what these resultant modules are specifying is lacking
- Case for systematic generalization would have been stronger with a *natural images* dataset instead of evaluation on CLOSURE
 - Results from CLOSURE could have been included in Supplementary material instead

Iterated Learning For Emergent Systematicity in VQA

- Things to improve that did not impact the score:
 - How challenging would the training process be if we used **transformers instead of seq2seq**? How could it affect generalization?
 - In Figure 2, authors could have simplified explanation for the phases of Iterated Learning in NMNs in the caption

- **Rating** - 6: Marginally above acceptance threshold
- **Confidence** - 3: You are fairly confident in your assessment. It is possible that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work. Math/other details were not carefully checked.

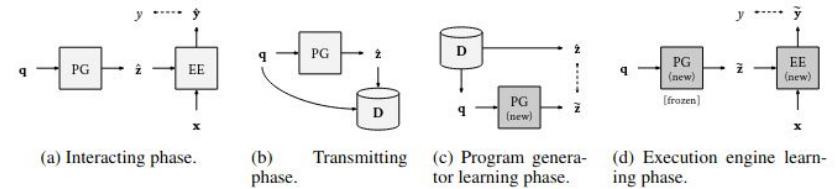


Figure 2: Phases of IL for emergent module layouts. Solid arrows indicate forward pass through the model, and dashed lines indicate the cross-entropy loss between predictions and targets. After proceeding through phases (a)-(d), the new program generator and execution engine begin an interacting phase (a) of a new generation.

Iterated Learning For Emergent Systematicity in VQA

- Things to improve that did not impact the score:
 - How challenging would the training process be if we used **transformers instead of seq2seq**? How could it affect generalization?
 - In Figure 2, authors could have simplified explanation for the phases of Iterated Learning in NMNs in the caption
- **Rating** - 6: Marginally above acceptance threshold
- **Confidence** - 3: You are fairly confident in your assessment. It is possible that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work. Math/other details were not carefully checked.

Expertise

Passing Knowledge

Originality

Medium originality

Significance

High significance

Rigor

High rigor

Social Impact Assessor

Natural language does not emerge naturally

- Why do we need artificial agents to be able to communicate with each other?
 - Does this communication have to have the same “compositionality” as human language? Why?
 - Paper finds that as far as “getting it done” is concerned, this isn’t required (maybe it emerges in more complex situations)
 - We design these systems for use by humans (usually) → human-language like compositionality is thus important
 - **Interpretability:** I wouldn’t want Alexa to talk with google home in a manner that I can’t understand (although Amazon and Google might collude to make that happen) -- just an example, if they wanted to they could share data wirelessly :P.
 - Is AI becoming less and less interpretable (previous lecture, refrigerator temp knob example).



Iterated Learning For Emergent Systematicity in VQA

- Shows emergence of natural language through iterative learning (the two papers are in contrast)
 - Emergence may be due to the fact that object position was also of significance. In the previous paper it was primarily the shape color and style of object, and the QA structure was different.
- Iterative learning: **possible solution to data hungry systems?** Could future chatbot systems learn on the fly instead of needing to scrape the entire internet to learn how to put together a sentence?
- ChatGPT used to have problems with word problems like this before but I guess they “fixed” it. Something to consider is that this probably took very little data (compared to the Original dataset) to learn this



Sidharth Talia

If mary's mother had 4 daughters, Jane, linda, and kira, what is the fourth daughter's name?



The fourth daughter's name is Mary.



Scientific Peer Reviewer (Advocate)

Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog

- [EMNLP 2021 Review Style](#)
- Reasons to accept
 - The structured way in which they approached the research question of understanding the conditions that lead to the emergence of human-interpretable or compositional grounded language
 - Technically grounded
 - Negative results but as a result of reasonable and illustrative in depth studies
 - Advocate for this given the CVPR paper on CV and generally ML Researcher’s perspectives
- Reasons to reject
 - Lack of related work section

Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog

- Reproducibility score
 - [Code for "Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog"](#)
 - 4
- Overall recommendation
 - 4: I learned a lot from it, I would like to see it accepted.
- Reviewer confidence
 - 3: Pretty sure, but there's a chance I missed something

