

COMMAI: EVALUATING THE FIRST STEPS TOWARDS A USEFUL GENERAL AI

Marco Baroni, Armand Joulin, Allan Jabri, Germán Kruszewski, Angeliki Lazaridou,*
Klemen Simonic & Tomas Mikolov

Facebook Artificial Intelligence Research

{mbaroni, ajoulin, ajabri, germank, angelikil, klemen, tmikolov}@fb.com

ABSTRACT

With machine learning successfully applied to new daunting problems almost every day, general AI starts looking like an attainable goal (LeCun et al., 2015). However, most current research focuses instead on very specific applications, such as image classification or machine translation. We believe this to be largely due to the lack of objective ways to measure progress towards broad machine intelligence. In order to fill this gap, we propose here a set of concrete desiderata for general AI, together with a platform to test machines on how well they satisfy such desiderata, while keeping all further complexities to a minimum.

1 DESIDERATA FOR THE EVALUATION OF MACHINE INTELLIGENCE

Rather than trying to define intelligence in abstract terms, we take a pragmatic approach: we would like to develop AIs that are *useful* for us. This naturally leads to the following desiderata.

Communication through natural language An AI will be useful to us only if we are able to communicate with it: assigning it tasks, understanding the information it returns, and teaching it new skills. Since natural language is by far the easiest way for us to communicate, we require our useful AI to be endowed with basic linguistic abilities. The language the machine is exposed to in the testing environment will inevitably be very limited. However, given that we want the machine to also be a powerful, fast learner (see next point), humans should later be able to teach it more sophisticated language skills as they become important to instruct the machine in new domains. In concrete, the environment should not only expose the machine to a set of tasks, but provide instructions and feedback about the tasks in simple natural language. The machine should rely on this form of linguistic interaction to efficiently solve the tasks.

Learning to learn A useful AI should be flexible. As our needs change, the AI should help us with the new challenges we face: from solving a scientific problem in the morning at work to stocking our fridge at night. Progress towards AI should thus be measured on the ability to master a continuous flow of new tasks, with data-efficiency in solving new tasks as a fundamental evaluation component, and without distinguishing train and test phases. We must distinguish this *learning to learn* ability, pertaining to generalization across tasks (Ring, 1997; Schmidhuber, 2015; Silver et al., 2013; Thrun & Pratt, 1997), from 1-shot learning, that is, the challenging but more limited ability to generalize to new classes within the same task (e.g., extending an object classifier to recognize unseen objects from just a few examples; Lake et al., 2015). It's generally agreed that, in order to generalize across tasks, a program should be capable of *compositional* learning, that is, of storing and re-combining solutions to sub-problems across tasks (Fodor & Lepore, 2002; Lake et al., 2016; Minsky, 1986). The testing environment should thus feature sets of *related* tasks, such that a compositional learner can bootstrap skills from one task to the other. Finally, mastering language skills might be a crucial component of learning to learn, since understanding linguistic instructions allows us to quickly learn how to accomplish tasks we have never performed before.

Feedback As we grow up, we learn to master complex tasks with decreasing amounts of explicit reward. A useful AI should possess similar capabilities. Consequently, in our testing environment,

*Now at DeepMind.

reward should decrease with time. Conversely, the machine should be able to learn from performance cues that are not directly linked to an explicit reward score, such as purely linguistic feedback (see also Weston, 2016), or observing other agents that are correctly performing a task (as in learning by demonstration, Argall et al., 2009). The testing environment should include such cues.

Interface The interface between the machine and the world should be maximally general. The machine itself should learn the best way to process different kinds of input and output streams, with no need for manual re-programming as we apply it to different domains. We thus assume the simplest possible interface. At each time step, the machine receives one bit and sends one bit, without any further structure imposed on the bit stream (a separate channel is used for reward in the initial stages of the simulation).

We do not claim that satisfying our desiderata will lead to a full-fledged intelligent machine, but we see them as prerequisites to be able to efficiently teach it more advanced skills.

