Avoiding distractions in parity games Tom van Dijk - t.vandijk@utwente.nl ISoLA 2024

Outline of the talk

- Parity games
- Distractions
- Tangle learning
- Recursive tangle learning (this work)
- Perspectives

Formal methods

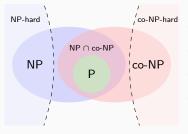
- The model checking problem of modal $\mu\text{-calculus}...$...is equivalent to the problem of solving a parity game.
- The synthesis problem of ω-regular specifications (LTL, etc)...
 ...translates to the problem of solving a parity game.

Formal methods

- The model checking problem of modal $\mu\text{-calculus}...$...is equivalent to the problem of solving a parity game.
- The synthesis problem of ω-regular specifications (LTL, etc)...
 ...translates to the problem of solving a parity game.
- Good news! We can solve large practical parity games very fast! (TvD, TACAS 2018)

WHY PARITY GAMES?

Famous open problem: P vs NP



- P: answer computed in polynomial time
- NP: proof checked in polynomial time
- co-NP: refutation checked in polynomial time
- NP-complete: can simulate every NP problem

Parity games

- Are in NP \cap co-NP; even in UP \cap co-UP
- Since 2017: quasi-polynomial solutions: strictly above polynomial time, and below exponential time
- Goal: a polynomial-time algorithm or a superpolynomial lower bound

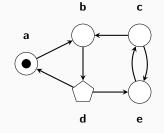
Main algorithm families

- Strategy improvement (policy iteration) little progress in years
- Value iteration underlying universal tree
- Decomposition-based underlying universal tree for some variants
- Universal trees have quasi-polynomial size! lower bound

Tangle learning (our approach)

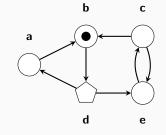
- Based on decomposition with "attractors" (controlled predecessor)
- Not as rigid as Zielonka's recursive algorithm
- Explicitly targets won subgames (tangles)

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



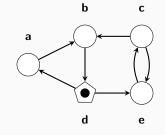
The play π : **a**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



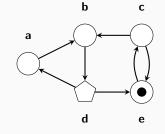
The play π : **a b**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



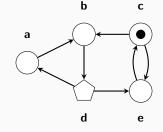
The play π : **a b d**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



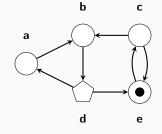
The play π : **a b d e**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



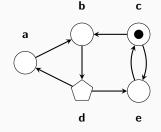
The play π : **a b d e c**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



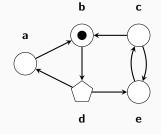
The play π : **a b d e c e**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



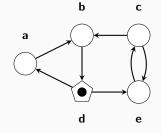
The play π : **a b d e c e c**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



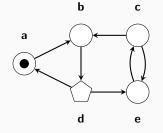
The play π : **a b d e c e c b**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



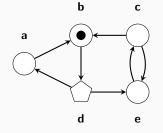
The play π : **a b d e c e c b d**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



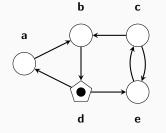
The play π : **a b d e c e c b d a**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



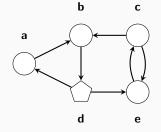
The play π : **a b d e c e c b d a b**

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



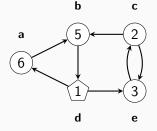
The play π : **a b d e c e c b** (**d a b**)^{ω}

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



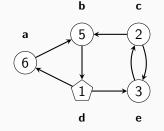
We need a winning condition...

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



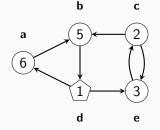
- Each vertex has a priority { 0, 1, 2, ..., d }
- Player Even wins if the highest priority seen infinitely often is even

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



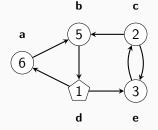
The play π : 6 5 1 3 2 3 2 5 (1 6 5)^{ω} Who wins this play?

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



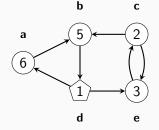
How do we determine who wins a vertex?

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



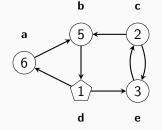
A player wins vertices V if they have a positional strategy $\sigma: V \to V$ such that every play in V consistent with σ is won by that player.

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



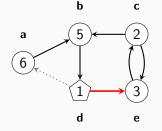
A player wins a (subgame) \mathcal{G} if they have a positional strategy σ such that all cycles in the induced graph $\mathcal{G}[\sigma]$ are won by that player

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



Which vertices are won by which player?