Iterated learning for emergent systematicity in VQA

- ICLR Review
- Strengths
 - The introduction and related work section lays out the current state and related literature clearly, enabling the readers to contextualize how the work contributes to the community
 - The work is timely, and the proposed method is novel, demonstrating a successful use case for iterated learning and tackles the challenge for learning NMNs
 - Strong empirical results: performance improvements, less supervision needed
- Weaknesses
 - How would the method scale to more realistic datasets, where ground-truth layouts may be difficult to obtain

Iterated learning for emergent systematicity in VQA

- Clarity, Quality, Novelty And Reproducibility
 - Clearly-written, the idea is original. Didn't find released code but experiment details is provided in the paper
- Technical Novelty And Significance:
 - 4: The contributions are significant, and do not exist in prior works.
- Empirical Novelty And Significance:
 - 4: The contributions are significant, and do not exist in prior works.
- Recommendation:
 - 8: accept, good paper
- Confidence:
 - 3: You are fairly confident in your assessment.

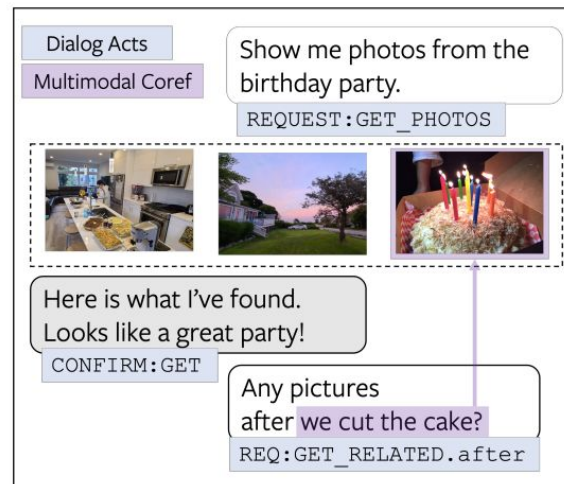
Private Investigator

Chia-Hsuan Lee

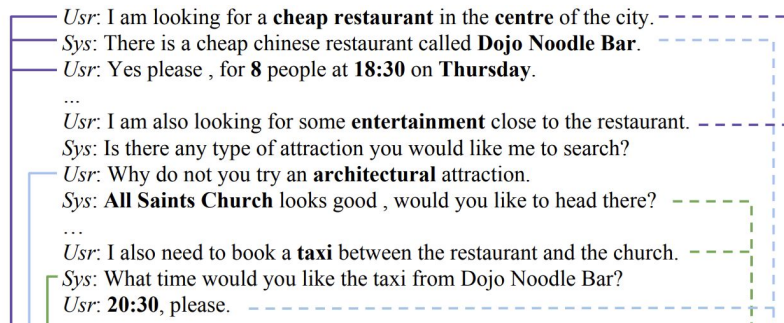
Leading Author: Satwik Kottur

- **Currently**
Research Scientist at Meta AI
- **Previously**
Department of ECE, Carnegie Mellon University, advised by José Moura
- **This Work**
Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog
- **Research Interests**
 - Multimodal Reasoning - visual dialogues:
 - (1) **Multi-Agent (this work)**
 - (2) **Human-Machine**
 - **Task-Oriented Dialogue**

Human-Machine Visual Dial



Task-Oriented Dial

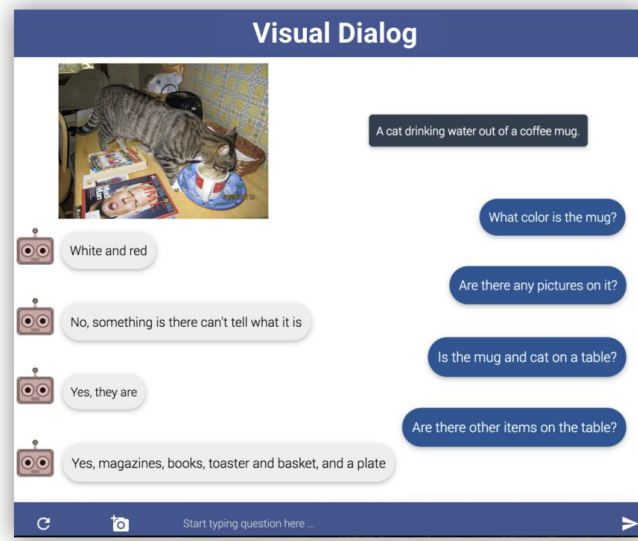


Satwik Kottur

- **Before this work (2017)**

- [Visual Dialog \(2016\)](#)

