



UNIVERSITY  
OF WARSAW



UNIVERSITY OF  
OXFORD

# Machine Learning and Logic

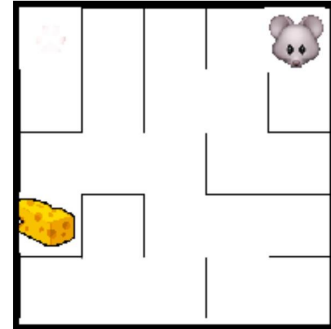
Henryk Michalewski

Google Brain, Mountain View, 28th of November 2018

# what is reinforcement learning?

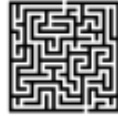
“... neuroscience was ... instrumental in erecting a **second pillar of contemporary AI**, stimulating the emergence of the field of **reinforcement learning (RL)**. RL methods address the problem of *how to maximize future reward by mapping states in the environment to actions* and are among the most widely used tools in AI research. ..., RL methods were originally inspired by research into **animal learning.**”

- Neuroscience-Inspired Artificial Intelligence, Demis Hassabis, Dharshan Kumaran, Christopher Summerfield and Matthew Botvinick, Neuron 2017



# RL crash course

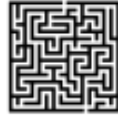
maximize future reward of the agent by mapping states in the ***environment*** to actions



*Environment*

# RL crash course

maximize future reward of the **agent** by mapping states in the **environment** to actions



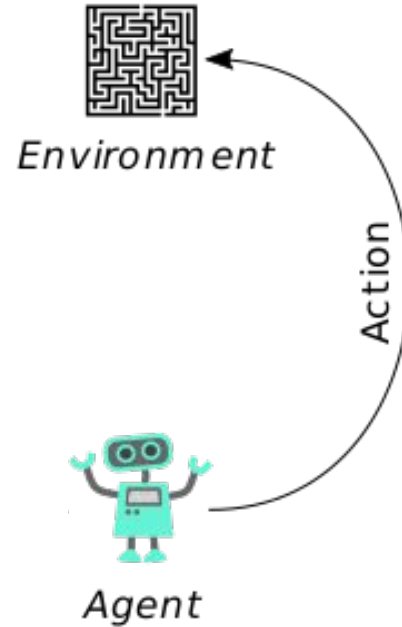
*Environment*



*Agent*

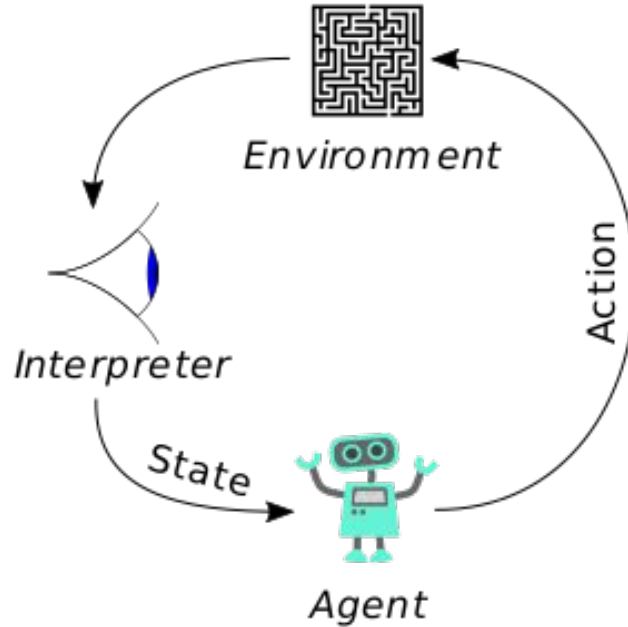
# RL crash course

maximize future reward of the *agent* by mapping states in the *environment* to *actions*



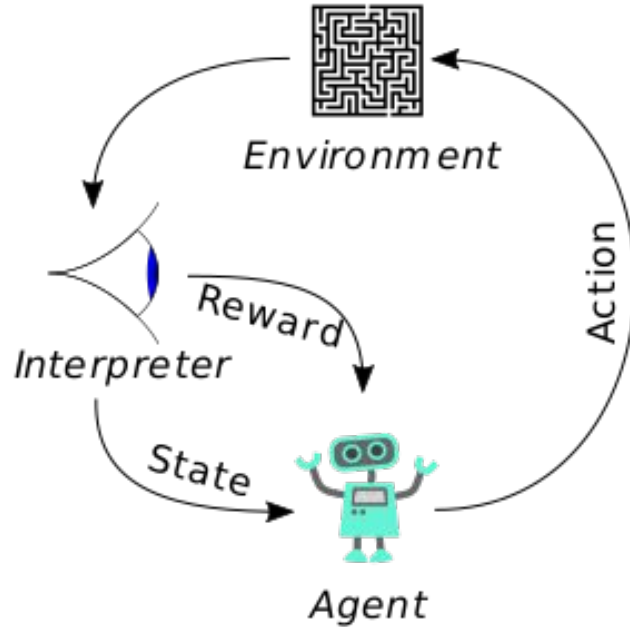
# RL crash course

maximize future reward of the *agent* by mapping *states* in the *environment* to *actions*



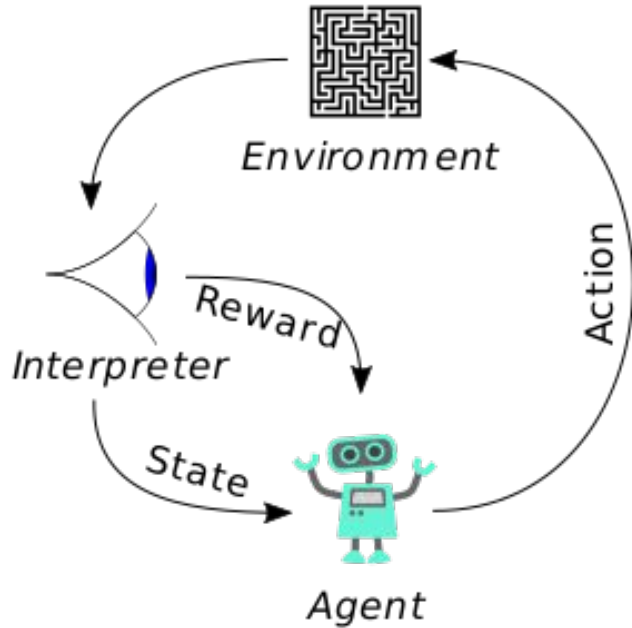
# RL crash course

maximize future **reward** of the **agent** by mapping **states** in the **environment** to **actions**



# RL crash course

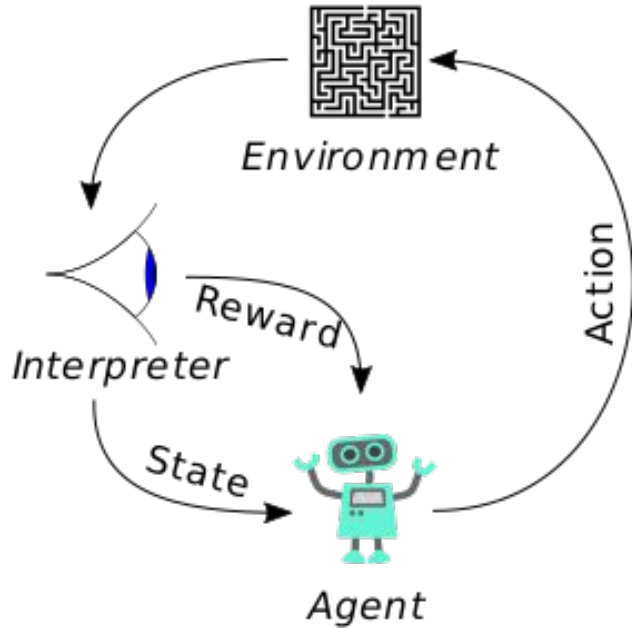
maximize future **reward** of the **agent** by mapping **states** in the **environment** to **actions**





# RL crash course - some math ;)

maximize future **reward** of the **agent** by mapping **states** in the **environment** to **actions**



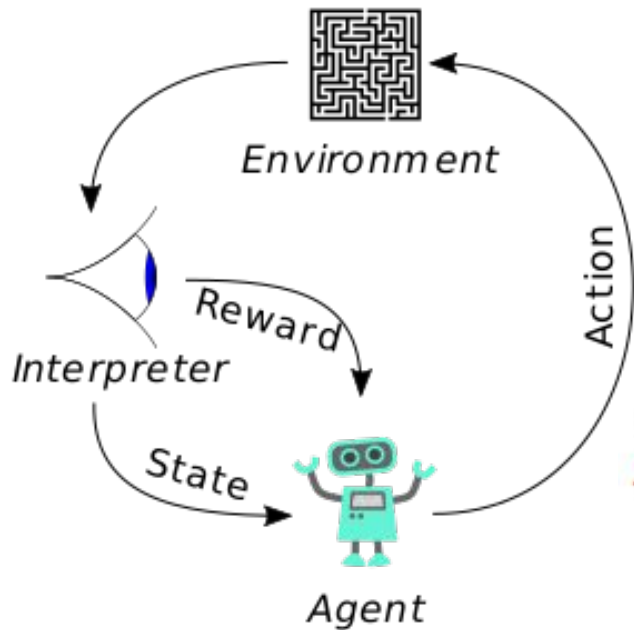
$S$  – state space

$A$  – action space

$\pi : S \mapsto A$  – policy

# RL crash course - some math ;)

maximize future **reward** of the **agent** by mapping **states** in the **environment** to **actions**



$S$  – state space

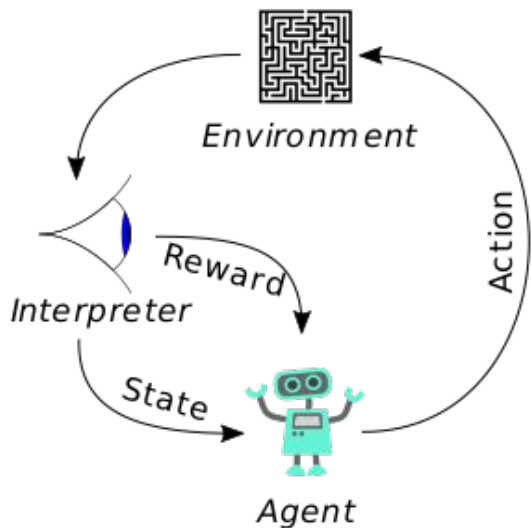
$A$  – action space

$\pi : S \mapsto A$  – policy

$$R_{\pi} = \mathbb{E}_{\pi} (r_1 + r_2 + r_3 + \dots)$$

# RL crash course - some math ;)

maximize future **reward** of the **agent** by mapping **states** in the **environment** to **actions**



$S$  – state space

$A$  – action space

$\pi : S \mapsto A$  – policy

$$R_{\pi} = \mathbb{E}_{\pi} (r_1 + r_2 + r_3 + \dots)$$

**Goal:** find  $\pi_0$  such that  $R_{\pi_0}$  is maximal

## RL - part for adults

**Goal:** find  $\pi_0$  such that  $R_{\pi_0}$  is maximal

$$R_{\pi} = \mathbb{E}_{\pi} (r_1 + r_2 + r_3 + \dots)$$

## RL - part for adults

**Goal:** find  $\pi_0$  such that  $R_{\pi_0}$  is maximal

$$R_{\pi} = \mathbb{E}_{\pi} (r_1 + r_2 + r_3 + \dots)$$

$$\max_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [r(\tau)]$$

## RL - part for adults

**Goal:** find  $\pi_0$  such that  $R_{\pi_0}$  is maximal

$$R_{\pi} = \mathbb{E}_{\pi} (r_1 + r_2 + r_3 + \dots)$$

$$\max_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [r(\tau)]$$

$$\nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [r(\tau)] = \mathbb{E}_{\tau \sim \pi_{\theta}} [\nabla_{\theta} \log(\pi_{\theta}(\tau)) r(\tau)].$$

