

Descending to Complementarity

Martin Suda*

Czech Technical University in Prague, Czech Republic

AITP, September 2023

Strategies in modern Automatic Theorem Provers

- The use of proving strategies is an essential element in high-performance ATPs such as E, iProver, or Vampire.

Strategies in modern Automatic Theorem Provers

- The use of proving strategies is an essential element in high-performance ATPs such as E, iProver, or Vampire.
- They are combined into schedules (sequential or parallel) empirically selected to have complementary powers

Strategies in modern Automatic Theorem Provers

- The use of proving strategies is an essential element in high-performance ATPs such as E, iProver, or Vampire.
- They are combined into schedules (sequential or parallel) empirically selected to have complementary powers

Neural network guided theorem proving

- In recent years, systems like ENIGMA or Deepire were able to dramatically improve the performance of an ATP by integrating neural networks and learning appropriate guiding heuristics

Strategies in modern Automatic Theorem Provers

- The use of proving strategies is an essential element in high-performance ATPs such as E, iProver, or Vampire.
- They are combined into schedules (sequential or parallel) empirically selected to have complementary powers

Neural network guided theorem proving

- In recent years, systems like ENIGMA or Deepire were able to dramatically improve the performance of an ATP by integrating neural networks and learning appropriate guiding heuristics
- However: always only improved a single base strategy!

Strategies in modern Automatic Theorem Provers

- The use of proving strategies is an essential element in high-performance ATPs such as E, iProver, or Vampire.
- They are combined into schedules (sequential or parallel) empirically selected to have complementary powers

Neural network guided theorem proving

- In recent years, systems like ENIGMA or Deepire were able to dramatically improve the performance of an ATP by integrating neural networks and learning appropriate guiding heuristics
- However: always only improved a single base strategy!

➡ Could we train a full host of heuristics using gradient descent?

Strategies in modern Automatic Theorem Provers

- The use of proving strategies is an essential element in high-performance ATPs such as E, iProver, or Vampire.
- They are combined into schedules (sequential or parallel) empirically selected to have complementary powers

Neural network guided theorem proving

- In recent years, systems like ENIGMA or Deepire were able to dramatically improve the performance of an ATP by integrating neural networks and learning appropriate guiding heuristics
 - However: always only improved a single base strategy!
- ➔ Could we train a full host of heuristics using gradient descent?
- Breed specialisations to problems, complementary by design?

- 1 How Does It Work?
- 2 Stateless Clause Selection Reinforcement
- 3 Neural Strategies in Practice
- 4 Conclusion

- 1 How Does It Work?
- 2 Stateless Clause Selection Reinforcement
- 3 Neural Strategies in Practice
- 4 Conclusion

Embedding problems into a latent space of strategies:

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p
- initially unknown: e.g., $v_p = \vec{0}$ for every $p \in \text{Train}$

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p
- initially unknown: e.g., $v_p = \vec{0}$ for every $p \in \text{Train}$
- eventually to represent the “ideal strategy for proving p ”

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p
- initially unknown: e.g., $v_p = \vec{0}$ for every $p \in \text{Train}$
- eventually to represent the “ideal strategy for proving p ”

Condition guidance on v_p :

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p
- initially unknown: e.g., $v_p = \vec{0}$ for every $p \in Train$
- eventually to represent the “ideal strategy for proving p ”

Condition guidance on v_p :

- Instead of $N_\theta(input)$, let's use $N_\theta(input, v_p)$ on p

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p
- initially unknown: e.g., $v_p = \vec{0}$ for every $p \in \text{Train}$
- eventually to represent the “ideal strategy for proving p ”

Condition guidance on v_p :

- Instead of $N_\theta(\text{input})$, let's use $N_\theta(\text{input}, v_p)$ on p
- Gradient formula becomes:

$$\nabla_{\theta, v_p} \text{Loss}(N_\theta(\text{input}, v_p), \text{target})$$

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p
- initially unknown: e.g., $v_p = \vec{0}$ for every $p \in \text{Train}$
- eventually to represent the “ideal strategy for proving p ”

Condition guidance on v_p :

- Instead of $N_\theta(\text{input})$, let's use $N_\theta(\text{input}, v_p)$ on p
- Gradient formula becomes:

$$\nabla_{\theta, v_p} \text{Loss}(N_\theta(\text{input}, v_p), \text{target})$$

- In training, v_p “travels” in the strategy space to encode a specialization of the general guiding heuristic suitable for p