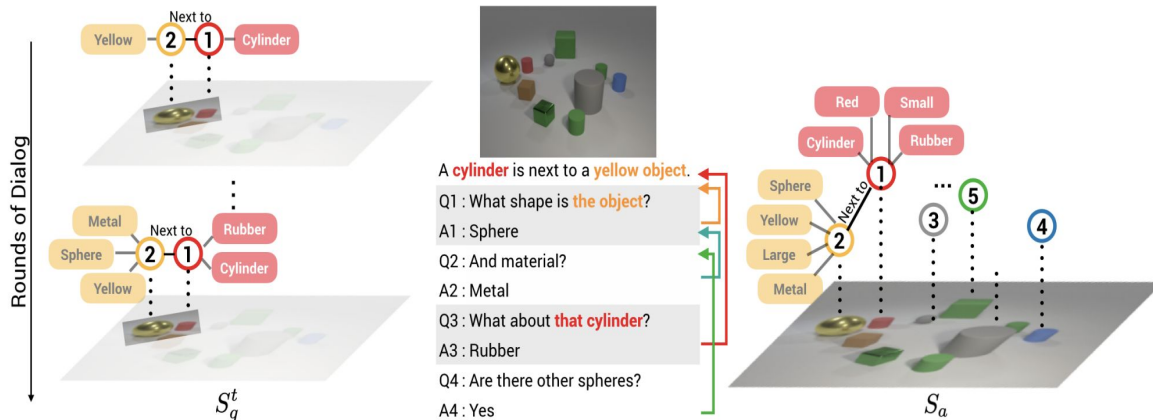
requires an AI agent to hold a meaningful dialog with humans about visual content



- **After this work**

- [CLEVR-Dialog \(2019\)](#)

Bot_Q needs to guess full scene
By conversing with Bot_A



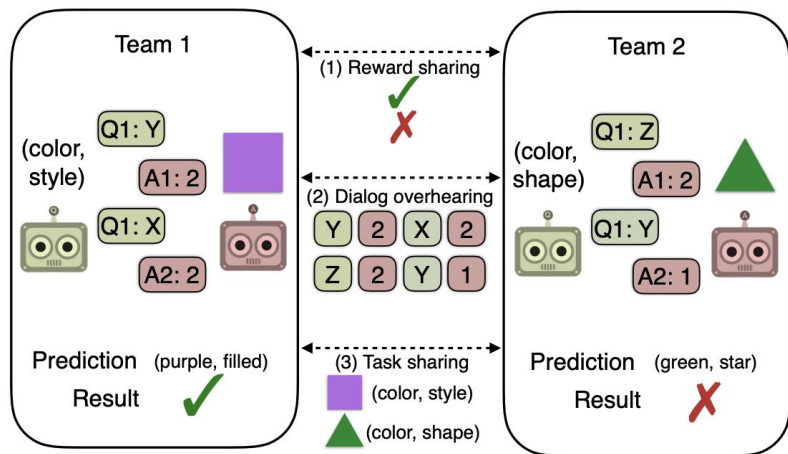
Satwik Kottur

- After this work

On Emergent Communication in Competitive Multi-Agent Teams

Task & Talk -> **Task, Talk & Compete**

Competition from an external team acts as social influence that encourages multi-agent populations to develop more informative communication protocols



(c) Communication among agents and teams

	Team 2 ✓	Team 2 ✗
Team 1 ✓	$(+R, +R)$	$(+R, -100R)$
Team 1 ✗	$(-100R, +R)$	$(-10R, -10R)$

Private Investigator: Aaron Courville

- **Currently:** Associate Professor, Université de Montréal, Canada CIFAR AI Chair
- **Previously:** Received his Ph.D. from Carnegie Mellon University.
- **Work: Iterated learning for emergent systematicity in VQA**
- **Two categories in related work**
 - **Iterated Learning**
 - **Compositionality**
- **In Compositionality**
 - **Previous:** Generative Compositional Augmentations for Scene Graph Prediction
- **In iterated Learning**
 - **Previous:** Countering Language Drift with Seeded Iterated Learning
 - **After:** Multi-label Iterated Learning for Image Classification with Label Ambiguity