## RL - part for adults

$$\nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [r(\tau)] = \mathbb{E}_{\tau \sim \pi_{\theta}} [\nabla_{\theta} \log(\pi_{\theta}(\tau)) r(\tau)].$$

$$\hat{g} = \nabla_{\theta} \mathbb{E}_{(s_t, a_t) \sim \pi_0} \left[ \frac{\pi_{\theta}(a_t | s_t)}{\pi_0(a_t | s_t)} \hat{A}_{\pi_0}(s_t, a_t) \right] = \mathbb{E}_{(s_t, a_t) \sim \pi_0} \left[ \frac{\nabla_{\theta} \pi_{\theta}(a_t | s_t)}{\pi_0(a_t | s_t)} \hat{A}_{\pi_0}(s_t, a_t) \right]$$

$$\hat{A}_t = \delta_t + (\gamma\lambda)\delta_{t+1} + \dots + \dots + (\gamma\lambda)^{T-t+1}\delta_{T-1},$$

where  $\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$

# RL - part for adults

An innovation from the PPO paper:

$$\max_{\theta} \mathbb{E}_{(s_t, a_t) \sim \pi} \left[ \min \left( \text{clip}(\rho_t, 1 - \varepsilon, 1 + \varepsilon) \hat{A}_{\pi}(s_t, a_t), \rho_t \hat{A}_{\pi}(s_t, a_t) \right) \right]$$

where

$$\rho_t = \frac{\pi_{\theta}(a_t | s_t)}{\pi(a_t | s_t)}$$



# RL - part for adults

---

## Algorithm 1 PPO, Actor-Critic Style

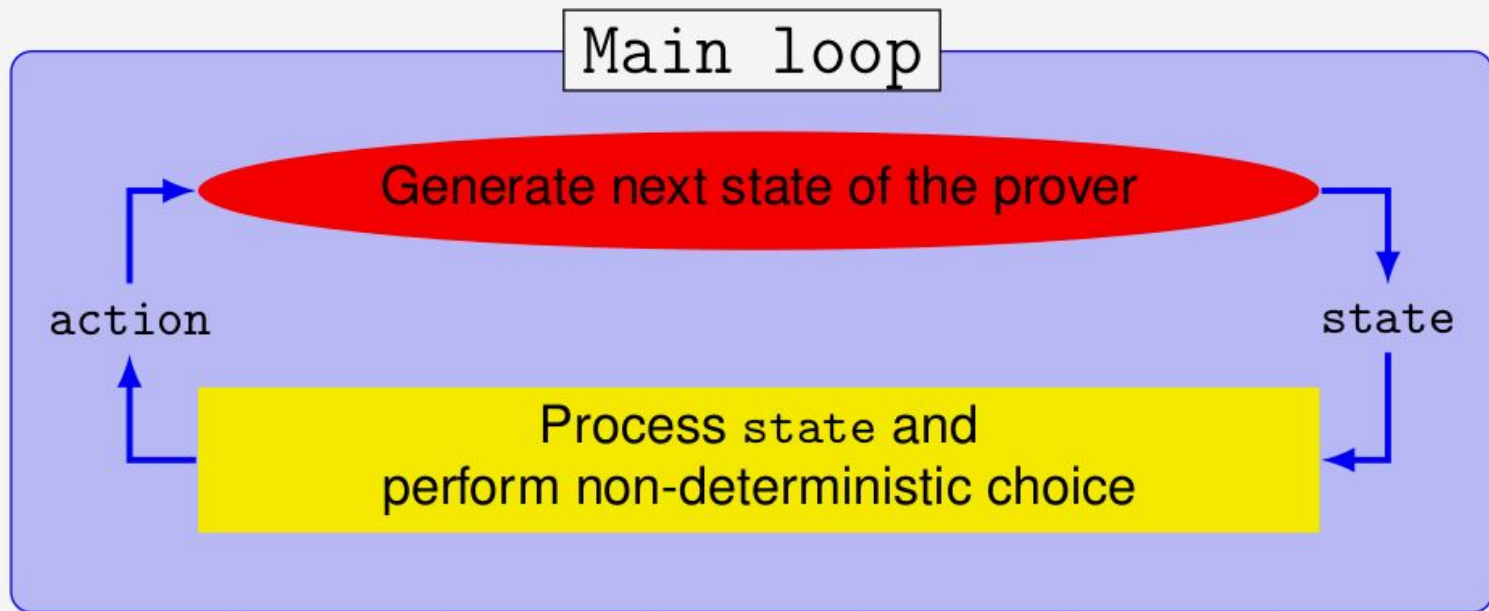
---

```
for iteration=1, 2, ... do  
  for actor=1, 2, ...,  $N$  do  
    Run policy  $\pi_{\theta_{\text{old}}}$  in environment for  $T$  timesteps  
    Compute advantage estimates  $\hat{A}_1, \dots, \hat{A}_T$   
  end for  
  Optimize surrogate  $L$  wrt  $\theta$ , with  $K$  epochs and minibatch size  $M \leq NT$   
   $\theta_{\text{old}} \leftarrow \theta$   
end for
```

---

# Curriculum learning and theorem proving

(joint work with Adrián Csiszárík, Cezary Kaliszyk,  
Josef Urban, Zsolt Zombori)



*Figure 1: The general prover loop.*

```
with tf.variable_scope("value"):
    x = flat_observations
    for size in config.value_layers:
        x = tf.contrib.layers.fully_connected(x, size, tf.nn.relu)
        value = tf.contrib.layers.fully_connected(x, 1, None)[..., 0]
mean = tf.check_numerics(mean, "mean")
logstd = tf.check_numerics(logstd, "logstd")
value = tf.check_numerics(value, "value")

policy = tfp.distributions.MultivariateNormalDiag(mean, tf.exp(logstd))

return NetworkOutput(policy, value, lambda a: tf.clip_by_value(a, -2., 2))
```

<https://github.com/tensorflow/tensor2tensor/blob/master/tensor2tensor/models/research/rl.py>

## How to translate theorem proving to games?

1.  $\forall x \neg s(x) = 0$ ;
2.  $\forall x, y s(x) = s(y) \Rightarrow x = y$ ;
3.  $\forall x x = 0 \vee \exists y x = s(y)$ ;
4.  $\forall x x + 0 = x$ ,
5.  $\forall x, y x + s(y) = s(x + y)$ ;
6.  $\forall x x \cdot 0 = 0$ ;
7.  $\forall x, y x \cdot s(y) = x \cdot y + x$ .

There is no induction scheme.

## How to translate theorem proving to games?

```
cnf(additionZero, axiom, (plus(X, o) = X)).  
cnf(additionSuccessor, axiom, (plus(X, s(Y)) = s(plus(X, Y)))).  
cnf(multiplicationZero, axiom, (mul(X, o) = o)).  
cnf(multiplicationSuccessor, axiom, (mul(X, s(Y)) = plus(mul(X, Y), X))).  
cnf(computeSquare, negated_conjecture, mul(s(s(s(s(s(s(s(s(o))))))))),  
s(s(o)) != s(s(s(s(s(s(s(s(s(s(s(s(s(s(s(s(o)))))))))))))))).
```

# Approach

The RL workshop at NIPS 2018 (joint work with Michał Garmulewicz and Piotr Miłoś)



Use curriculum learning to counterbalance sparse rewards

- Take a few given proof trajectories
- Start exploration from the end of the trajectory
- Gradually move backwards towards the beginning of the proof

Handling the the discrete action space (**two options**):

- Use fixed action space: restricted to a narrow type of problems. **A quick route to strong results, however in general our action space is infinite discrete.**
- Action selector model: policy receives all available actions and returns probabilities. **Can generalize to yet unseen actions.**

## How to translate theorem proving to games?

```
299 ACTION_MEANING = [  
300   (-2062024981253889485, '(B=A=>(D=C=>mul(B,D)=mul(A,C)))'),  
301   (3690526387366737463, '(s(B)=s(A)|~B=A)'),  
302   (4050761576753868900, '(B=A=>(plus(B,D)=plus(A,C)|~D=C))'),  
303   (3472613293128430178, '(B=A=>(C=A|~C=B))'),  
304   (-2002073908694207900, '(B=A=>(C=B=>C=A))'),  
305   (-2003758548611157658, 'plus(mul(A,B),A)=mul(A,s(B))'),  
306   (3833747664017044787, '(B=A|~B=A)'),  
307   (3763098583011768420, 'plus(A,o)=A'),  
308   (3545286421226010726, 'plus(A,s(B))=s(plus(A,B))'),  
309   (3907207153334761009, 'o=mul(A,o)'),  
310   (4121696788395669346, '(B=A=>(D=C=>plus(B,D)=plus(A,C)))'),  
311   (3472663870629359922, '(B=A=>B=A)'),  
312   (3833514601877824305, '(B=A=>(mul(D,B)=mul(C,A)|~D=C))'),  
313   (3618979168708342629, '(B=A=>(C=B|~C=A))'),  
314   (3545284204167390566, 'A=A'),  
315   (3617859672460256866, '(B=A=>(plus(D,B)=plus(C,A)|~D=C))'),  
316   (3990523741287298352, '(B=A=>(mul(B,D)=mul(A,C)|~D=C))'),  
317   (-2147880350953228237, '(B=A=>s(B)=s(A))')  
318 ]
```



# Benefits of action selector model

- Returns a probability distribution (vs scoring function)
- Cannot generalize to yet unseen actions (vs fixed action space model)
- Can handle large action spaces (provided they are not all available together)
- Can make decisions based on other available actions
- Exploits the features of actions
  - Actions and spaces are embedded into the same feature space (manual features, widely used in the ATP community)
  - The network can meaningfully process the concatenation of state-action feature vectors.

# Results

Trained and tested on Robinson arithmetic problems

RL method - variations of the Actor-Critic meta-algorithm:

- Fixed action space model trained on  $7*2=14$  (51 steps) managed to prove all multiplications up to  $9*9$  (1201 steps)
- Action selector model trained on  $7+5=12$  (17 steps) managed to prove all additions up to  $9+9$

## What is curriculum learning in the context of reinforcement learning?

We have a proof trajectory:

```
('/home/henrykmichalewski/problems/gen_mul_new/pe  
ano_mul_7_2_14.p', [0, 1, 0, 2, 4, 0, 2, 4, 0, 2, 4, 0, 2,  
4, 0, 2, 4, 0, 2, 4, 0, 2, 4, 0, 3, 0, 1, 0, 2, 4, 0, 2, 4, 0, 2,  
4, 0, 2, 4, 0, 2, 4, 0, 2, 4, 0, 2, 4, 0, 3, 5]),
```

### Curriculum learning recipe:

1. Start close to the end of trajectory
2. Gradually raise the bar
3. Check whether it generalize to other problems





# SQL queries and machine learning

**Joint project with Michael Benedikt** (work in progress)

**Input:** a query  $Q$ , a set of views/view definitions

**Output:** a query using only the views that is equivalent to  $Q$

**Input:** a query  $Q$ , a set of views/view definitions

**Output:** a query using only the views that is equivalent to  $Q$

Query:

```
SELECT DISTINCT c.custname
FROM Customer c, Orders o
WHERE c.custkey=o.custkey
AND NOT EXISTS
(SELECT li.quantity FROM
LineItem li WHERE
li.quantity<2
AND li.orderkey=o.orderkey)
```

Views:

```
V1 (orderkey, ...) = SELECT * FROM
LineItem li WHERE li.quantity<2

V2 (custkey, custname, orderkey) =
SELECT * FROM Customers c, Orders o
WHERE c.custkey=o.custkey
```



```

fof(q1, axiom,
(? [Ckey, Okey] : customer(Ckey, mycname) & order(Okey,Ckey)
& ~ ( ? [Li, Qty] : ( lineitem(Li, Okey, Qty) & le2(Qty)) )
)
).

fof(const1,axiom,
(
! [Li, Okey, Qty] :
(v1(Okey) <=> ( lineitem(Li, Okey, Qty) & le2(Qty) ) )
)
).

fof(const2, axiom,
(
! [Ckey, Cname, Okey] :
(v2(Ckey, Cname, Okey) <=> ( customer(Ckey, Cname) & order(Okey, Ckey) )
)
)
).

fof(const1p,axiom,
(
! [Li, Okey, Qty] :
(v1(Okey) <=> ( lineitem(Li, Okey, Qty) & le2(Qty)) )
)
).

fof(const2p, axiom,
(
! [Ckey, Cname, Okey] :
(v2(Ckey, Cname, Okey) <=> ( customerp(Ckey, Cname) & orderp(Okey, Ckey) )
)
)
).

fof(q1p, conjecture,
(? [Ckey, Okey] : customerp(Ckey, mycname) & orderp(Okey,Ckey)
& ~ ( ? [Li, Qty] : ( lineitem(Li, Okey, Qty) & le2(Qty)) )
)
).