- A parity game is played on a directed graph
- Two players: Even and Odd •
- A play is an infinite path along the edges
- The owner of each vertex chooses the successor



Player Odd wins with strategy $\{ \mathbf{d} \rightarrow \mathbf{e} \}$ Only two (Odd) cycles remain

Solving a parity game

- Compute the winning regions $W_{m 0}$ and $W_{m 0}$
- Compute the winning positional strategies σ_{ullet} and σ_{ullet}

- Systematically investigate possible paths (cycles)
- Keep good paths (even cycles)
- Reject bad paths (odd cycles)

- Systematically investigate possible paths (cycles)
- Keep good paths (even cycles)
- Reject bad paths (odd cycles)
- Sounds simple?

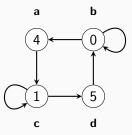
- Systematically investigate possible paths (cycles)
- Keep good paths (even cycles)
- Reject bad paths (odd cycles)
- Sounds simple?
- Follow edges to (good) high even vertices
- Avoid edges to (bad) high odd vertices

- Systematically investigate possible paths (cycles)
- Keep good paths (even cycles)
- Reject bad paths (odd cycles)
- Sounds simple?
- Follow edges to (good) high even vertices
- Avoid edges to (bad) high odd vertices
- ... and (somehow) generalize from edges to paths

- Systematically investigate possible paths (cycles)
- Keep good paths (even cycles)
- Reject bad paths (odd cycles)
- Sounds simple?
- Follow edges to (good) high even vertices
- Avoid edges to (bad) high odd vertices
- ... and (somehow) generalize from edges to paths
- ... and (somehow) generalize from paths to (sets of) cycles

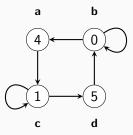
- Systematically investigate possible paths (cycles)
- Keep good paths (even cycles)
- Reject bad paths (odd cycles)
- Sounds simple?
- Follow edges to (good) high even vertices
- Avoid edges to (bad) high odd vertices
- ... and (somehow) generalize from edges to paths
- ... and (somehow) generalize from paths to (sets of) cycles
- BUT...

DISTRACTIONS



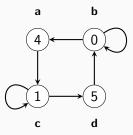
• Find the winning strategy for player Even.

DISTRACTIONS

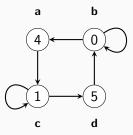


- Find the winning strategy for player Even.
- Pick a random strategy inside the winning region?

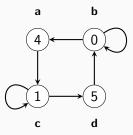
DISTRACTIONS



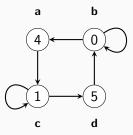
- Find the winning strategy for player Even.
- Pick a random strategy inside the winning region? If you play from **c** to **c** you lose.



- Find the winning strategy for player Even.
- Pick a random strategy inside the winning region? If you play from **c** to **c** you lose.
- Play to nice high Even vertices?



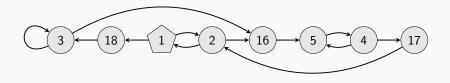
- Find the winning strategy for player Even.
- Pick a random strategy inside the winning region? If you play from **c** to **c** you lose.
- Play to nice high Even vertices? If you play from **b** to **a** you lose.

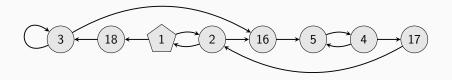


- Find the winning strategy for player Even.
- Pick a random strategy inside the winning region? If you play from **c** to **c** you lose.
- Play to nice high Even vertices? If you play from **b** to **a** you lose.
- Vertex **a** is a distraction for player Even

Intuition

- A distraction is a vertex that makes algorithms require more steps to solve the parity game why? because the algorithm tries paths (assumptions) that turn out to be unproductive
- A distraction is a vertex that *seems* a good target to play to but is actually bad (or less good than other targets)
- A distraction for α is a vertex with an α -priority that the opponent $\overline{\alpha}$ can win if player α always tries to visit it.





 $\begin{array}{l} \mbox{Vertex 16 is a distraction.} \\ \mbox{Play $2 \rightarrow 1$ instead of $2 \rightarrow 16$ to win!} \\ \mbox{But also play $3 \rightarrow 16!} \end{array}$

Wrong assumption: play to 18, 16, 4, 2To win: play to 2, 18 > rest

Identifying distractions

Every algorithm somehow overcomes distractions.

- 1 By showing that distractions lead to bad paths
 - Example: paths from 16 reach 17, so avoid 16
 - Decomposition-based algorithms
- 2 By showing that other vertices reach better paths
 - i.e. give a higher "value" to vertices along good paths (to good vertices)
 - Example: paths from 2 reach 16, so play to 2
 - Value iteration algorithms

Normal tangle learning

- Pure decomposition based
- shows that a distraction leads to bad paths
- does not give higher value to vertices on good paths

Recursive tangle learning (this work)

- Further decomposes each "region" recursively
- shows that a distraction leads to bad paths
- does give higher value to vertices on good paths
 "I now try to play to a vertex inside a region rather than its top vertex"

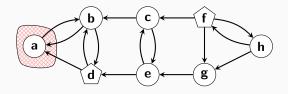
Playing from ${\boldsymbol{\mathsf{A}}}$ to ${\boldsymbol{\mathsf{B}}}$

- From which vertices **A** must a play eventually reach **B**?
- What is the highest vertex that player α can reach?
- Which vertices cannot be avoided by the opponent $\overline{\alpha}$?