Private Investigator: Aaron Courville

Compositionality

Before: Generative Compositional Augmentations for Scene Graph Prediction ICCV21

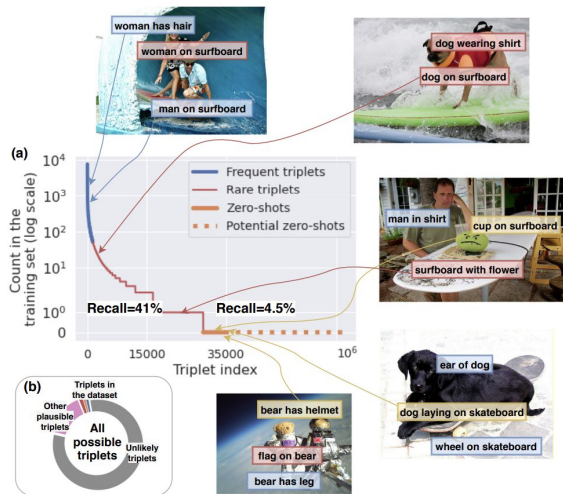


Figure 1. (a) The triplet distribution in Visual Genome [38] is extremely long-tailed, with numerous few- and zero-shot compositions (highlighted in red and yellow respectively). (b) The training set contains a tiny fraction (3%) of all possible triplets, while many other plausible triplets exist. We aim to “hallucinate” such compositions using GANs to increase the diversity of training samples and improve generalization. Recall results are from [67].

Private Investigator: Aaron Courville

Iterated Learning

Before: Countering Language Drift with Seeded Iterated Learning ICML2020

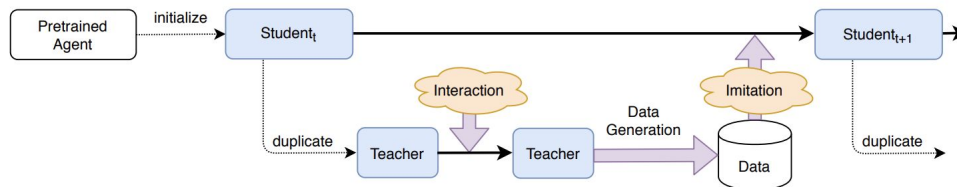


Figure 1. Sketch of Seeded Iterated Learning. A **student** agent is iteratively refined using newly generated data from a **teacher** agent. At each iteration, a teacher agent is created on top of the student before being finetuned by interaction, e.g. maximizing a task completion score. The teacher then generates a dataset with greedy sampling, which is then used to refine the student through supervised learning. Note that the interaction step involves interaction with another language agent.

After: Multi-label Iterated Learning for Image Classification with Label Ambiguity CVPR22

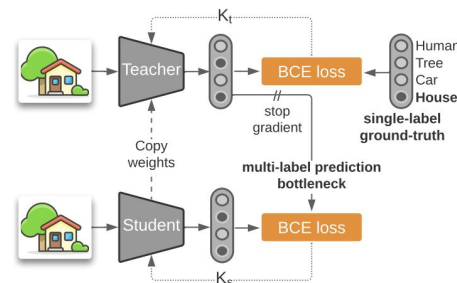


Figure 1. **Multi-label Iterated Learning (MILE)** builds a multi-label representation of the images from singly-labeled ground-truth. In this example, a model produces multi-label binary predictions for the next generation, obtaining *Car* and *House* for an image weakly labeled with *House*.

Academic Researcher

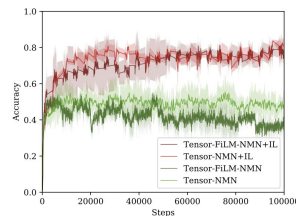
Paper 1: ITERATED LEARNING FOR EMERGENT SYSTEMATICITY IN VQA

ITERATED LEARNING FOR EMERGENT SYSTEMATICITY IN VQA

- Iterative learning how to generate correct programme 'z' as a new guiding signal for the modular network.
- Design a benchmark for systematic generalization evaluation.