```

Query:

```

SELECT DISTINCT c.custname
FROM Customer c, Orders o
WHERE c.custkey=o.custkey
AND NOT EXISTS
  (SELECT li.quantity FROM
   LineItem li WHERE
   li.quantity<2
   AND li.orderkey=o.orderkey)

```

Views:

```

V1 (orderkey, ...) = SELECT * FROM
LineItem li WHERE li.quantity<2

V2 (custkey, custname, orderkey) =
SELECT * FROM Customers c, Orders o
WHERE c.custkey=o.custkey

```

Query:

```
SELECT DISTINCT c.custname
FROM Customer c, Orders o
WHERE c.custkey=o.custkey
AND NOT EXISTS
(SELECT li.quantity FROM
LineItem li WHERE
li.quantity<2
AND li.orderkey=o.orderkey)
```

Rewriting:

```
SELECT DISTINCT v2.custname FROM V2 v2
WHERE v2.orderkey NOT IN
(SELECT orderkey FROM V1)
```

Views:

```
V1 (orderkey, ...) = SELECT * FROM
LineItem li WHERE li.quantity<2
```

```
V2 (custkey, custname, orderkey) =
SELECT * FROM Customers c, Orders o
WHERE c.custkey=o.custkey
```

Query:

```
SELECT DISTINCT c.custname
FROM Customer c, Orders o
WHERE c.custkey=o.custkey
AND NOT EXISTS
(SELECT li.quantity FROM
LineItem li WHERE
li.quantity<2
AND li.orderkey=o.orderkey)
```

Views:

```
V1 (orderkey, ...) = SELECT * FROM
LineItem li WHERE li.quantity<2
```

```
V2 (custkey, custname, orderkey) =
SELECT * FROM Customers c, Orders o
WHERE c.custkey=o.custkey
```

**Craig interpolation theorem.** Let  $\varphi_1, \varphi_2$  be first-order sentences over the vocabularies  $\mathcal{L}_1, \mathcal{L}_2$  such that  $\varphi_1 \models \varphi_2$ . Then there is a sentence  $\theta$  over the common vocabulary  $\mathcal{L}_1 \cap \mathcal{L}_2$ , the *interpolant*, such that  $\varphi_1 \models \theta$  and  $\theta \models \varphi_2$ .

```
SELECT DISTINCT v2.custname FROM V2 v2
WHERE v2.orderkey NOT IN
(SELECT orderkey FROM V1)
```

# Neural architectures for SAT

Joint work with Sebastian Jaszczur, Michał Łuszczczyk

---

**Algorithm 1** Simplified propositional SAT algorithm

---

```
1: function DPLL( $\Phi$ )
2:   if  $\Phi$  is empty then return True
3:   if  $\Phi$  contains an empty clause then return False
4:    $l \leftarrow$  choose-literal( $\Phi$ )
5:   return DPLL( $\Phi \wedge l$ ) or DPLL( $\Phi \wedge \text{not } l$ )
```

---

In this experiment we deal with formulas given in the conjunctive normal form (CNF):

$$(x_1 \vee x_2 \vee \dots \vee x_n) \wedge$$

$$(y_1 \vee x_2 \vee \dots \vee x_n) \wedge$$

$$(x_1 \vee y_2 \vee \dots \vee x_n) \wedge$$

$$(y_1 \vee y_2 \vee \dots \vee x_n) \wedge \dots \wedge$$

$$(x_1 \vee x_2 \vee \dots \vee y_n) \wedge$$

$$(y_1 \vee x_2 \vee \dots \vee y_n) \wedge$$

$$(x_1 \vee y_2 \vee \dots \vee y_n) \wedge$$

$$(y_1 \vee y_2 \vee \dots \vee y_n);$$

$$(l_1 \vee l_2 \vee x_2) \wedge$$

$$(\neg x_2 \vee l_3 \vee x_3) \wedge$$

$$(\neg x_3 \vee l_4 \vee x_4) \wedge \dots \wedge$$

$$(\neg x_{n-3} \vee l_{n-2} \vee x_{n-2}) \wedge$$

$$(\neg x_{n-2} \vee l_{n-1} \vee l_n)$$

---

**Algorithm 1** Simplified propositional SAT algorithm

---

```
1: function DPLL( $\Phi$ )
2:   if  $\Phi$  is empty then return True
3:   if  $\Phi$  contains an empty clause then return False
4:    $l \leftarrow$  choose-literal( $\Phi$ )
5:   return DPLL( $\Phi \wedge l$ ) or DPLL( $\Phi \wedge \text{not } l$ )
```

---

In this experiment we want to do things which are sound and simple from engineering point of view:

**Done:**

- Very simple implementation of DPLL for prototyping (allows for various heuristics)
- Generators random, but non-trivial instances
- Various neural architectures compared using the above infrastructure
- Generic distributed training on a cluster

**Not done:**

- Integration with a more sophisticated solver
- Training/testing on large and very large instances
- Implementation of DPLL in tf (aka on GPU)
- TPU-compliant implementation of NeuroDPLL

**Goal:** Learned policy or embedding which outdoes basic heuristics and generalizes to large instances than these in training

## Algorithm 1 Simplified propositional SAT algorithm

```
1: function DPLL( $\Phi$ )
2:   if  $\Phi$  is empty then return True
3:   if  $\Phi$  contains an empty clause then return False
4:    $l \leftarrow$  choose-literal( $\Phi$ )
5:   return DPLL( $\Phi \wedge l$ ) or DPLL( $\Phi \wedge \text{not } l$ )
```



In this experiment we want to do things which are sound and simple from engineering point of view:

### Done:

- Very simple implementation of DPLL for prototyping (allows for various heuristics)
- Generators random, but trivial instances
- Various neural architectures compared using the above infrastructure
- Generic distributed training on a cluster

deep**sense**.ai  
BIG DATA SCIENCE

Start with AI | AI solutions



### Running distributed TensorFlow on Slurm clusters

June 26, 2017 / 0 Comments / in Data science, Deep learning, Machine learning / by Tomasz Grel

In this post, we provide an example of how to run a TensorFlow experiment on a Slurm cluster. Since TensorFlow doesn't yet officially support this task, we developed a simple Python module for automating the configuration. It parses the environment variables set by Slurm and creates a TensorFlow cluster configuration based on them. We're sharing this code along with a simple image recognition example on CIFAR-10. You can find it in our [github repo](#).

**Goal:** Learned policy or embedding which outdoes basic heuristics and generalizes to large instances than these in training

---

**Algorithm 1** Simplified propositional SAT algorithm

---

```
1: function DPLL( $\Phi$ )
2:   if  $\Phi$  is empty then return True
3:   if  $\Phi$  contains an empty clause then return False
4:    $l \leftarrow$  choose-literal( $\Phi$ )
5:   return DPLL( $\Phi \wedge l$ ) or DPLL( $\Phi \wedge \text{not } l$ )
```

---

In this experiment we want to do things which are sound and simple from engineering point of view:

**Goal:** Learned policy or embedding which outdoes basic heuristics and generalizes to large instances than these in training

**Long-term goal:** learn the distributions in the random SAT benchmark

*Proceedings of SAT Competition 2017: Solver and Benchmark Descriptions*, volume B-2017-1 of Department of Computer Science Series of Publications B, University of Helsinki 2017.

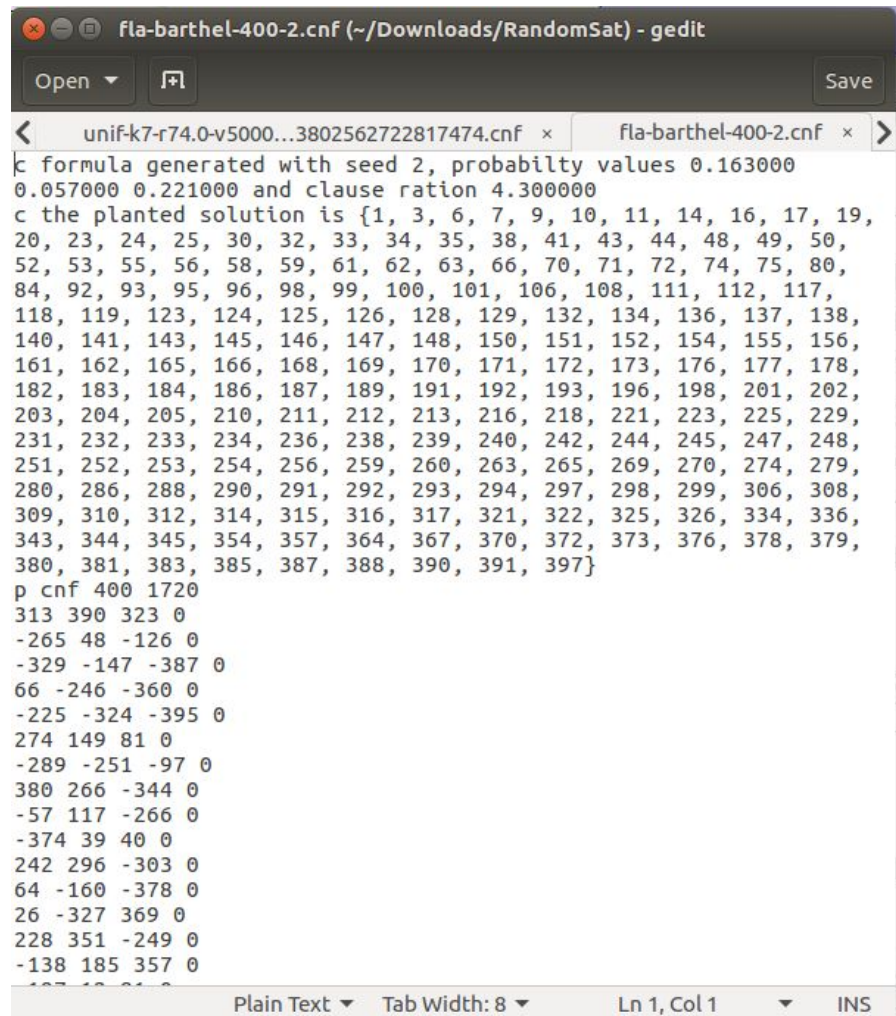
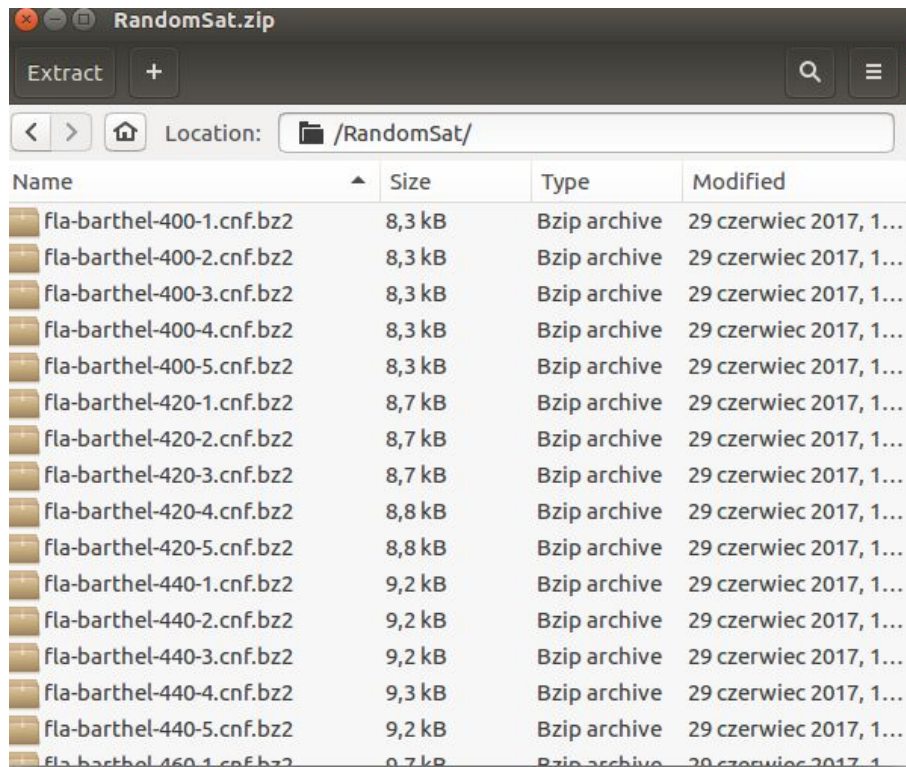
## Generating the Uniform Random Benchmarks

Marijn J. H. Heule

Department of Computer Science,  
The University of Texas at Austin, United States

*Abstract*—The uniform random k-SAT instances described here, together with the hard satisfiable random instances described on pages XXX of this compilation, constitute the benchmark set of the Random Track of SAT Competition 2017.





---

## Learning a SAT Solver from Single-Bit Supervision

---

Daniel Selsam<sup>1</sup> Matthew Lamm<sup>2</sup> Benedikt Bünnz<sup>1</sup> Percy Liang<sup>1</sup> Leonardo de Moura<sup>3</sup> David L. Dill<sup>1</sup>

### Abstract

We present NeuroSAT, a message passing neural network that learns to solve SAT problems after only being trained as a classifier to predict satisfiability. Although it is not competitive with state-of-the-art SAT solvers, NeuroSAT can solve problems that are substantially larger and more difficult than it ever saw during training by simply running for more iterations. Moreover, NeuroSAT generalizes to novel distributions; after training only on random SAT problems, at test time it can solve SAT problems encoding graph coloring, clique detection, dominating set, and vertex cover problems, all on a range of distributions over small random graphs.

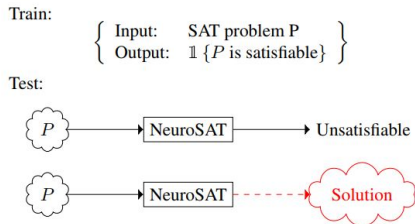


Figure 1. We train NeuroSAT to predict whether SAT problems are satisfiable, providing only a single bit of supervision for each problem. At test time, when NeuroSAT predicts *satisfiable*, we can almost always extract a satisfying assignment from the network’s activations. The problems at test time can also be substantially

## CAN NEURAL NETWORKS UNDERSTAND LOGICAL ENTAILMENT?

Richard Evans\*

David Saxton\*

David Amos

Pushmeet Kohli

Edward Grefenstette\*

DeepMind

{richardevans, saxton, davidamos, pushmeet, etg}@google.com

### ABSTRACT

We introduce a new dataset of logical entailments for the purpose of measuring models’ ability to capture and exploit the structure of logical expressions against an entailment prediction task. We use this task to compare a series of architectures which are ubiquitous in the sequence-processing literature, in addition to a new model class—PossibleWorldNets—which computes entailment as a “convolution over possible worlds”. Results show that convolutional networks present the wrong inductive bias for this class of problems relative to LSTM RNNs, tree-structured neural networks outperform LSTM RNNs due to their enhanced ability to exploit the syntax of logic, and PossibleWorldNets outperform all benchmarks.