Embedding problems into a latent space of strategies:

- a latent “tweak” variable v_p for every training problem p
- initially unknown: e.g., $v_p = \vec{0}$ for every $p \in \text{Train}$
- eventually to represent the “ideal strategy for proving p ”

Condition guidance on v_p :

- Instead of $N_\theta(\text{input})$, let's use $N_\theta(\text{input}, v_p)$ on p
- Gradient formula becomes:

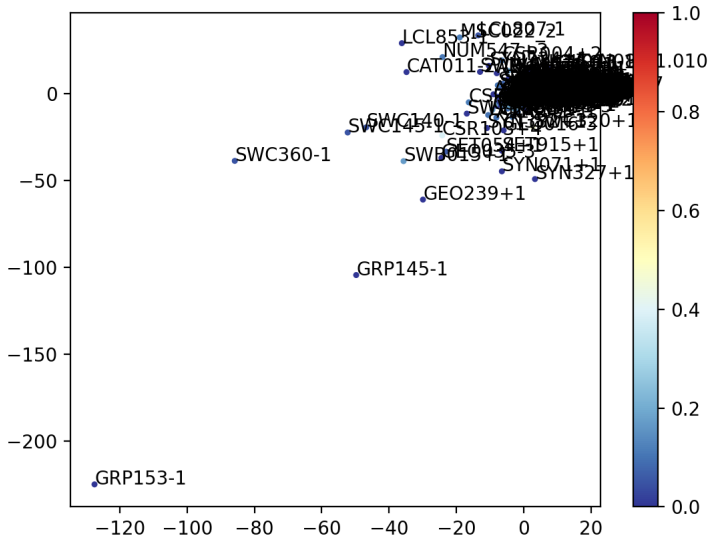
$$\nabla_{\theta, v_p} \text{Loss}(N_\theta(\text{input}, v_p), \text{target})$$

- In training, v_p “travels” in the strategy space to encode a specialization of the general guiding heuristic suitable for p

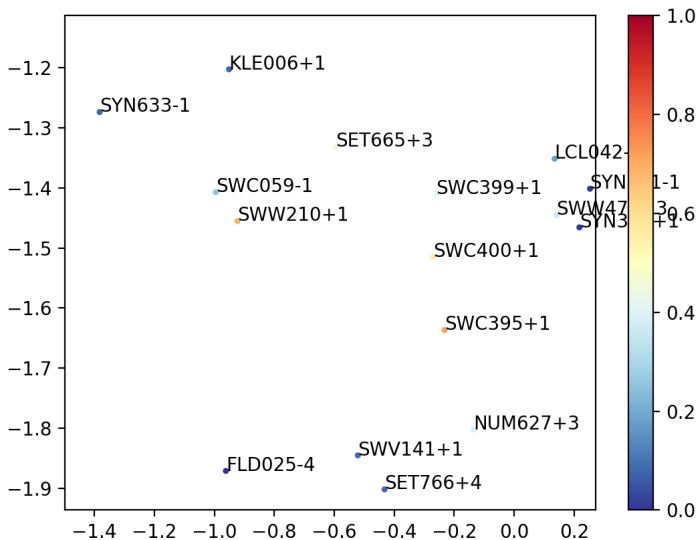
Pick the embedding's dimension well!

A Strategy Map of TPTP?

A Strategy Map of TPTP?



A Strategy Map of TPTP?



- 1 How Does It Work?
- 2 Stateless Clause Selection Reinforcement
- 3 Neural Strategies in Practice
- 4 Conclusion

Clause selection

Clause selection

- arguably the most important choice point in saturation-based proving

Passive → ?givenClause? → *Active*

Clause selection

- arguably the most important choice point in saturation-based proving

Passive → ?givenClause? → *Active*

“What is the next likeliest clause to participate in the proof?”

Clause selection

- arguably the most important choice point in saturation-based proving

Passive → ?givenClause? → *Active*

“What is the next likeliest clause to participate in the proof?”

Traditional clause selection heuristics

- simple criteria: age, weight, ...
- have a priority queue ordering *Passive* for each criterion
- alternate between selecting from the queues using a fixed ratio

Clause selection

- arguably the most important choice point in saturation-based proving

Passive → ?givenClause? → *Active*

“What is the next likeliest clause to participate in the proof?”