2 THE COMMAI FRAMEWORK

We call the evaluation framework satisfying the desiderata above **CommAI** (communication-based AI), given the prominence we give to communication skills. We have developed the open source **CommAI-env** platform¹ to implement sets of CommAI tasks. As a concrete example of a set of simple tasks already satisfying many of the requirements above, we briefly present here the **CommAI-mini** tasks (described in more detail in the Supplementary Materials).

In a CommAI-mini task, the environment presents a (simplified) regular expression to the learner. It then asks it to either recognize or produce a string matching the expression. The environment listens to the learner response and it provides linguistic feedback on the learner’s performance (possibly assigning reward). All exchanges take place at the bit level. Some examples follow:

<i>Environment:</i>	<i>Learner:</i>
description: C or D; verify: CCCC.	
wrong; correct: true.	fkljfd.
	...
description: HL and RM and BT; verify: RMBTBTHLHLBT.	
right. (+1 reward)	true.
	...
description: not AB; verify: ADFCFHGHADDDDB.	
	...
description: C; produce.	
wrong; example correct: CCC.	ABCD.
	...
description: C; produce two distinct strings.	
	...

Learning-to-learn is a must, since the learner will rarely, if ever, be tested on exactly the same target grammar, grammars become more complex with time, and the learner is asked to use the same grammar in different ways (e.g., for recognition or production). Compositionality plays an important role at multiple levels: (i) Skills such as chunking bit sequences into characters and parsing messages into predictable parts (the description, the test string, the delimiters, etc.) will greatly help the learner to generalize across tasks. (ii) Succeeding at the recognition tasks should help in solving the equivalent production tasks (and *vice versa*). (iii) We control moreover the complexity of the stringsets and their descriptions, by incrementally adding operators to the regular expressions.

¹<https://github.com/facebookresearch/CommAI-env/>

For example, checking if a string contains all n-grams in a set requires checking the presence of the n-grams in the string, so a compositional learner should be faster at a task involving n-gram conjunction after it has solved the task disjunction tasks.

The tasks are different from standard artificial grammar learning (Reber, 1967), in that the learner is given *explicit instructions* in simplified English about the target stringset and what to do with it (*description*., *verify*., *produce*...), as well as verbal feedback on its performance. We thus satisfy our linguistic communication desideratum.

Importantly, although the CommAI-mini tasks are fully “linguistic”, in the sense that they pertain to character string recognition and production, we chose the regular grammar domain just because it’s simple and well-understood (incidentally, one could also think of the test strings as denoting non-verbal acoustic or visual stimulus sequences). What satisfies the language desideratum is not the nature of the tasks, but the fact that the environment provides meta-information about them (instructions, feedback) in simplified English. Other CommAI task sets could, for example, be based on simple physics tasks, where sensory information would be passed through the bit-based channel, together with instructions and feedback that would *still* be expressed in simplified English (e.g., *move the red block over the blue block*; see Andreas et al., 2016 for somewhat related ideas).

Despite their simplicity, we conjecture that solving the CommAI-mini tasks without astronomical amounts of training examples is out of the scope of current machine learning methods (more advanced task examples can be found on the CommAI-env site). We hope the CommAI-mini challenge is at the right level of complexity to stimulate researchers to develop genuinely new models.

3 RELATED WORK

We can identify two broad approaches to benchmarking general AI. Some researchers, like us, take a top-down view, deriving their requirements from psychological or mathematical considerations (for example, Adams et al., 2012, Lake et al., 2016, and see the extensive review in Hernández-Orallo, 2017). We sympathize with this principled approach, but we are not aware of others having emphasized the same set of practical desiderata that we outlined above, nor proposing a concrete framework for evaluation like we do with CommAI-env.

Others focus on existing applications that are considered of sufficient complexity to measure progress towards general-purpose intelligence. For example, games such as Go (Silver et al., 2016) and StarCraft (Ontañón et al., 2013) require sophisticated planning skills intuitively associated with intelligence. While current results in these domains are impressive, we think this approach is at the same time too simple and too complex as a general AI benchmark. On the one hand, the focus shifts from domain-independent skills to more limited game-specific strategies. On the other, raw input pre-processing and adapting to game-specific dynamics might require heavy computational resources and advanced domain-specific *know-how*, with high entry cost for researchers that are not already working in the target domains. These issues are partially addressed by platforms that provide a unified interface to multiple games and other programs.² However, simply pooling a large number of existing applications will make for a ragtag collection of benchmarks, with no clear unified goal in terms of evaluating general intelligence.