# DPLL + some embedding

## Algorithm DPLL

Input: A set of clauses  $\Phi$ .

Output: A Truth Value.

## function DPLL( $\Phi$ )

**if**  $\Phi$  is empty

**then return** true;

**if**  $\Phi$  contains an empty clause

**then return** false;

// **for every** unit clause  $l$  **in**  $\Phi$

//  $\Phi \leftarrow \text{unit-propagate}(l, \Phi)$ ;

// **for every** literal  $l$  that occurs pure **in**  $\Phi$

//  $\Phi \leftarrow \text{pure-literal-assign}(l, \Phi)$ ;

$l \leftarrow \text{choose-literal}(\Phi)$ ;

**return** DPLL( $\Phi \wedge l$ ) **or** DPLL( $\Phi \wedge \text{not}(l)$ );

## function choose-literal( $\Phi$ ):

**for each variable**  $v$

**calculate embedding of**  $\Phi[v=\text{true}]$

**calculate embedding of**  $\Phi[v=\text{false}]$

**take literal for which embedding is farthest from False**

# DPLL + some trained policy choose-literal

## Algorithm DPLL

Input: A set of clauses  $\Phi$ .

Output: A Truth Value.

## function DPLL( $\Phi$ )

**if**  $\Phi$  is empty

**then return** true;

**if**  $\Phi$  contains an empty clause

**then return** false;

// **for every** unit clause  $l$  **in**  $\Phi$

//  $\Phi \leftarrow \text{unit-propagate}(l, \Phi)$ ;

// **for every** literal  $l$  that occurs pure **in**  $\Phi$

//  $\Phi \leftarrow \text{pure-literal-assign}(l, \Phi)$ ;

$l \leftarrow \text{choose-literal}(\Phi)$ ;

**return** DPLL( $\Phi \wedge l$ ) **or** DPLL( $\Phi \wedge \text{not}(l)$ );

# Axiomatization table for CNF

Axiom	EqNet	LSTM	NeuroSAT
“Variable renaming” - independent			
“Permutation of literals in clause” - independent			
“Permutation of clauses in formula” - independent			
“Negation of all occurrences of variable” - independent			

In this presentation I will conflate terms NeuroSAT, NeuroDPLL, GraphSAT, GraphDPLL ...  
- they all refer to a SAT policy based on a NeuroSAT-style architecture.

# EQNET

## Learning Continuous Semantic Representations of Symbolic Expressions

Goal: learn a mapping formula ->  
vector space

Semantically equivalent formulas  
should be close.

---

### Learning Continuous Semantic Representations of Symbolic Expressions

---

Miltiadis Allamanis<sup>1</sup> Pankajan Chanthirasegaran<sup>2</sup> Pushmeet Kohli<sup>3</sup> Charles Sutton<sup>2,4</sup>

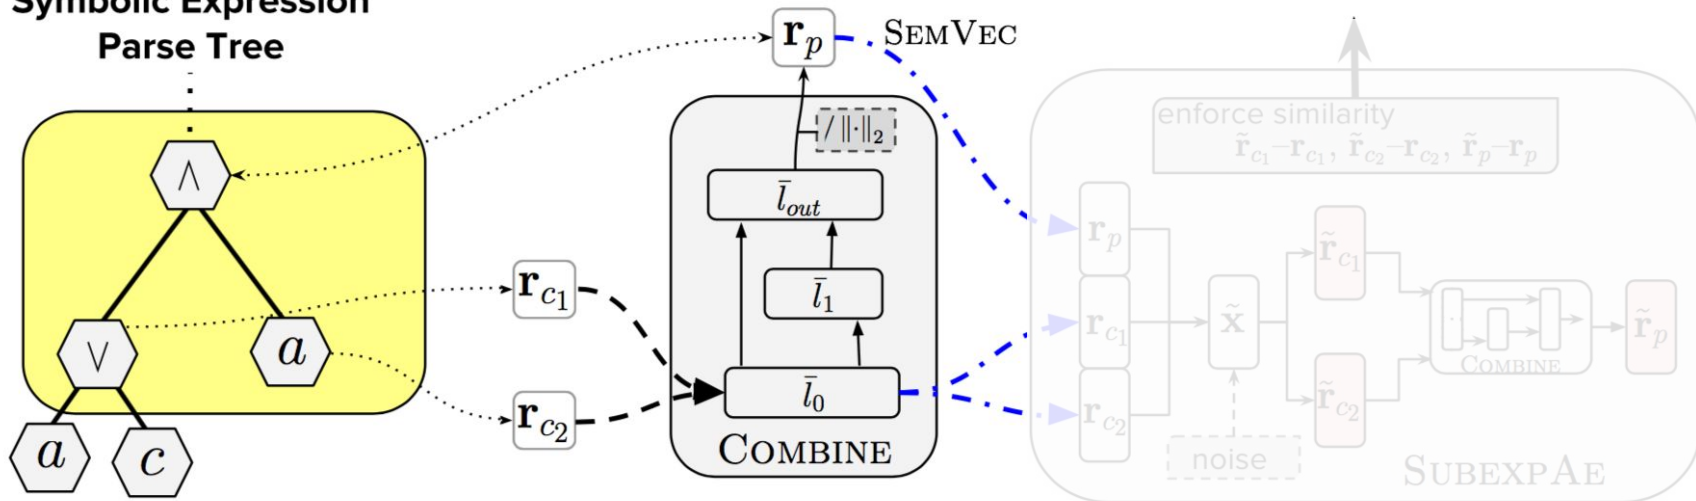
#### Abstract

Combining abstract, symbolic reasoning with continuous neural reasoning is a grand challenge of representation learning. As a step in this direction, we propose a new architecture, called *neural equivalence networks*, for the problem of learning continuous semantic representations of algebraic and logical expressions. These networks are trained to represent semantic equivalence, even of expressions that are syntactically very different. The challenge is that semantic representations must be computed in a syntax-directed manner, because semantics is compositional, but at the same time, small changes in syntax can lead to very large changes in semantics, which can be difficult for continuous neural architectures. We perform an exhaustive evaluation on the task of checking equivalence on a highly diverse class of symbolic algebraic and boolean expression types, showing that our model significantly outperforms existing architectures.

ships. However, apart from some notable exceptions (Alemi et al., 2016; Loos et al., 2017; Zaremba et al., 2014), this area has received relatively little attention in machine learning. In this work, we explore the direction of learning continuous *semantic* representations of symbolic expressions. The goal is for expressions with similar semantics to have similar continuous representations, even if their syntactic representation is very different. Such representations have the potential to allow a new class of symbolic reasoning methods based on heuristics that depend on the continuous representations, for example, by guiding a search procedure in a symbolic solver based on a distance metric in the continuous space. In this paper, we make a first essential step of addressing the problem of learning continuous semantic representations (SEMVECs) for symbolic expressions. Our aim is, given access to a training set of pairs of expressions for which semantic equivalence is known, to assign continuous vectors to symbolic expressions in such a way that semantically equivalent, but syntactically diverse expressions are assigned to identical (or highly similar) continuous vectors. This is an important but hard problem; learning composable SEMVECs of symbolic expressions requires that we learn about the semantics of symbolic elements and operators

## Symbolic Expression

### Parse Tree



(a) Architectural diagram of EQNETs. Example parse tree shown is of the boolean expression  $(a \vee c) \wedge a$ .

TREENN (current node  $n$ )

**if**  $n$  **is not a leaf then**

$$\mathbf{r}_n \leftarrow \text{COMBINE}(\text{TREENN}(c_0), \dots, \text{TREENN}(c_k), \tau_n),$$

where  $(c_0, \dots, c_k) = \text{ch}(n)$

**else**

$$\mathbf{r}_n \leftarrow \text{LOOKUPLEAFEMBEDDING}(\tau_n)$$

**return**  $\mathbf{r}_n$

**COMBINE** ( $\mathbf{r}_{c_0}, \dots, \mathbf{r}_{c_k}, \tau_p$ )

$$\bar{l}_0 \leftarrow [\mathbf{r}_{c_0}, \dots, \mathbf{r}_{c_k}]$$

$$\bar{l}_1 \leftarrow \sigma(W_{i, \tau_p} \cdot \bar{l}_0)$$

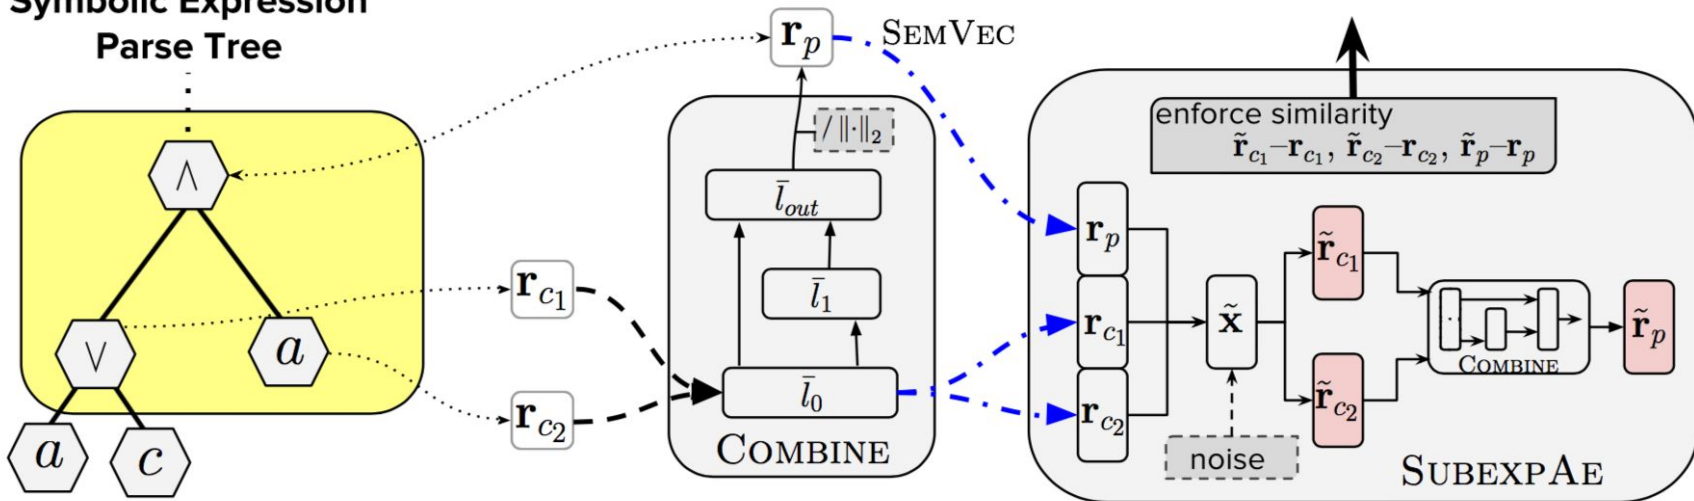
$$\bar{l}_{out} \leftarrow W_{o0, \tau_p} \cdot \bar{l}_0 + W_{o1, \tau_p} \cdot \bar{l}_1$$

$$\mathbf{return} \bar{l}_{out} / \|\bar{l}_{out}\|_2$$

(b) **COMBINE** of EQNET.

## Symbolic Expression

### Parse Tree



(a) Architectural diagram of EQNETs. Example parse tree shown is of the boolean expression  $(a \vee c) \wedge a$ .

SUBEXPAE ( $\mathbf{r}_{c_0}, \dots, \mathbf{r}_{c_k}, \mathbf{r}_p, \tau_p$ )

$$\mathbf{x} \leftarrow [\mathbf{r}_{c_0}, \dots, \mathbf{r}_{c_k}]$$

$$\tilde{\mathbf{x}} \leftarrow \tanh(W_d \cdot \tanh(W_{e, \tau_p} \cdot [\mathbf{r}_p, \mathbf{x}] \cdot \mathbf{n}))$$

$$\tilde{\mathbf{x}} \leftarrow \tilde{\mathbf{x}} \cdot \|\mathbf{x}\|_2 / \|\tilde{\mathbf{x}}\|_2$$

$$\tilde{\mathbf{r}}_p \leftarrow \text{COMBINE}(\tilde{\mathbf{x}}, \tau_p)$$

$$\text{return} - (\tilde{\mathbf{x}}^\top \mathbf{x} + \tilde{\mathbf{r}}_p^\top \mathbf{r}_p)$$

- goal: reversibility
- bottleneck model aka autoencoder
- “not used to compute the parent representation, regulariser only”

(c) Loss function used for subexpression autoencoder



# EQNET as a SAT oracle

Experiment:

1. trained eqnet: size 10, 5 vars. Operators AND, OR and NOT.  
Accuracy (claimed and reproduced): 85%.
2. calculate embeddings on a freshly generated set.
3. calculate distances from embedding of False
4. plot one histogram for satisfiable formulas and one from unsatisfiable formulas.