Traditional clause selection heuristics

- simple criteria: age, weight, ...
- have a priority queue ordering *Passive* for each criterion
- alternate between selecting from the queues using a fixed ratio

“Thus, there is real synergy in the interleaved strategies, which beat not only the individual components but also their union.”

— Schulz & Möhrmann, 2016

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

Architecture design

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

Architecture design

- simple clause features: age, weight, pos/neg-length, justEq/justNeq, varOcc, goalDist, numSplits

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

Architecture design

- simple clause features: age, weight, pos/neg-length, justEq/justNeq, varOcc, goalDist, numSplits
- the neural part: $\text{MLP}(\text{features}_C) \rightarrow \text{logit}$

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

Architecture design

- simple clause features: age, weight, pos/neg-length, justEq/justNeq, varOcc, goalDist, numSplits
- the neural part: $\text{MLP}(\text{features}_C) \rightarrow \text{logit}$
- stateless: no conjecture dependence, no proof planning

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

Architecture design

- simple clause features: age, weight, pos/neg-length, justEq/justNeq, varOcc, goalDist, numSplits
- the neural part: $\text{MLP}(\text{features}_C) \rightarrow \text{logit}$
- stateless: no conjecture dependence, no proof planning
- yet learning from traces: $\text{Passive}_1, \text{Passive}_2, \dots, \text{Passive}_n$

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

Architecture design

- simple clause features: age, weight, pos/neg-length, justEq/justNeq, varOcc, goalDist, numSplits
- the neural part: $\text{MLP}(\text{features}_C) \rightarrow \text{logit}$
- stateless: no conjecture dependence, no proof planning
- yet learning from traces: $\text{Passive}_1, \text{Passive}_2, \dots, \text{Passive}_n$
- at each step (e.g., step i), the agent is thought to sample:

$$\text{softmax}([\text{MLP}(\text{features}_C)]_{C \in \text{Passive}_i})$$

AITP 2022: Can we, through deep reinforcement learning, somehow re-discover this age/weight-queue alternation?

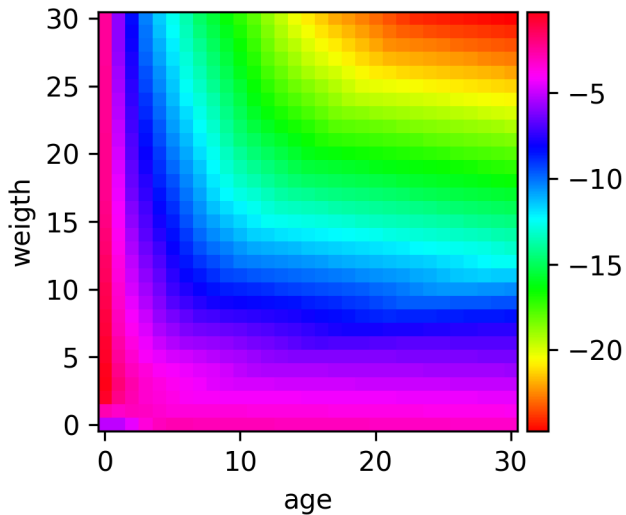
Architecture design

- simple clause features: age, weight, pos/neg-length, justEq/justNeq, varOcc, goalDist, numSplits
- the neural part: $\text{MLP}(\text{features}_C) \rightarrow \text{logit}$
- stateless: no conjecture dependence, no proof planning
- yet learning from traces: $\text{Passive}_1, \text{Passive}_2, \dots, \text{Passive}_n$
- at each step (e.g., step i), the agent is thought to sample:

$$\text{softmax}([\text{MLP}(\text{features}_C)]_{C \in \text{Passive}_i})$$

- “reward”: learn from proof clauses at each step

One queue to rule them all!



- 1 How Does It Work?
- 2 Stateless Clause Selection Reinforcement
- 3 Neural Strategies in Practice**
- 4 Conclusion

Early stopping in RL?

- keep some traces aside for validation purposes
- suddenly, a more careful learning rate does not hurt (reevaluation with the updated policy is the costly bit)

Early stopping in RL?

- keep some traces aside for validation purposes
 - suddenly, a more careful learning rate does not hurt (reevaluation with the updated policy is the costly bit)
- ➔ Huge speedups through this!