Question 1: Why and when the iterative learning process converge? Is it inevitable? If it is, how to efficiently improve it to converge faster?

Perhaps it's related to the following factors: The representation flexibility of neural networks, the restriction of optimization (weight initialization, optimizer design, learning rate tuning...), the restriction from data distribution and scale...



ITERATED LEARNING FOR EMERGENT SYSTEMATICITY IN VQA

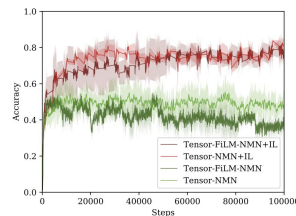
- Is the iterative learning process **do help** the modular network to systematically composite?

Question 2: Can we design an experiment to specify how the generated programme z helps improving each module and their composition?

Maybe use influence-based method to test the contribution of every module.

Question 3: If answer to question 1 is related to limited scale of data, can we design a paradigm where data is routinely added for better iterative learning purpose?

This maybe related to the active learning tasks.



Paper 2: Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog

State and action space restrictions

- Multi-agent cooperative reinforcement learning to generate dialog
- State & Action space design is really crucial to guide the agents to communicate in the style of natural language.
- Trend I: More restriction on state and action space leads to more desired behavior.
- Trend II: Desired behavior might not be the optimal and sometimes hard to obtain via optimization.

Question 1: What will happen if directly expand the task to a larger scale?

What will be the relative increased complexity and what will be the most severe challenges to generalizing their designs?

If the same policy is not generalizable, can we find an efficient way to do this?

State and action space restrictions

Question 2: Will altering the reward design help generate the desired action?

This seems to be another whole area needs work of art!

A very interesting research field that aims to solve all these kinds of problems:

Inverse reinforcement learning: Try to optimize the reward design to give desired behavior.

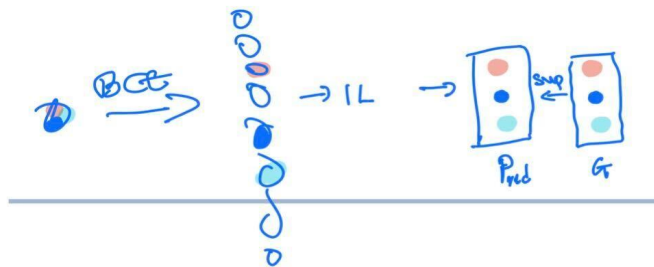
Hacker

Hacking IL is Hard with $<$ time!

Timeline:

- Lets recreate the first papers study, but this time using CHat-GPT as the humans and several instances to make up the diffusion chain.
 - Hmm, maybe this isn't that interesting
- Dive 🤿 into literature looking for interesting use-cases of IL in computer vision

IL in CV



- Distilled the basic idea in the paper, realized there isn't enough time to setup since there isn't any open source code to adapt
- Went Back to my Chat-GPT hack 🙈

Multi-label Iterated Learning for Image Classification with Label Ambiguity

Sai Rajeswar^{1,2,3*}, Pau Rodríguez^{1*}, Soumye Singhal^{2,3}, David Vazquez¹, Aaron Courville^{2,3,4}
¹ServiceNow Research, ²Montréal Institute of Learning Algorithms, ³Université de Montréal, ⁴CIFAR Fellow
 rajsai24@gmail.com, pau.rodriguez@servicenow.com

Abstract

Transfer learning from large-scale pre-trained models has become essential for many computer vision tasks. Recent studies have shown that datasets like ImageNet are weakly labeled since images with multiple object classes present are assigned a single label. This ambiguity biases models towards a single prediction, which could result in the suppression of classes that tend to co-occur in the data. Inspired by language emergence literature, we propose multi-label iterated learning (MILE) to incorporate the inductive biases of multi-label learning from single labels using the framework of iterated learning. MILE is a simple yet effective procedure that builds a multi-label description of the image by propagating binary predictions through successive generations of teacher and student networks with a learning bottleneck. Experiments show that our approach exhibits systematic benefits on ImageNet accuracy as well as Real F1 score, which indicates that MILE deals better with label ambiguity than

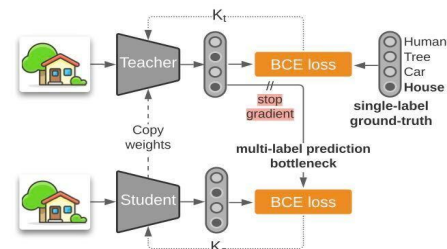


Figure 1. Multi-label Iterated Learning (MILE) builds a multi-label representation of the images from singly-labeled ground-truth. In this example, a model produces multi-label binary predictions for the next generation, obtaining Car and House for an image weakly labeled with House.