# DPLL + eqnet

## Algorithm DPLL

Input: A set of clauses  $\Phi$ .

Output: A Truth Value.

## function DPLL( $\Phi$ )

```
if  $\Phi$  is empty
  then return true;
if  $\Phi$  contains an empty clause
  then return false;
// for every unit clause  $l$  in  $\Phi$ 
//    $\Phi \leftarrow \text{unit-propagate}(l, \Phi)$ ;
// for every literal  $l$  that occurs pure in  $\Phi$ 
//    $\Phi \leftarrow \text{pure-literal-assign}(l, \Phi)$ ;
 $l \leftarrow \text{choose-literal}(\Phi)$ ;
return DPLL( $\Phi \wedge l$ ) or DPLL( $\Phi \wedge \text{not}(l)$ );
```

## function choose-literal( $\Phi$ ):

**for each variable**  $v$

calculate embedding of  $\Phi[v=\text{true}]$

calculate embedding of  $\Phi[v=\text{false}]$

**take literal for which embedding is farthest from False**

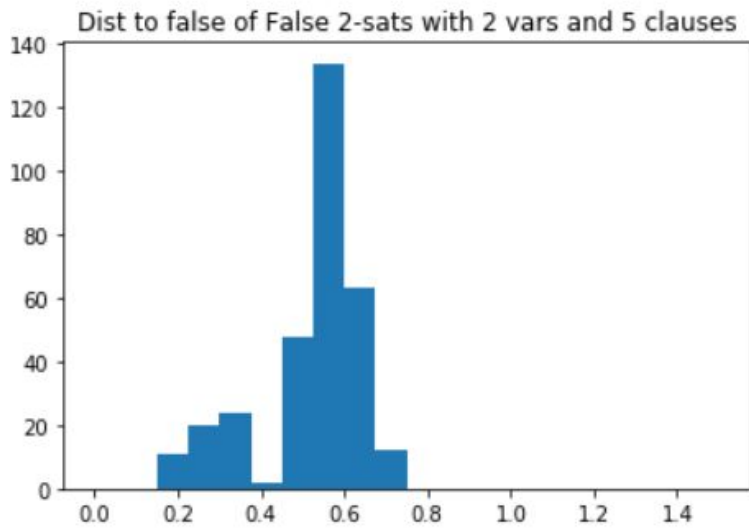
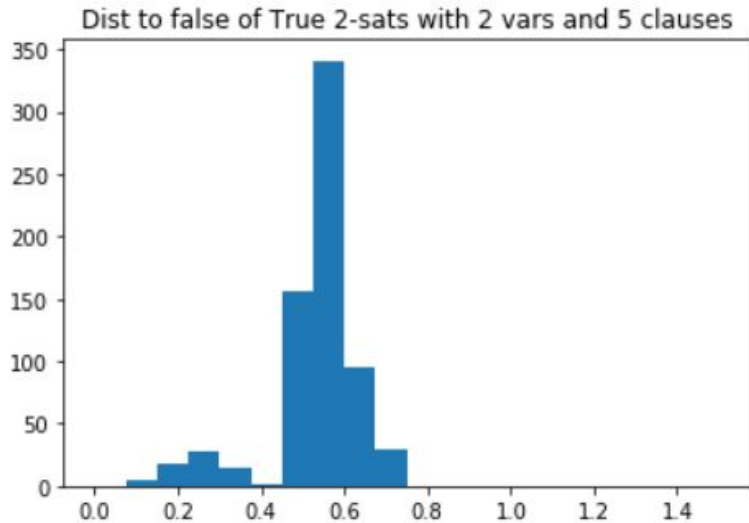
## Experiment 2. [Notebook link.](#)

Dataset with random k-SATs:

2-SAT, 2 variables, 5 clauses

#Sat: 686; avg dist: 0.5334; stdev dist: 0.103

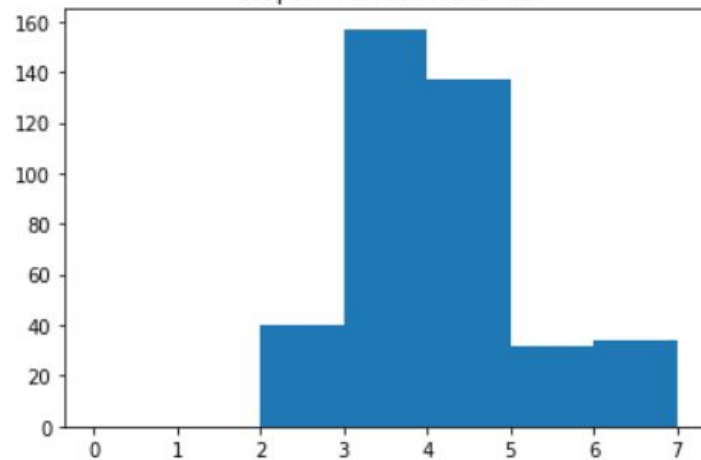
#Unsat: 314; avg dist: 0.5207; stdev dist: 0.124



DPLL  
results: [link](#)  
  
2-SAT,  
2 variables,  
3 clauses

#Sats: 400; avg step: 3.6925; stdev step: 1.13707

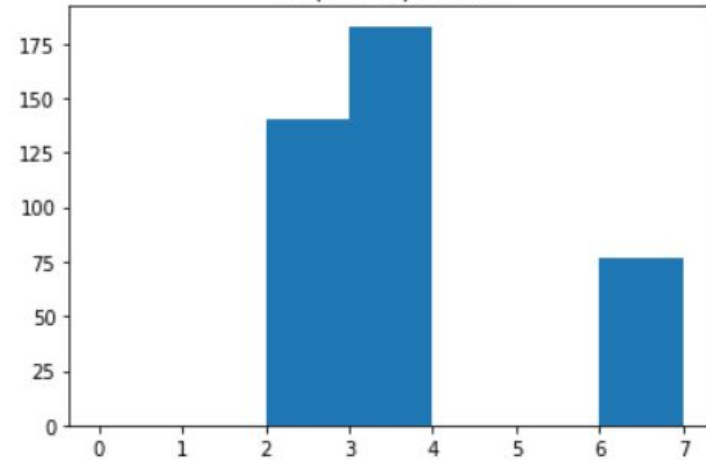
Steps of RandomVarDPLL



#Sats: 400; avg step: 3.0025; stdev step: 0.75663

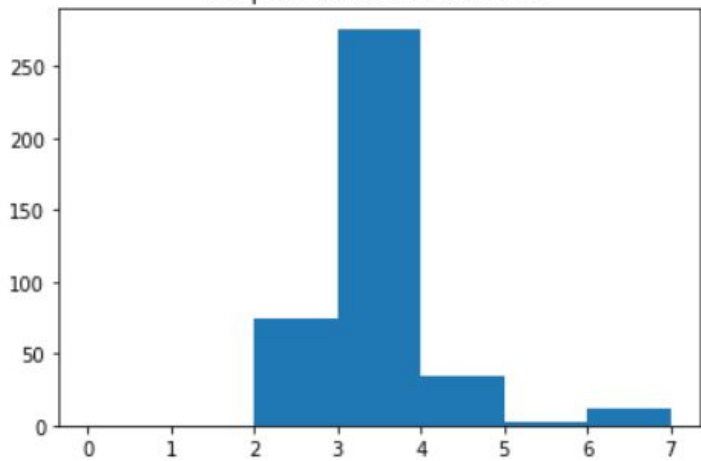
#Sats: 400; avg step: 3.2275; stdev step: 1.42504

Steps of EqnetDPLL

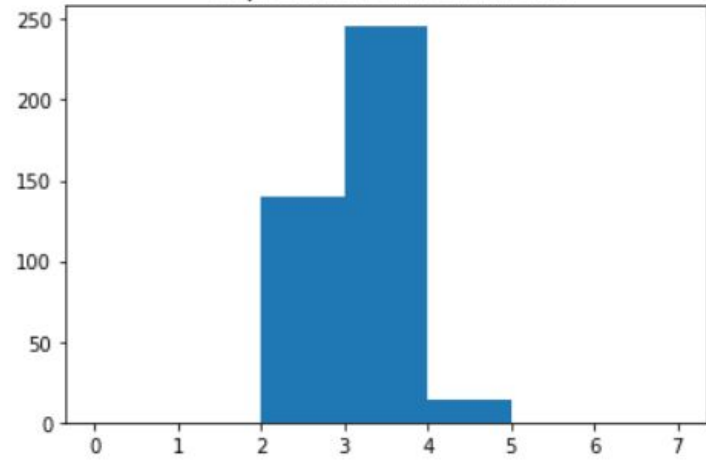


#Sats: 400; avg step: 2.685; stdev step: 0.534579

Steps of RandomClauseDPLL

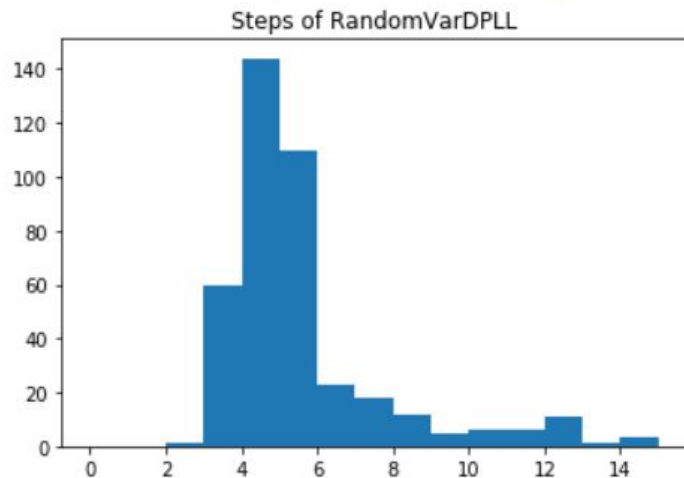


Steps of MostCommonVarDPLL

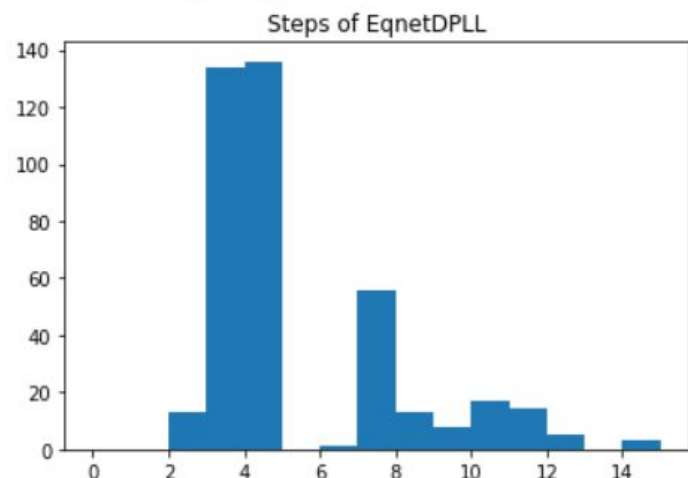


DPLL  
results: [link](#)  
  
3-SAT,  
3 variables,  
6 clauses

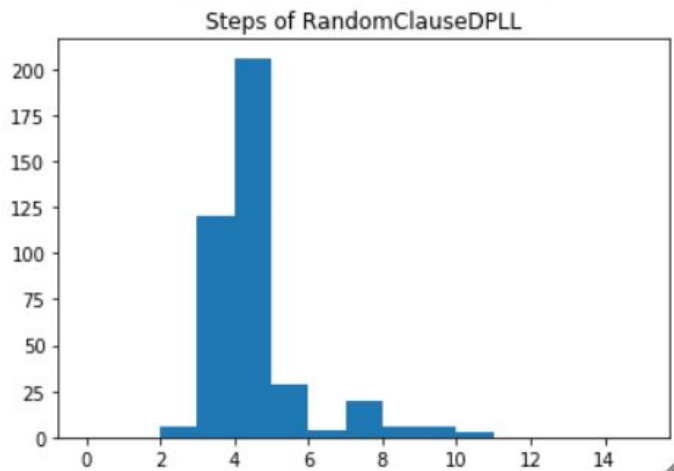
#Sats: 400; avg step: 5.0675; stdev step: 2.182416