The bAbI tasks (Weston et al., 2015) are superficially similar to CommAI tasks, but they evaluate general text understanding phenomena, rather than compositional learning-to-learn abilities.

ACKNOWLEDGMENTS

This abstract summarizes and refines ideas we originally presented in an unpublished manuscript (Mikolov et al., 2016). We thank Gemma Boleda, Stan Dehaene, Emmanuel Dupoux, Jan Feyereisl, Amaç Herdağdelen, José Hernández-Orallo, Iasonas Kokkinos, Martin Poliak, Marek Rosa, our FAIR colleagues and the participants of the MAIN@NIPS 2016 workshop for feedback.

²E.g.: <https://openai.com/blog/universe/>, <https://github.com/deepmind/lab>, <http://www.ggp.org>, <http://www.gvgai.net/>

REFERENCES

- Sam Adams, Itamar Arel, Joscha Bach, Robert Coop, Rod Furlan, Ben Goertzel, Storrs Hall, Alexei Samsonovich, Matthias Scheutz, Matthew Schlesinger, Stuart Shapiro, and John Sowa. Mapping the landscape of human-level artificial general intelligence. *AI Magazine*, 33(1):25–41, 2012.
- Jacob Andreas, Dan Klein, and Sergey Levine. Modular multitask reinforcement learning with policy sketches. <https://arxiv.org/abs/1611.01796>, 2016.
- Brenna Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- Jerry Fodor and Ernest Lepore. *The Compositionality Papers*. Oxford University Press, Oxford, UK, 2002.
- José Hernández-Orallo. *The Measure of All Minds*. Cambridge University Press, Cambridge, UK, 2017.
- Gerhard Jäger and James Rogers. Formal language theory: refining the Chomsky hierarchy. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 367(1598):1956–1970, 2012.
- Brenden Lake, Ruslan Salakhutdinov, and Joshua Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Brenden Lake, Tomer Ullman, Joshua Tenenbaum, and Samuel Gershman. Building machines that learn and think like people. <https://arxiv.org/abs/1604.00289>, 2016.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436–444, 2015.
- Robert McNaughton and Seymour Papert. *Counter-Free Automata*. MIT Press, Cambridge, MA, 1971.
- Tomas Mikolov, Armand Joulin, and Marco Baroni. A roadmpap towards machine intelligence. <http://arxiv.org/abs/1511.08130/>, 2016.
- Marvin Minsky. *The Society of Mind*. Simon & Schuster, New York, 1986.
- Santiago Ontañón, Gabriel Synnaeve, Alberto Uriarte, Florian Richoux, David Churchill, and Mike Preuss. A survey of real-time strategy game AI research and competition in StarCraft. *IEEE Transactions on Computational Intelligence and AI in Games*, 5(4):293–311, 2013.
- Arthur Reber. Implicit learning of artificial grammars. *Verbal Learning and Verbal Behavior*, 5(6): 855–863, 1967.
- Mark Ring. CHILD: A first step towards continual learning. *Machine Learning*, 28:77–104, 1997.
- James Rogers, Jeffrey Heinz, Margaret Fero, Jeremy Hurst, Dakotah Lambert, and Sean Wibel. Cognitive and sub-regular complexity. In Glyn Morrill and Mark-Jan Nederhof (eds.), *Formal Grammar: 17th and 18th International Conferences*, pp. 90–108. Springer, Berlin, Germany, 2013.
- Jürgen Schmidhuber. On learning to think: Algorithmic information theory for novel combinations of reinforcement learning controllers and recurrent neural world models. <http://arxiv.org/abs/1511.09249>, 2015.
- Daniel Silver, Qiang Yang, and Lianghao Li. Lifelong machine learning systems: Beyond learning algorithms. In *Proceedings of the AAAI Spring Symposium on Lifelong Machine Learning*, pp. 49–55, Stanford, CA, 2013.
- David Silver, Aja Huang, Christopher Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529:484–503, 2016.