Early stopping in RL?

- keep some traces aside for validation purposes
- suddenly, a more careful learning rate does not hurt (reevaluation with the updated policy is the costly bit)

➡ Huge speedups through this!

Maintaining the strategy space embeddings v_p :

- for $p \in Train$: updated with gradient descent

Early stopping in RL?

- keep some traces aside for validation purposes
- suddenly, a more careful learning rate does not hurt (reevaluation with the updated policy is the costly bit)

➡ Huge speedups through this!

Maintaining the strategy space embeddings v_p :

- for $p \in \text{Train}$: updated with gradient descent
- for $p \in \text{Valid}$:

Early stopping in RL?

- keep some traces aside for validation purposes
- suddenly, a more careful learning rate does not hurt (reevaluation with the updated policy is the costly bit)

➡ Huge speedups through this!

Maintaining the strategy space embeddings v_p :

- for $p \in \text{Train}$: updated with gradient descent
- for $p \in \text{Valid}$: try finding best v_p before contributing to loss

Early stopping in RL?

- keep some traces aside for validation purposes
- suddenly, a more careful learning rate does not hurt (reevaluation with the updated policy is the costly bit)

➔ Huge speedups through this!

Maintaining the strategy space embeddings v_p :

- for $p \in Train$: updated with gradient descent
- for $p \in Valid$: try finding best v_p before contributing to loss
- for $p \in Unseen$: ?

Early stopping in RL?

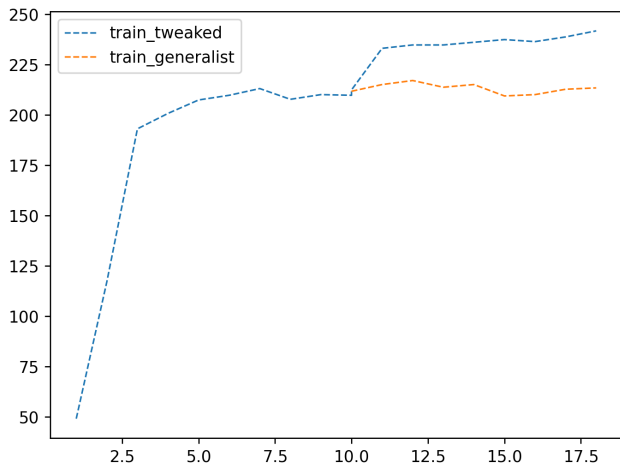
- keep some traces aside for validation purposes
- suddenly, a more careful learning rate does not hurt (reevaluation with the updated policy is the costly bit)

➡ Huge speedups through this!


Maintaining the strategy space embeddings v_p :

- for $p \in Train$: updated with gradient descent
- for $p \in Valid$: try finding best v_p before contributing to loss
- for $p \in Unseen$: ?
 - ➡ Unconditional schedule like with Snake?

How Well Does It Work?



How Well Does the Reward/Loss Capture the “True Game”?



A trace collected from solving the TPTP problem GRA002+1

Summary:

- Neural guiding models for ATPs can be “tweaked” to create targeted proving strategies
- This is basically just gradient descent, but a few aspects require extra care (picking the right dimension, validation, . . .)

Summary:

- Neural guiding models for ATPs can be “tweaked” to create targeted proving strategies
- This is basically just gradient descent, but a few aspects require extra care (picking the right dimension, validation, ...)

Related:

- If we were to add an inference model (i.e., a $IM : p \rightarrow v_p$) then that's what, e.g., ENIGMA is already doing with conjecture-conditioned guidance

Summary:

- Neural guiding models for ATPs can be “tweaked” to create targeted proving strategies
- This is basically just gradient descent, but a few aspects require extra care (picking the right dimension, validation, ...)

Related:

- If we were to add an inference model (i.e., a $IM : p \rightarrow v_p$) then that's what, e.g., ENIGMA is already doing with conjecture-conditioned guidance \Rightarrow These are strategies too!

Summary:

- Neural guiding models for ATPs can be “tweaked” to create targeted proving strategies
- This is basically just gradient descent, but a few aspects require extra care (picking the right dimension, validation, ...)

Related:

- If we were to add an inference model (i.e., a $IM : p \rightarrow v_p$) then that's what, e.g., ENIGMA is already doing with conjecture-conditioned guidance \Rightarrow These are strategies too!

Thank you!