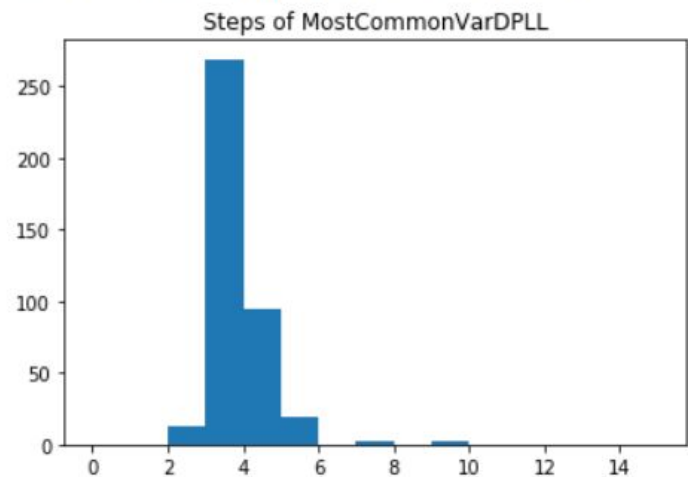
#Sats: 400; avg step: 4.93; stdev step: 2.5855560



#Sats: 400; avg step: 4.0925; stdev step: 1.34124

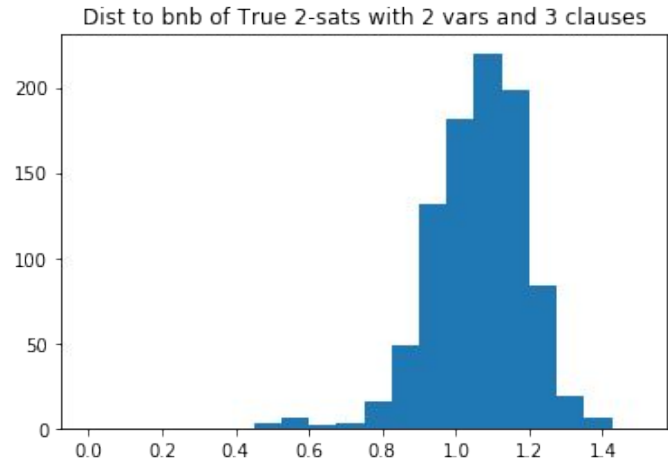
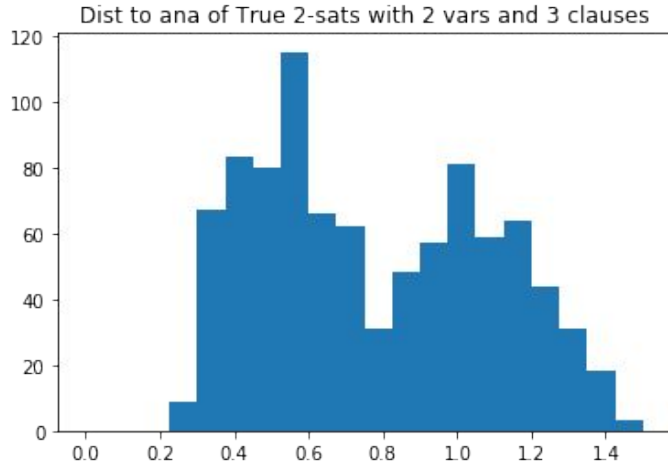


#Sats: 400; avg step: 3.35; stdev step: 0.7729812



# Experiment: $a \& \sim a$ versus $b \& \sim b$

[Notebook link](#)



distance from/to	$a \& \sim a$	$b \& \sim b$	$c \& \sim c$
$a \& \sim a$	0.0	1.14	1.07
$b \& \sim b$	1.14	0.0	0.9
$c \& \sim c$	1.07	0.9	0.0

# Standard solutions: why not LSTM?

Treat like a sequence:  $(A \vee \sim C \vee B) \wedge (\sim B \wedge C) \wedge (\sim A \vee B \vee D)$

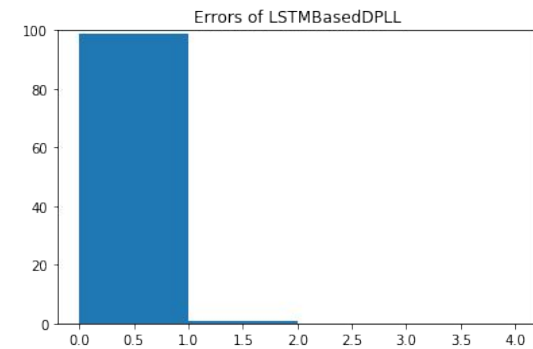
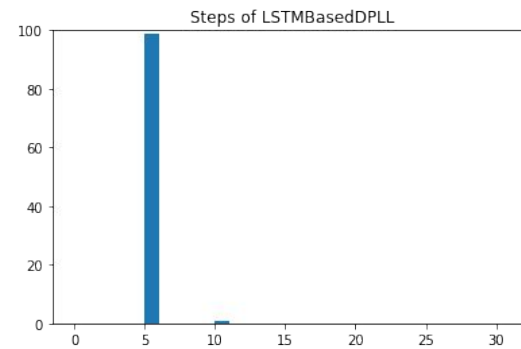
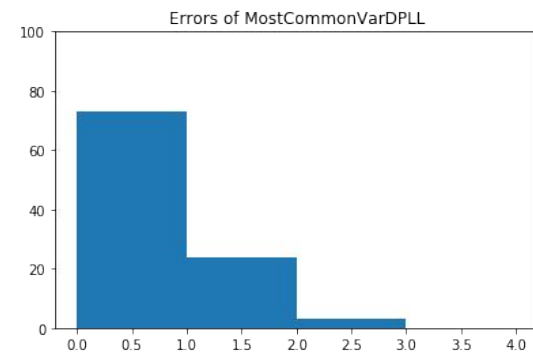
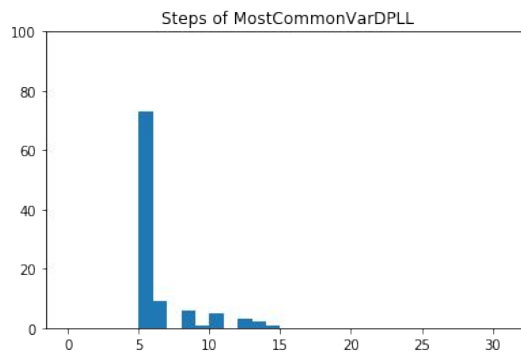
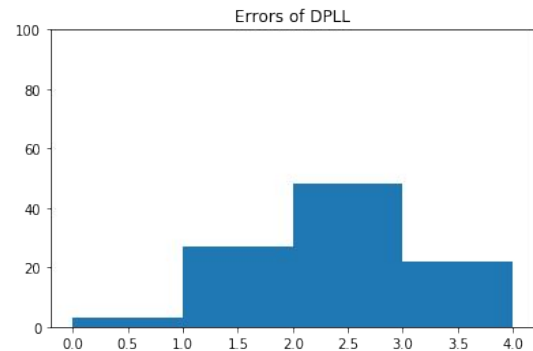
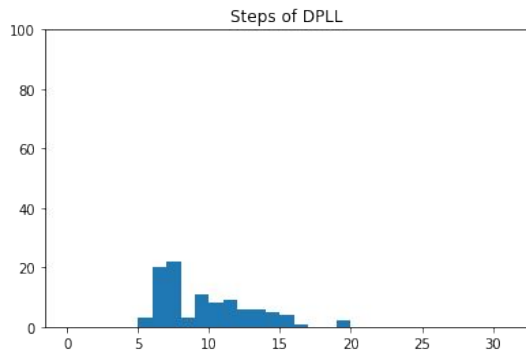
- Problems with LSTMs and similar:
  - 😞 Permutation of variables inside a clause changes the result.
  - 😞 Permutation of clauses changes the result.
  - 😞 Negation of all occurrences of a variable changes the result.
  - 😞 Definite limit on embedding size, regardless of the size of the formula.

# Experiment 1.

#vars: 4

trained on up to 40 clauses

test on up to 70 clauses

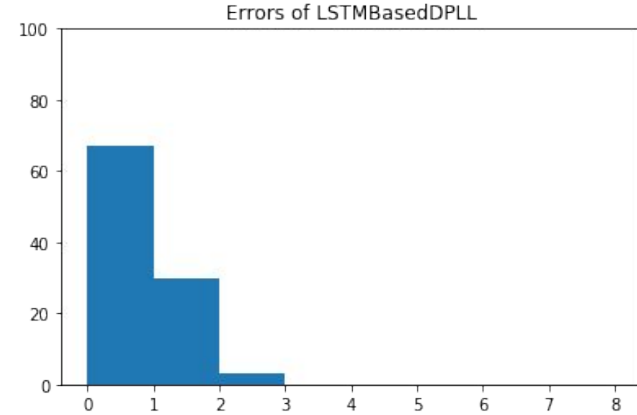
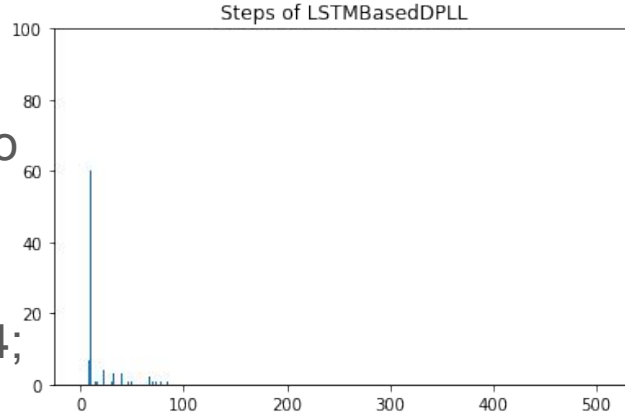




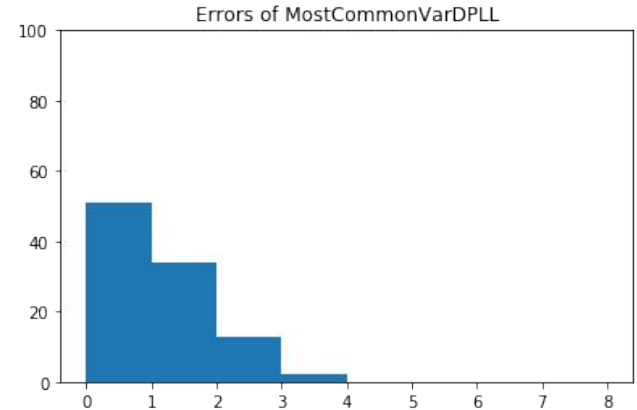
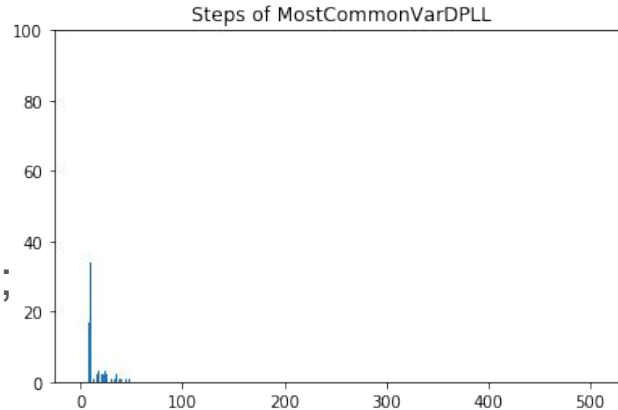
# Experiment 3.

#vars: 8, #clauses: up to 50

#Sats: 100; avg step: 19.73; stdev step: 19.94;  
avg error: 0.36; stdev error: 0.54



#Sats: 100; avg step: 15.27; stdev step: 11.05;  
avg error: 0.66; stdev error: 0.78



---

# Learning a SAT Solver from Single-Bit Supervision

---

Daniel Selsam<sup>1</sup> Matthew Lamm<sup>2</sup> Benedikt Bünz<sup>1</sup> Percy Liang<sup>1</sup> Leonardo de Moura<sup>3</sup> David L. Dill<sup>1</sup>

## Abstract

We present NeuroSAT, a message passing neural network that learns to solve SAT problems after only being trained as a classifier to predict satisfiability. Although it is not competitive with state-of-the-art SAT solvers, NeuroSAT can solve problems that are substantially larger and more difficult than it ever saw during training by simply running for more iterations. Moreover, NeuroSAT generalizes to novel distributions; after training only on random SAT problems, at test time it can solve SAT problems encoding graph coloring, clique detection, dominating set, and vertex cover problems, all on a range of distributions over small random graphs.