Sebastian Thrun and Lorien Pratt (eds.). *Learning to Learn*. Kluwer, Dordrecht, 1997.

Jason Weston. Dialog-based language learning. In *Proceedings of NIPS*, pp. 829–837, Barcelona, Spain, 2016.

Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. Towards AI-complete question answering: A set of prerequisite toy tasks. <http://arxiv.org/abs/1502.05698>, 2015.

SUPPLEMENTARY MATERIALS: THE COMMAI-MINI TASKS

The CommAI-mini tasks are based on the hierarchy of sub-regular languages, in turn a subset of the regular languages (Jäger & Rogers, 2012; McNaughton & Papert, 1971). Sub-regular languages are useful to characterize pattern recognition skills of humans and other animals (for example, constraints on the distribution of stressed syllables in the world languages, Rogers et al., 2013).

Strictly local languages, the simplest class of sub-regular languages, can be recognized with an n-gram lookup table only. For example, the stringset accepted by the regular expression $(AB)^+$ is strictly local, because a lookup table containing the bigram AB is sufficient to recognize the strings in it (we’re ignoring here technicalities regarding begin- and end-of-string conditions, and the low-level implementation of the actual “scanner”). The $(AB|C)^+$ language is also strictly local, because it can be recognized through a lookup table containing the n-grams AB and C .

The next class ascending the hierarchy is that of locally testable languages. The latter can be recognized by imposing logical constraints (union, conjunction, complement) on the n-grams that occur, or do not occur, in a string. For example, $A^*(BA^*)^+$, the “*at-least-one-B*” language, can be recognized by using a lookup table containing the unigrams A and B , plus a checking device that verifies that the B unigram occurred at least once. The strictly local languages are a strict subset of the locally testable languages.

Locally testable languages are not the most complex kind of sub-regular languages, and they are still far from exploiting the full expressive power of regular languages (that are in turn the simplest class in the Chomsky hierarchy). Yet, by combining (subsets of) strictly local and locally testable languages, we already obtain an interesting challenge for CommAI learners. Importantly, an efficient solution to the CommAI-mini tasks does not just involve stringset recognition/production, but learning the description language that specifies the rules about legal strings.

All CommAI-mini tasks have the same structure (where communication flows through the bit-level interface):

1. The environment presents the description of a target stringset;
2. the environment tells the learner whether it is a recognition or a production task;
3. if it is a recognition task, the environment produces the string to be recognized;
4. the environment listens to the learner, and records the string produced by the latter until a period occurs, or a maximum number of bits has been emitted;
5. the environment checks the string produced by the learner;
6. if the string is correct, the environment issues reward and states that the answer is correct;
7. if the string is wrong
 - in a recognition task, the environment states that the answer is wrong, and produces the right answer (`true` or `false`);
 - in a production task, the environment states that the answer is wrong, and produces a sample correct string.

The tasks are organized into task sets, with each set constituting a videogame-like “level”. Tasks in the same set are presented in random order. Each recognition-based set below could also be seen as a single recognition task for the relevant class of stringsets (more generally, we could think of *all* recognition tasks as a single task). We prefer the granular structure we are outlining below, because it will facilitate analysis. For example, learning strictly local unigram languages (A or B or C) is a special case of learning strictly local maximally-5-gram languages ($ANFJG$ or CED or $KPQR$ or ZM or S). However, treating these as separate tasks should make it easier to check if a learner has memory limitations, such that it doesn’t scale up to n-grams beyond a certain length.

Each task is defined by the structure of the description (maximum n-gram length, number of terms, permitted operators), but the actual symbols defining an acceptable stringset will change from exposure to exposure. For example, the second task in set #1 below consists in recognizing any $(XY)^+$ string, where X and Y are arbitrary upper-case letters: `description AB` and `description LK` are two different instances of this task.

In what follows, the tasks are illustrated by the string produced by the environment at the beginning of a task instance (corresponding to the first 3 steps in the enumeration above). As we just remarked, the target language (the actual stringset) will change from instance to instance of the same task. We will moreover only show a few illustrative tasks for each set. Further tasks can be generated by varying the maximum n-gram size and, except in set #1, the number of n-gram terms present in the description.

