
Reinforcement Learning of Theorem Proving

Josef Urban <josef.urban@gmail.com>
To: Stephen Ornes <stephen.ornes@gmail.com>

Thu, Aug 20, 2020 at 10:25 AM

Dear Stephen,

I have found more comments that I sent to another journalist last year - see below.

I understand that there is a push by some groups (paid by big US companies) to claim that deep learning is the next big thing in AR and to promote their work - e.g. <https://www.newscientist.com/article/2200707-google-has-created-a-maths-ai-that-has-already-proved-1200-theorems/> .

If you check the facts below, you will see that the largest recent breakthroughs were done so far without deep learning and that claiming DL-based breakthroughs is quite premature.

This is nothing against DL, which is interesting and we have done many experiments with it, rigorously comparing it to other ML methods. But facts are facts and it would be bad if science itself was steered by hype and fake news instead of rigorous evaluation.

Let me perhaps also note that AR is one of the pinnacles of science in the sense that it is about establishing truth through a rigorous process called "logic" and "machine-verified proof". You may want to read John McCarthy's vision about it: <http://www-formal.stanford.edu/jmc/future/objectivity.html> and my short vision from 2013: https://people.ciirc.cvut.cz/~urbanjo3/arw2013_submission_21.pdf .

As an AR/AI researcher inspired (among others) by John McCarthy, I very much hope that one day we will be able to use AR methods to automatically debunk hyped, biased and false claims made by influential politicians and companies. One of the immediate benefits would be to prevent disproportionate hype and spread of fake news in science itself.

Best,
Josef

----- Forwarded message -----

From: **Josef Urban** <josef.urban@gmail.com>
Date: Wed, Jul 10, 2019, 09:20
Subject: Re: Reinforcement Learning of Theorem Proving
To: Matthew Hutson <matt@silverjacket.com>
Cc: Cezary Kaliszyk <cezary.kaliszyk@uibk.ac.at>

Hi Matthew,

On Wed, Jul 10, 2019, 07:05 Matthew Hutson <matt@silverjacket.com> wrote:
Hi Cezary and Josef,

I'm a freelance science reporter who covers AI, and I'm working on a potential story for *Scientific American* about getting neural networks to use logical reasoning.

I would not fall for the hype. Neural nets certainly brought many improvements recently, but they are far from omnipotent and there are many other ML approaches. Some are much faster and work today practically much better when combined with reasoning algorithms. An honest comparison and evaluation is:
<https://arxiv.org/abs/1903.03182> .

Also compare the 70% real time improvement of E prover done on today's largest math benchmark (Mizar) in:

<https://arxiv.org/abs/1904.01677>

with practically no improvement of E due to the slowness (and so far little other advantages) of the neural guidance in:

<https://arxiv.org/abs/1701.06972>

The same holds for comparing tactical guidance:

E.g. the first major breakthrough in <https://arxiv.org/abs/1804.00596> does not need to go neural to beat the state of the art. Whereas the neural systems are yet far from comparable (and explicitly avoid comparison) with state-of-the-art baselines - e.g. <https://arxiv.org/abs/1904.03241>.

I have strongly believed for the last 20 years that combinations of learning and reasoning are very important for building AI. Neural nets are useful in many areas, but I don't think they have yet seriously proved themselves in combination with reasoning. And it's unclear if they will be the last word. Symbolic ML methods may need to be combined with the statistical ones. See also Ben Goertzel's talk: <https://www.youtube.com/watch?v=Zt2HSTuGBn8> .

I came across your paper "Reinforcement Learning of Theorem Proving" (<http://papers.nips.cc/paper/8098-reinforcement-learning-of-theorem-proving>), and I'm wondering if you think that's good example.

Definitely. It is the first real breakthrough in low-level guidance of a very simple (tabula rasa), yet universally applicable reasoning system by ML and AlphaGo/Zero-style RL.

If so, I have a few questions:

-Where in the paper might I find this comparison: "The trained system solves within the same number of inferences over 40% more problems than a baseline prover". I didn't find it on first skim.

In the evaluation tables towards the end.

-How close are we to having AI that outperforms humans?

See my interview for MIRI: <http://intelligence.org/2013/12/21/josef-urban-on-machine-learning-and-automated-reasoning/>

-Is ML the way to go?

ML has shown a great potential and in my opinion it is the most promising approach to combine with reasoning methods to get us to strong assistance of math and science. ML should be understood much more broadly than just deep learning though. There is a lot of hype created by "deep"-invested startups like DeepMind, their counterparts in Mountain View, etc., in the last years. It's good for media, but not as good for science.

Also, what other recent examples of learning to reason stand out? (I'm also looking at SATNet <http://proceedings.mlr.press/v97/wang19e.html> and possibly NS-CL <https://openreview.net/forum?id=rJgMlhRctm>.)

See the papers above.

<https://arxiv.org/abs/1904.01677> finally beats state-of-the-art ATPs with a very large margin. <https://arxiv.org/abs/1804.00596> paves the way toward very strong automated tactical systems.

Two more are:

<https://arxiv.org/abs/1805.06502> - shows that we have a good shot at automatically turning informal math knowledge into formal and verified one - here neural nets may be very useful.

And https://link.springer.com/chapter/10.1007%2F978-3-319-66107-0_2#authorsandaffiliations shows the benefits of combining such translation systems with strong reasoning.

For more info, you can also look at aitp-conference.org - the talks and slides.

Best,
Josef

Thanks!
Matt

.....
Matthew Hutson
Freelance Science Writer
Book: MagicalThinkingBook.com
Science: ScienceMag.org/author/Matthew-Hutson
Twitter: [@SilverJacket](https://twitter.com/SilverJacket)