Train:

$$\left\{ \begin{array}{l} \text{Input: SAT problem } P \\ \text{Output: } \mathbb{1} \{P \text{ is satisfiable}\} \end{array} \right\}$$

Test:

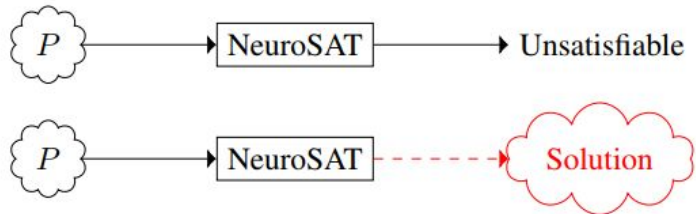


Figure 1. We train NeuroSAT to predict whether SAT problems are satisfiable, providing only a single bit of supervision for each problem. At test time, when NeuroSAT predicts *satisfiable*, we can almost always extract a satisfying assignment from the network’s activations. The problems at test time can also be substantially

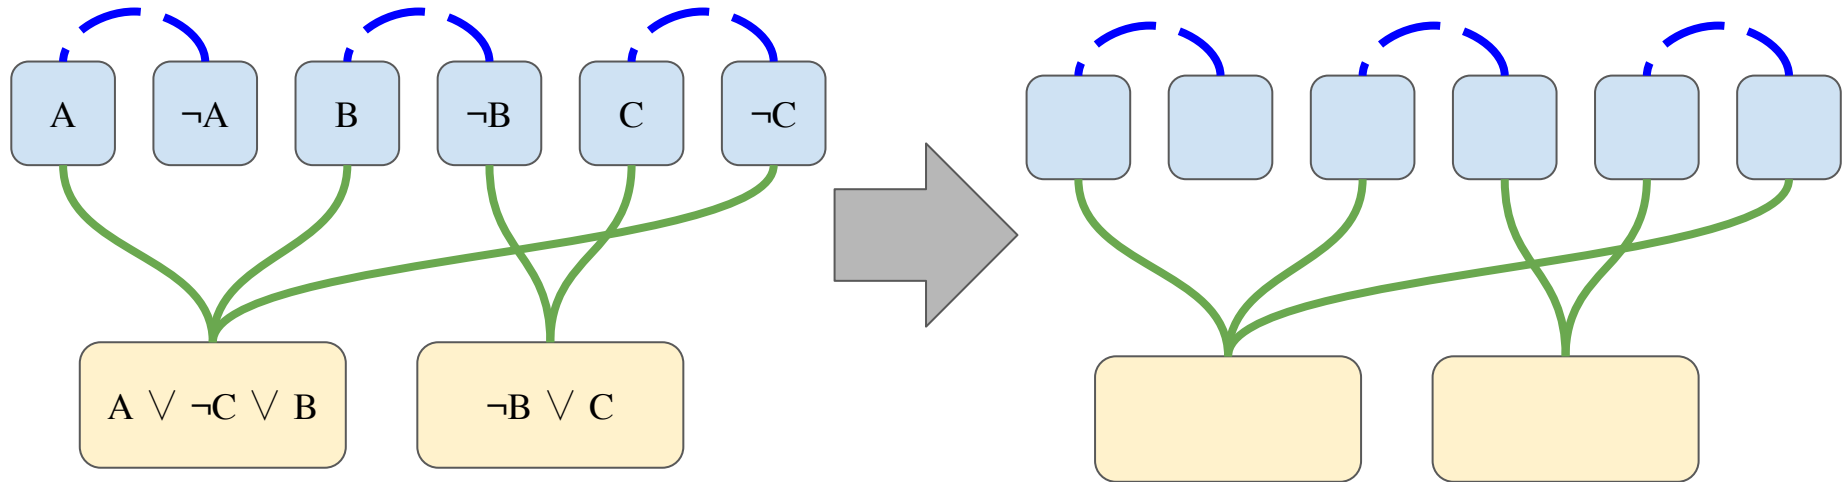
# Axiomatization table for CNF

Axiom	EqNet	LSTM	NeuroSAT
“Variable renaming” - independent	No	No	Yes
“Permutation of literals in clause” - independent	No	No*	Yes
“Permutation of clauses in formula” - independent	No	No	Yes
“Negation of all occurrences of variable” - independent	No	No	Yes

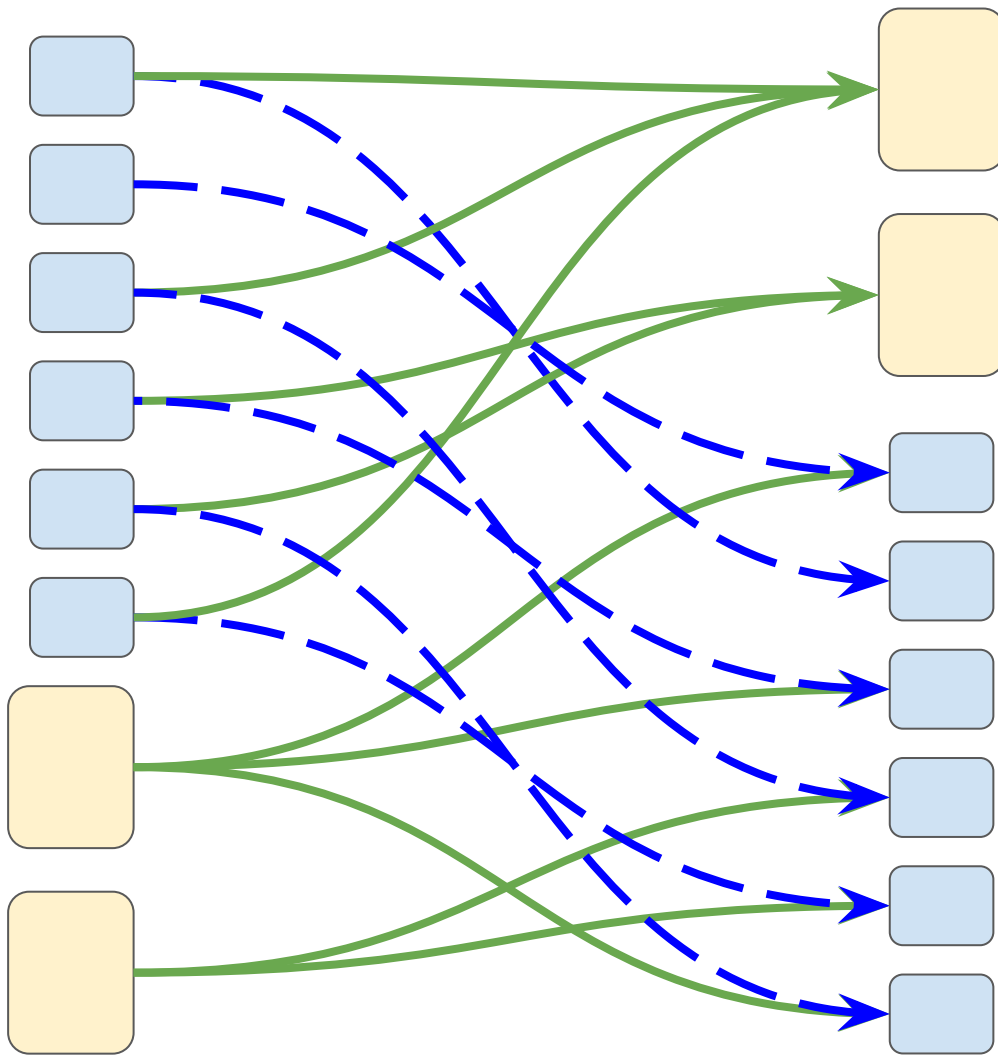
\*Unless Bag of Words is used for in-clause literals aggregation.

# We may forget the names

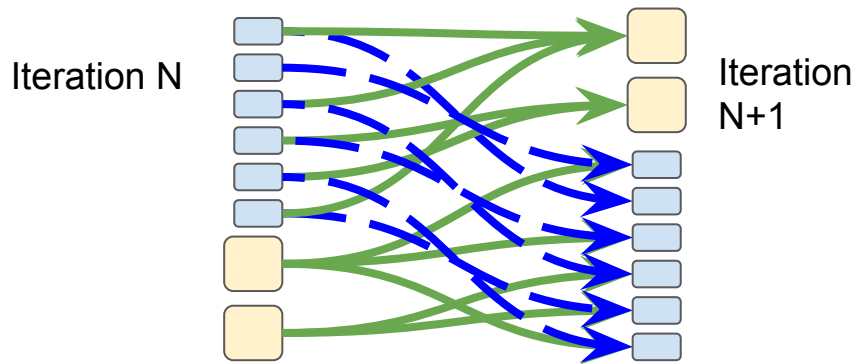
A graph representation without names still has all important information!



Iteration N



Iteration N+1



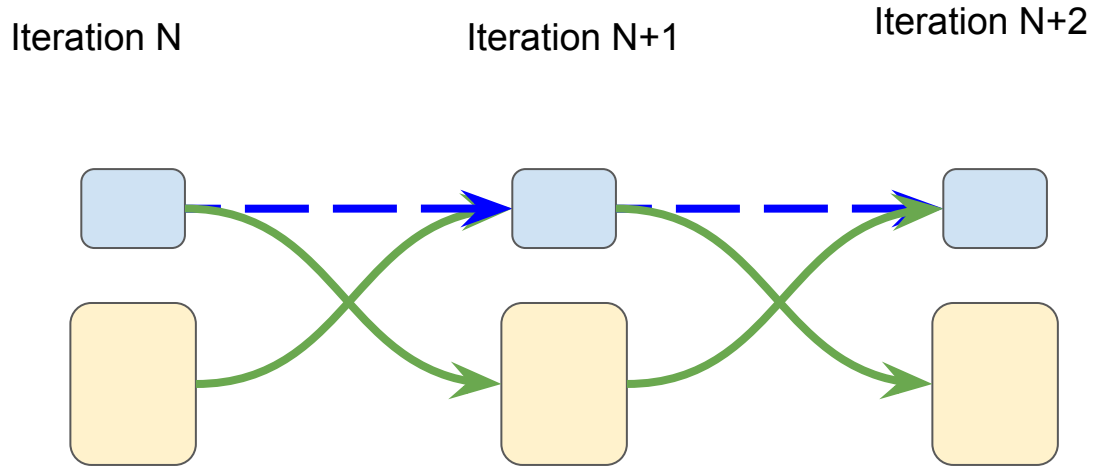
```

elif level >= 1:
    assert_shape(positive_literal_embeddings, [BATCH_SIZE, None, EMBEDDING_SIZE])
    assert_shape(negative_literal_embeddings, [BATCH_SIZE, None, EMBEDDING_SIZE])

    clause_preembeddings = tf.concat([
        tf.divide(
            tf.matmul(positive_connections, positive_literal_embeddings, transpose_a=True) +
            tf.matmul(negative_connections, negative_literal_embeddings, transpose_a=True)
        ), tf.expand_dims(variables_per_clause, -1) + 1.0),
        clause_embeddings], axis=-1)
    assert_shape(clause_preembeddings,
                 [BATCH_SIZE, None, EMBEDDING_SIZE * 2])

    positive_literal_preembeddings = tf.concat([
        tf.divide(
            tf.matmul(positive_connections, clause_embeddings)
        ), tf.expand_dims(clauses_per_variable, -1) + 1.0),
        positive_literal_embeddings,
        negative_literal_embeddings], axis=-1)
    negative_literal_preembeddings = tf.concat([
        tf.divide(
            tf.matmul(negative_connections, clause_embeddings)
        ), tf.expand_dims(clauses_per_variable, -1) + 1.0),
        negative_literal_embeddings,
        positive_literal_embeddings], axis=-1)
    assert_shape(positive_literal_preembeddings, [BATCH_SIZE, None, EMBEDDING_SIZE * 3])

```



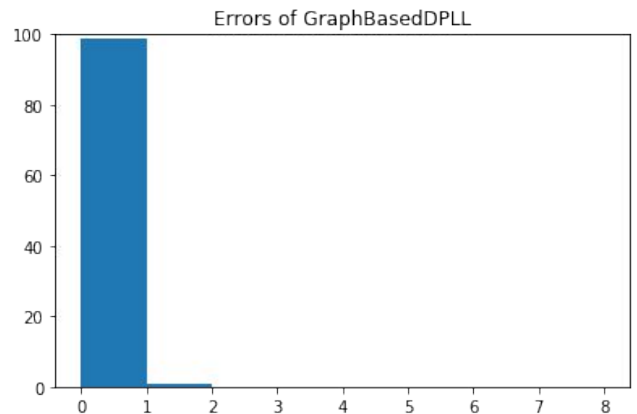
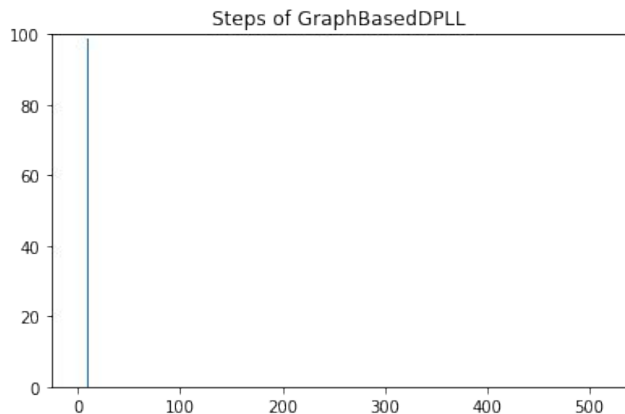
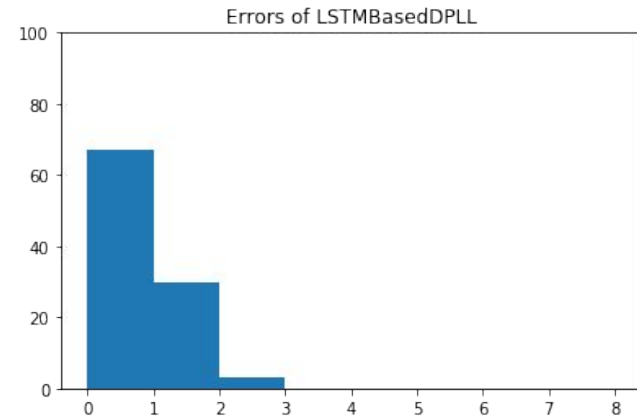
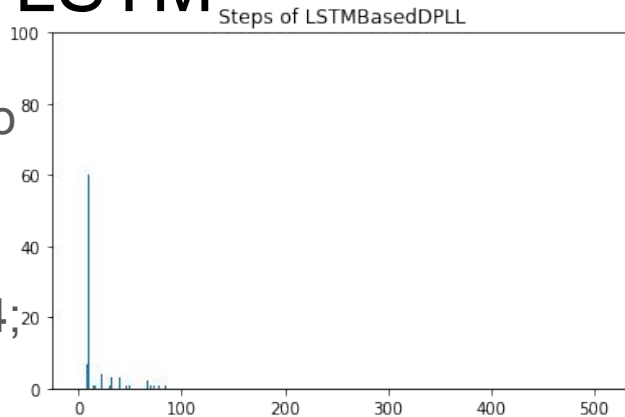
- Architecture idea: embeddings get copied from previous iteration
- + There are blue connections between negating literals
  - + There are green connections between clauses and literals.