Attractor computation

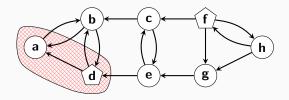
- "Backward reachability" with an opponent
- Given target set Z and a player α , compute all vertices from which player α can ensure arrival in Z
- Add to Z (exhaustively):
 - all vertices of the player α that can play to Z
 - all vertices of opponent $\overline{\alpha}$ that *must* play to Z

Computing the \bigcirc -attractor to **a**



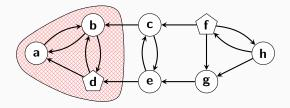
 $\begin{array}{l} \mbox{Attractor set: } \{a\} \\ \mbox{Can attract: } d \mbox{ but not } b \end{array}$

Computing the \bigcirc -attractor to **a**



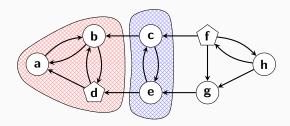
Attractor set: $\{a\}$, $\{a, d\}$ Can attract: **b** but not **e**

Computing the **•**-attractor to **a**



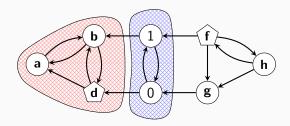
Attractor set: $\{a\}$, $\{a, d\}$, $\{a, b, d\}$ Can attract: neither c nor e

Computing the \bigcirc -attractor to **a**



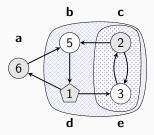
Attractor set: $\{a\}$, $\{a, d\}$, $\{a, b, d\}$ What if player \bullet cannot stay in $\{c, e\}$??

Computing the \bigcirc -attractor to **a**



Attractor set: $\{a\}$, $\{a, d\}$, $\{a, b, d\}$ Every play inside the blue area is won by \bigcirc !

A tangle is a (strongly connected) subgame for which one player has a strategy to win all plays that stay in the subgame.



A game with a 5-tangle and a 3-tangle

A tangle is a (strongly connected) subgame for which one player has a strategy to win all plays that stay in the subgame.

Definition

- A tangle is
 - a pair $T = (U, \sigma)$ where
 - $U \subseteq V$ is a nonempty set of vertices
 - $\sigma: U_{\alpha} \to U$ is a strategy for player $\alpha := \operatorname{pr}(U) \mod 2$ such that
 - player lpha wins all cycles in the induced subgame $\mathcal{G}[U,\sigma]$
 - the induced subgame $\mathcal{G}[U,\sigma]$ is strongly connected

A tangle is a (strongly connected) subgame for which one player has a strategy to win all plays that stay in the subgame.

Properties

- The opponent must escape (or lose inside the tangle)
- The opponent can freely choose any escape (strongly connected)
- A closed tangle (no escapes) is a winning region
- Note: any "won subgame" can be decomposed into tangles

A tangle is a (strongly connected) subgame for which one player has a strategy to win all plays that stay in the subgame.

Tangles are fundamental

- All algorithms implicitly reason about tangles:
 - every algorithm must deal with cycles and nested cyclic structures
 - if you conclude that the opponent cannot 'hide' in some subgame
 - then this **must** be because of tangles
- Most algorithms often explore the same tangle many times! (This can lead to repetitive behavior and exponential blowup!)

A tangle is a (strongly connected) subgame for which one player has a strategy to win all plays that stay in the subgame.

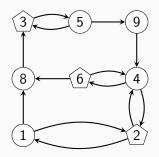
Tangle attractor

Because the opponent $\overline{\alpha}$ must escape the tangle, we can use tangles to attract all vertices of a tangle simultaneously.

Add to Z (exhaustively):

- all vertices of the player α that can play to Z
- all vertices of opponent $\overline{\alpha}$ that *must* play to Z
- all vertices in an α -tangle where all escapes lead to Z

EXAMPLE PARITY GAME



Tangle learning (two-player basic variation)

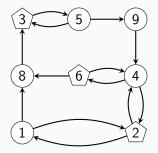
- Assume that every vertex has a unique priority (just to simplify the presentation)
- Partition game into regions with tangle attractor
 - Attract to highest priority
 - Repeat with the remainder until nothing left

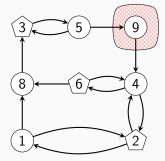
Tangle learning (two-player basic variation)