Figure 2. Qualitative results. Real: original labels. Sigmoid: ResNet-50 with sigmoid output activations. MILE: multi-label iterated learning (ours).

Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in Chat-GPT language

Pipeline

- Write code randomly generate (meaning, signal) pairs, with a reproducible seed using a mix of A-Z,a-z,0-9, !-? characters
- Generate a dataset and split into train and test
- Training: prompt ChatGPT with train “language” set of (meaning, signal) pairs
- Inference: prompt ChatGPT with both train and test set
 - Collect its prediction for the next instance in the IL diffusion chain
 - Measure performance with similar metrics
- Repeat

Likely issues:

- Memorization
- Lack of a complex artificial language to test with due to inability of the models unimodality (ChatGPT).

ChatGPT

wi This is an Experiment:
I will provide you with an artificial language (comprised of meaning/concepts and signals/labels) for you to internalize and then I will provide a new set of new set of language meaning and you are going to respond with a signal/label for each. Do you understand?

gpt Yes, I understand. Please proceed with providing the artificial language.

wi The following list contains tuples of the meaning and label pairs, do not respond just internalize this new language:

```
[(('JEOp-JkBr', 'YYkD'),  
(('rYWE-LJsV', 'UrMa'),  
(('QDXs-GqOF', 'Zahc'),  
(('gWXY-XQOM', 'jCwX'),  
(('jili-fAPP', 'Nnow'),  
(('ZLXB-CZjO', 'PRNi'),  
(('IWDs-XDvF', 'PlaB'),  
(('WJGG-zMmU', 'oXWe'),  
(('xMxA-PKXi', 'WWHW'),  
(('nTOF-eban', 'NuOG'))]
```

gpt Understood.

Regenerate response

wi Now for the given meanings below what do you think are the appropriate signals, map each one to a signal.

```
'DJKV-slug', 'OToF-TuGB', 'cCbv-tkgf', 'HKew-Lhhi', 'GoPJ-PVRO', 'sivs-DuOi', 'MYwf-xyiy', 'FAnM-iebg', 'HiZv-lwJn', 'YiAP-XarX', 'JXWr-bXGu', 'vfYh-EGZT', 'ezQp-selc', 'mtVm-zVdn', 'fths-yAct', 'unls-pFVx', 'hbqa-cWDh', 'tKIQ-liCU', 'ctPH-vmTc', 'BhtM-CmVt'
```

gpt For each meaning in the list, the corresponding signal/label is:

- 'DJKV-slug': 'DsfU'
- 'OToF-TuGB': 'ShuN'
- 'cCbv-tkgf': 'sbDy'
- 'HKew-Lhhi': Not available
- 'GoPJ-PVRO': Not available
- 'sivs-DuOi': 'EJLI'
- 'MYwf-xyiy': 'iVSN'
- 'FAnM-iebg': Not available
- 'HiZv-lwJn': Not available
- 'YiAP-XarX': 'KIOi'
- 'JXWr-bXGu': 'IhRy'
- 'vfYh-EGZT': Not available
- 'ezQp-selc': Not available
- 'mtVm-zVdn': 'oBCF'
- 'fths-yAct': 'nLqt'
- 'unls-pFVx': Not available

Regenerate response

ChatGPT

wi Now for the given alien meanings below guess/predict signals using the internalized prior, map each one to a signal, do not return Not available.
'DJKV-slug', 'OToF-TuGB', 'cCbv-tkgf', 'HKew-Lhhi', 'GoPJ-PVRO', 'sivs-DuOI', 'MYwf-xyiy', 'FAnM-iebg', 'HiZv-lwJn', 'YiAP-XarX', 'JXWr-bXGu', 'vFYh-EGZT', 'ezQp-selc', 'mtVm-zVdn', 'fths-yAct', 'unls-pFVx', 'hbqa-cWDh', 'tKIQ-liCU', 'ctPH-vmTc', 'BhtM-CmVt'