TASK SET #1

The following examples illustrate set #1 (here and below, we only show positive examples, where the learner should answer `true`):

description: C; verify: CCCC.

description: AB; verify: ABAB.

description: FJG; verify: FJG.

Tasks in set #1 involve strictly local descriptions. There is a natural hierarchy within the set in terms of the length of the n-grams that must be memorized: verifying the (AB) + language requires less memory than verifying the (FJG) + language.

TASK SET #2

Examples:

description: anything; verify: ANFHG.

description: AB or CD; verify: ABAB.

description: FAB or GH or MIL; verify FABFAB.

The tasks in set #2 are also based on strictly local descriptions. However, because of the `or` operator, solving them requires storing multiple n-grams in memory. The #2 tasks thus imply the abilities necessary to solve #1 tasks (storing n-grams and checking their presence in a string), but they generalize them (to storing and using multiple n-grams).

The #2 tasks vary in terms of the number of disjoint n-grams that comprise the description and the maximum length of the n-grams in the description.

We introduce `anything` as a special symbol matching any (byte-level) sequence. Recognizing `anything` is strictly local.

TASK SET #3

Examples:

description: AB and CF; verify: ABCFABAB.

description: HL and RM and BT; verify: RMBTBTHLHLBT.

description: AB and anything; verify: FKGABJJKJKSD.

description: AB and CF and anything; verify: FJGKJKJKJKJDCFDJKJKJKSJAB.

Set #3 tasks involve locally testable languages. Verifying that the target string only contains n-grams from the description no longer suffices. The learner must check whether *all* n-grams in the description have been used. These tasks thus generalize #2 tasks. They also require storing multiple n-grams in memory, but they further need some device to check that all the n-grams in the lookup table have been used.

There is again an obvious hierarchy in terms of how many distinct n-grams must be stored in memory and their length. Tasks can also be distinguished in terms of whether they include the `anything` operator or not.

We are not considering tasks mixing conjunction and disjunction, except for the implicit `anything`-denoted disjunction. We exclude the more general case to avoid having to implement complex scope conventions.

TASK SET #4

Examples:

description: not AB and anything; verify: ADFCFHGHADDDDB.

description: not AB and CF and anything; verify: DJFKJKJSCFDSFG.

description: not AB and not CF and anything; verify: DJFKJKJSCEFDSFG.

Tasks in set #4 also involve locally testable languages. However, on top of conjunction, they include a negation operator. In our setup, conjunction always takes scope over negation, to avoid the need for overt bracketing in the descriptions. The fact that a negated n-gram is equivalent to the affirmation of its complement is explicitly expressed by always adding the `and anything` condition. For the time being, we do not consider more general combinations of conjunction, disjunction and negation.

TASK SET #5

Examples:

description: C; produce.

description: AB; produce.

description: FJG; produce.

description: anything; produce.

description: AB or CD; produce.

description: FAB or GH or MIL; produce.

description: AB and CF; produce.

description: HL and RM and BT; produce.

description: AB and anything; produce.

description: AB and CF and anything; produce.

description: not AB and anything; produce.

description: not AB and CF and anything; produce.

description: not AB and not CF and anything; produce.

We consider the production counterparts of all recognition tasks. The learner is asked to generate one string matching the conditions in the description. We expect a compositional learner to solve the production tasks much faster if it has already been exposed to the recognition tasks (and *vice versa*).

FURTHER TASKS

The production tasks in set #5 can be solved by always generating the shortest string in the description. For the simpler tasks (without conjunction), this amounts to producing the first upper-case string in the description. We can force the learner out of this strategy by asking it to produce *two* distinct strings matching the description, e.g.:

```
description: C; produce two distinct strings.
```

Tasks of this sort would obviously build on skills acquired in set #5, adding the requirement that the learner stores its own past productions in memory, and uses them when planning what to produce next.

It's easy to think of further tasks that a learner could solve fast by exploiting skills acquired through sets #1-5, e.g., switching the capitalization conventions:

```
DESCRIPTION: ab; PRODUCE.
```

More ambitiously, the learner could be provided with a sample of strings, and asked to formulate a description accepting them (even producing extremely loose descriptions, such as `anything`, would constitute an impressive achievement).