# Comparison vs LSTM

#vars: 8, #clauses: up to 50

#Sats: 100; avg step: 19.73; stdev step: 19.94; avg error: 0.36; stdev error: 0.54

#Sats: 100; avg step: 9.01; stdev step: 0.10; avg error: 0.01; stdev error: 0.10



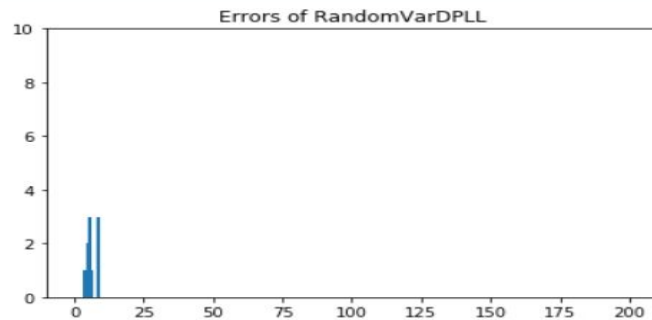
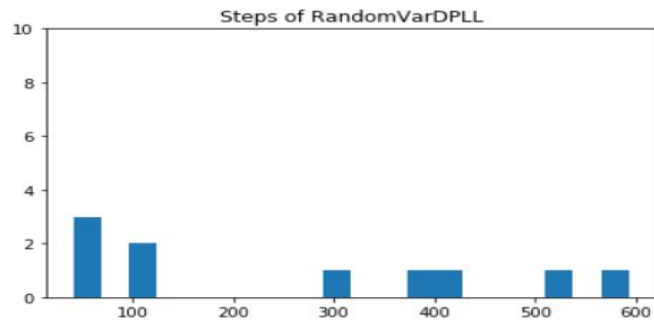


```
In [27]: print_all(10, 200, 12)
```

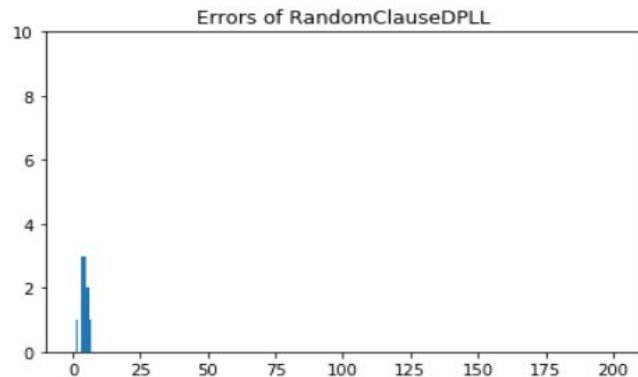
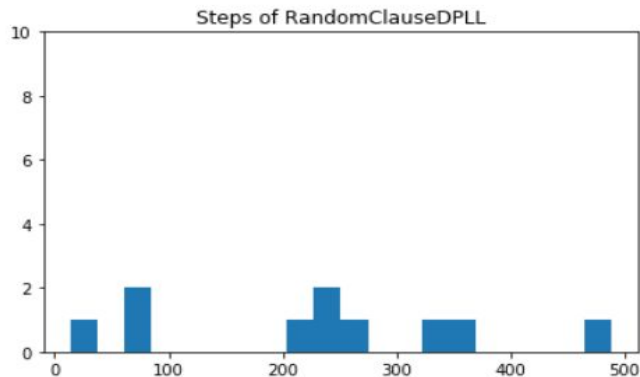
```
100%|██████████| 10/10 [00:00<00:00, 120.11it/s]
```

We have generated 10 formulas

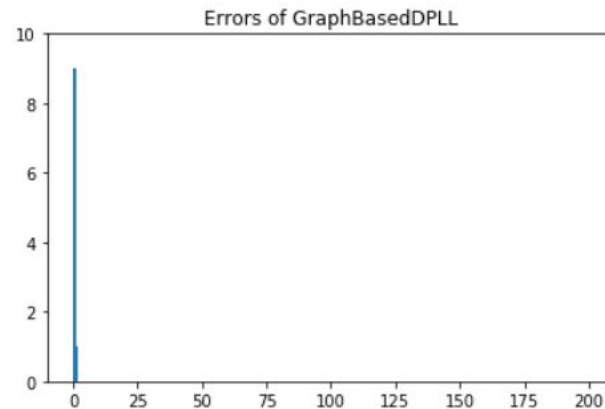
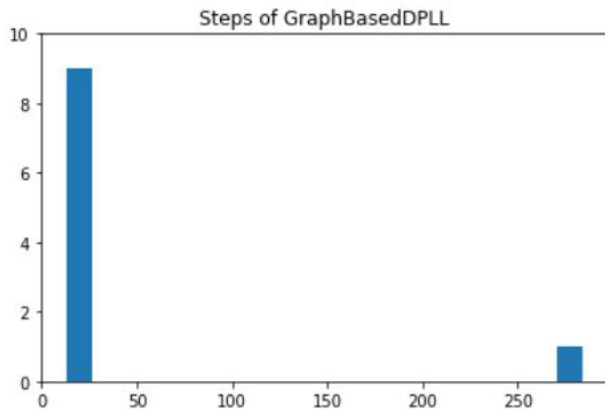
#Sats: 10; avg step: 260.60; stdev step: 200.45; avg error: 5.60; stdev error: 1.74



#Sats: 10; avg step: 227.50; stdev step: 138.15; avg error: 3.80; stdev error: 1.33

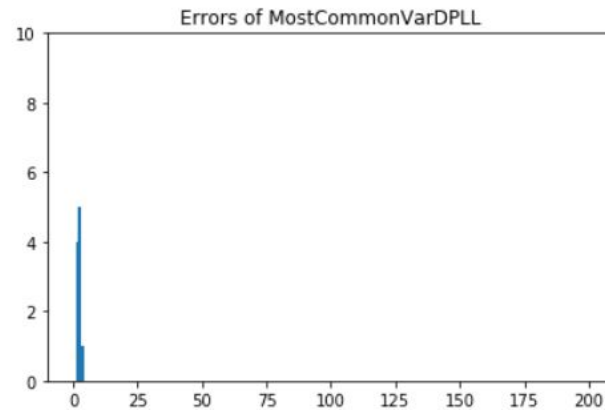
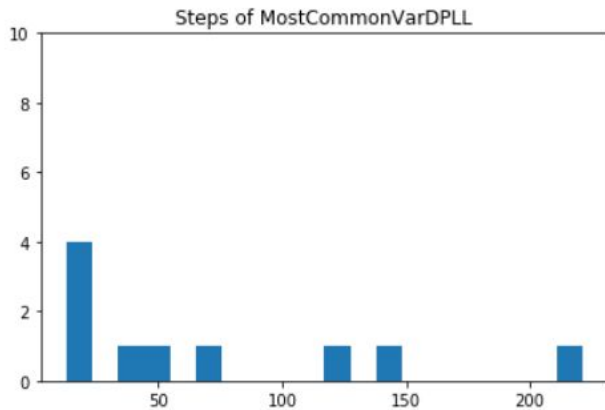


#Sats: 10; avg step: 40.10; stdev step: 81.30; avg error: 0.10; stdev error: 0.30



100% | ██████████ | 10/10 [00:00<00:00, 280.82it/s]

#Sats: 10; avg step: 69.20; stdev step: 65.82; avg error: 1.70; stdev error: 0.64



nodes	samples per hour	samples per node	days to 1e8
8	28693	3587	145
16	58080	3630	72
32	108672	3396	38
64	189504	2961	22
128	352683	2755	12
256	691328	2701	6

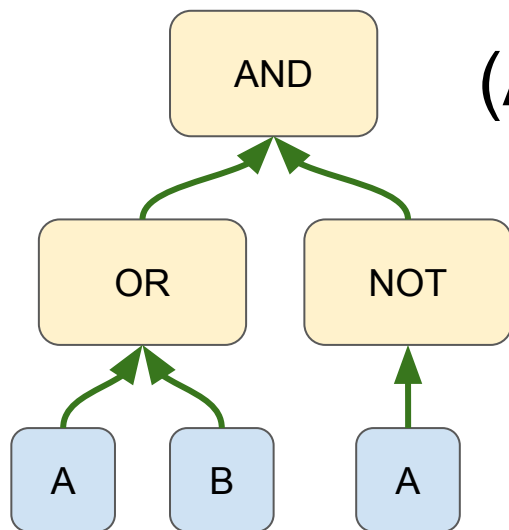
problem	Glucose 3	MiniSat 2.2
SR(10)	0.003	0.002
SR(30)	0.009	0.007
SR(50)	0.021	0.015
SR(70)	0.031	0.024
SR(90)	0.052	0.053
SR(110)	0.171	0.137
SR(130)	0.1238	0.504
SR(150)	6.091	3.406

Figure 2: Performance of our multinode training architecture tested using up to 256 servers each equipped with Xeon E5-2680v3@2,5 GHz. The days column shows how many days of computations are required to generate and train the message passing architecture on 100 million samples drawn from the SR(100) distribution. The table on the right shows the amount of time in seconds required by Glucose 3 [AS09] and Minisat 2.2 [ES03] to solve a randomly sampled problem from SR(n) family (see Section A for more details about SR(n) problems). Solvers were accessed using the PySAT interface [IMM18].

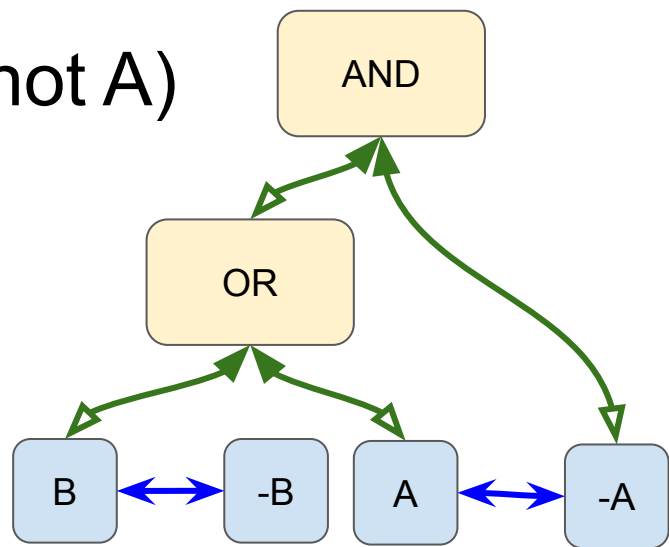
guidance	SR(16)	SR(18)	SR(20)	SR(25)	SR(30)	SR(40)	SR(50)	SR(70)	SR(90)
hybrid SR(8)	<b>99</b>	83	<b>76</b>	57	49	22	23	11	2
hybrid SR(30)	93	89	71	58	42	25	12	6	0
hybrid SR(50)	93	<b>90</b>	65	<b>61</b>	<b>53</b>	<b>46</b>	<b>30</b>	<b>19</b>	<b>16</b>
hybrid SR(100)	94	85	68	58	31	25	6	4	2
mostcommon	<b>99</b>	<b>90</b>	60	26	10	5	3	0	0

*Table 1: Problems solved in 1000 steps using various DPLL heuristics. In the first row we consider guidance trained on the SR(8) dataset and test it against SR( $n$ ) up to  $n = 90$ . In subsequent row we test other instances of guidance. Compare with Table [4](#) in the Appendix.*

# Generalisation of GraphSAT to non-CNF



$(A \text{ or } B) \text{ and } (\text{not } A)$



Eqnet/TreeNN:  
one-way edges.

GraphSAT: two-way edges,  
**symmetric** or **asymmetric**.

Current graphSAT - AND across clauses is implicit.