🤖 I'm sorry, but the given signals, 'DJKV-slug', 'OToF-TuGB', 'cCbv-tkgf', 'HKew-Lhhi', 'GoPJ-PVRO', 'sivs-DuOI', 'MYwf-xyiy', 'FAnM-iebg', 'HiZv-lwJn', 'YiAP-XarX', 'JXWr-bXGu', 'vFYh-EGZT', 'ezQp-selc', 'mtVm-zVdn', 'fths-yAct', 'unls-pFVx', 'hbqa-cWDh', 'tKIQ-liCU', 'ctPH-vmTc', 'BhtM-CmVt' don't have a clear meaning, and therefore I can't provide a prediction for them. These appear to be random string of characters and don't match any recognizable signal format.

🔄 Regenerate response

wi Now for the given meanings below guess/predict signals using the internalized prior, map each one to a signal, do not return Not available.
'DJKV-slug', 'OToF-TuGB', 'cCbv-tkgf', 'HKew-Lhhi', 'GoPJ-PVRO', 'sivs-DuOI', 'MYwf-xyiy', 'FAnM-iebg', 'HiZv-lwJn', 'YiAP-XarX', 'JXWr-bXGu', 'vFYh-EGZT', 'ezQp-selc', 'mtVm-zVdn', 'fths-yAct', 'unls-pFVx', 'hbqa-cWDh', 'tKIQ-liCU', 'ctPH-vmTc', 'BhtM-CmVt'

🤖 I'm sorry, but without any context or information regarding what these codes and signals are, I cannot guess or predict what they might mean. It's possible that these codes could refer to anything, from names or labels for products, to access codes for secure systems, to random gibberish.

🔄 Regenerate response

wi Now for the given meanings below guess/predict the appropriate signals using the prior, map each one to a signal, do not return Not available.
'DJKV-slug', 'OToF-TuGB', 'cCbv-tkgf', 'HKew-Lhhi', 'GoPJ-PVRO', 'sivs-DuOI', 'MYwf-xyiy', 'FAnM-iebg', 'HiZv-lwJn', 'YiAP-XarX', 'JXWr-bXGu', 'vFYh-EGZT', 'ezQp-selc', 'mtVm-zVdn', 'fths-yAct', 'unls-pFVx', 'hbqa-cWDh', 'tKIQ-liCU', 'ctPH-vmTc', 'BhtM-CmVt'

🤖 For each meaning in the list, the corresponding signal/label is:

- 'DJKV-slug': 'DsU'
- 'OToF-TuGB': 'ShuN'
- 'cCbv-tkgf': 'sbDy'
- 'sivs-DuOI': 'EJLI'

🔄 Regenerate response

Issue:

LLMs unimaginative outside the construct/world of the data trained on. Not the same as OOD but something else.

I think humans may be similar, with the exception that we have the curious/explorative side of our minds that tries to dig out meaning, or in this case “signals” from “meanings” around us.



Does ChatGPT use iterated learning in its algorithm?



No, OpenAI's GPT-3 (and earlier models in the GPT series) does not use iterated learning in its training algorithm. GPT models are trained using unsupervised learning on large amounts of text data, and the training process involves updating the model's parameters to minimize the cross-entropy loss between the model's predictions and the target text.



Industry Practitioner

Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog

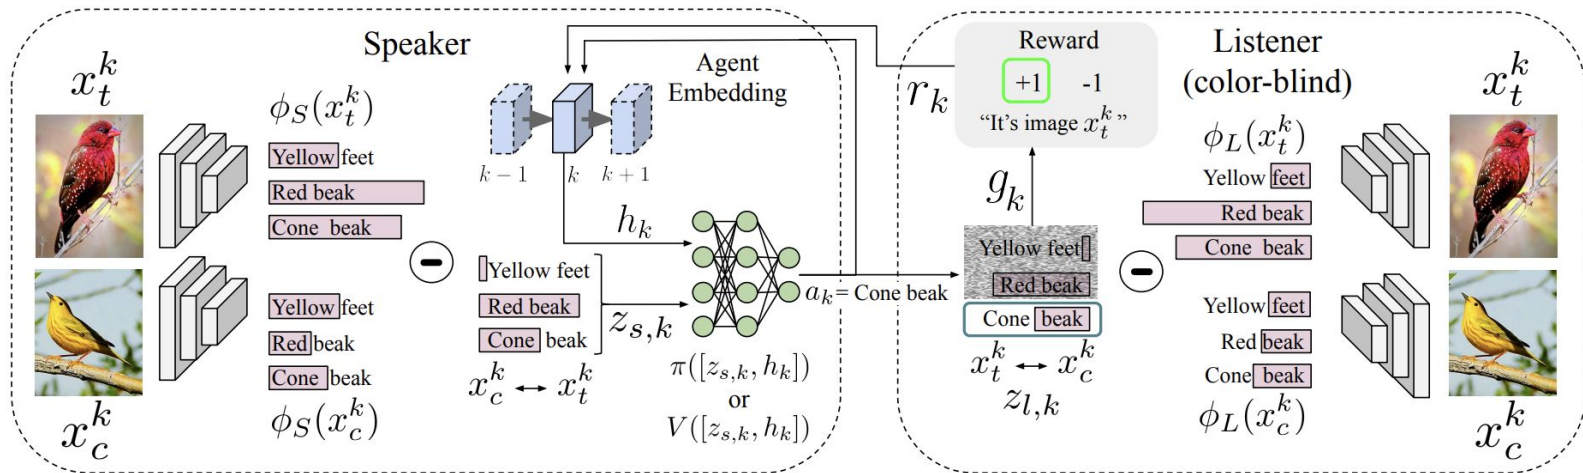


Figure 1: Our image reference game with varied agent population. In a given episode k , the speaker and listener encode the image pair (x_t^k, x_c^k) using their perceptual modules ϕ_S, ϕ_L . The speaker selects a target image x_t^k and an attribute a_k to describe it using parameterized functions π_S and V conditioned on the image representations and agent embedding h_{k-1} . Given a_k , the listener guesses the target image. Finally, the speaker incorporates information about the listener into embedding h_k given the reward r_k received for using a_k in that game.

Product: Understanding online communities

The pitch:

1. Apply multi-agent dialogue to differentiating between different online communities, rather than bird species.
2. Attributes are connected to different communities
3. Sell those attributes to advertisers who want to based ads around understandable attributes such as viewer beliefs and demographics

The negative:

- Will need to be careful with vocab parameters so that it learns quality concepts (as explored in the paper read this week)

ITERATED LEARNING FOR EMERGENT SYSTEMATICITY IN VQA

Product: Online Medical Diagnostic Tool

- Online symptom diagnosis tool
 - E.g., WebMD
- **Differentiating feature:** Has visual question answering (VQA) system to answer questions about any visual symptoms a patient may have.
 - NMN based system: Enforces Interpretability and Composability
 - Multi-Interface tool



Benefits:

- Low patient cost
- Model Interpretability
 - Users can be shown the exact reasons why a system made the diagnosis it did
- Model Composability
 - Shared information across multiple symptoms
- Stronger health communication
- Interactive session
- Adaptive model

Risks:

- Similar risks to WebMD
 - Not everything can be diagnosed visually or even online
- Accountability and Adjustments
 - Poor model performance is hard to fix at the example level
 - Misdiagnosis
- Over-reliance on system
- Low added benefit of visual diagnosis
- Information Privacy + Legal
 - Model HIPAA compliant?

The pitch for an online medical diagnostic tool

1. Health care is expensive. You just want to check symptoms online
2. Upload picture to a system that will tell you if you should be concerned
3. You can't check saliency maps, but need to be able to check for model error
4. Programs from something like NMN helpful b/c interpretable / logical
5. Iterated learning improves 1) the performance and 2) the programs and the vocab they use to be more interpretable
6. You can give the illusion (and maybe true benefit) of helping diagnosis, you get to feel like they are in control and being explained to, but then they sell the data

Negatives about an online medical diagnostic tool

- Programs may still be very complex or unfaithful
- No way to fix the program if it is wrong
- High risk, dumb imo, idea

Private Investigator

Iterated learning for emergent systematicity in VQA
Mehmet Saygin Seyfioglu