- Assume that every vertex has a unique priority (just to simplify the presentation)
- Partition game into regions with tangle attractor
 - Attract to highest priority
 - Repeat with the remainder until nothing left
- Every locally closed region contains a <u>new</u> tangle
 - Locally closed if the top vertex is attracted to the region
 - It is a won subgame, but it may not (yet) be strongly connected
 - Run Tarjan's SCC algorithm restricted to the attractor strategy
 - The result is the new tangle

Tangle learning (two-player basic variation)

- Assume that every vertex has a unique priority (just to simplify the presentation)
- Partition game into regions with tangle attractor
 - Attract to highest priority
 - Repeat with the remainder until nothing left
- Every locally closed region contains a <u>new</u> tangle
 - Locally closed if the top vertex is attracted to the region
 - It is a won subgame, but it may not (yet) be strongly connected
 - Run Tarjan's SCC algorithm restricted to the attractor strategy
 - The result is the new tangle
- Every closed tangle is a winning region (dominion)
 - After finding a dominion, maximize it with another attractor
 - Remove maximized dominions from the game

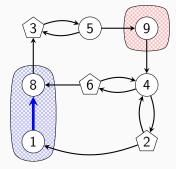




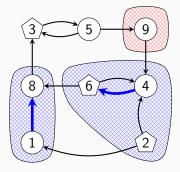
Regions:







Regions:	
• 9	(open)
• 81	(open)



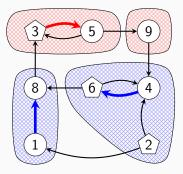
Regions: • 9 • 8 1

• 642

(open) (open)

(closed)

Learned tangles: $\{6, 4 \rightarrow 6\}$



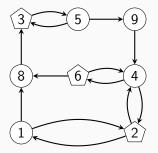
Regions: • 9

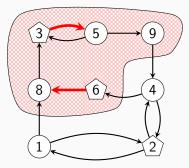
- 81
- 642
- 53

(open) (open)

(closed)

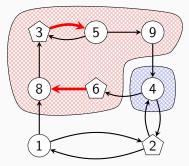
(closed)





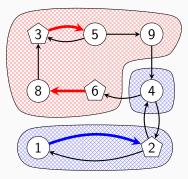
Regions:

• 95386 (open)



Regions:

95386 (open)
4 (open)



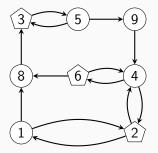
Regions: • 9 5 3 8 6

• 4

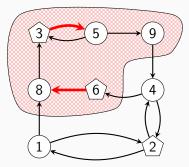
• 21

(open) (open) (closed)

Learned tangles: $\{6, 4 \rightarrow 6\}$, $\{5, 3 \rightarrow 5\}$, $\{2, 1 \rightarrow 2\}$



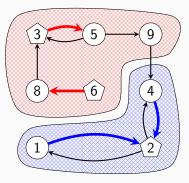
Learned tangles: $\{6, 4 \rightarrow 6\}$, $\{5, 3 \rightarrow 5\}$, $\{2, 1 \rightarrow 2\}$



Regions:

• 9 5 3 8 6 (open)

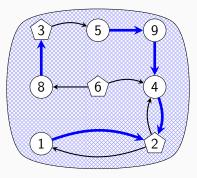
Learned tangles: $\{6, 4 \rightarrow 6\}$, $\{5, 3 \rightarrow 5\}$, $\{2, 1 \rightarrow 2\}$



Regions:

95386 (open)
421 (closed)

Learned tangles: $\{6, 4 \rightarrow 6\}$, $\{5, 3 \rightarrow 5\}$, $\{2, 1 \rightarrow 2\}$, $\{4 \rightarrow 2, 2, 1 \rightarrow 2\}$

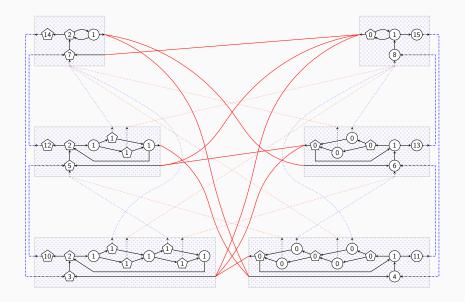


Learned tangles: $\{6, 4 \rightarrow 6\}$, $\{5, 3 \rightarrow 5\}$, $\{2, 1 \rightarrow 2\}$, $\{4 \rightarrow 2, 2, 1 \rightarrow 2\}$

Tangle learning (two-player basic variation)

- Problem! This algorithm has an exponential lower bound.
- Defeated by the two binary counters game

EXAMPLE 3-BIT GAME



EXAMPLE 3-BIT GAME

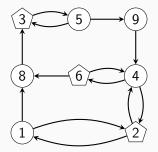
Iteration	Even	Odd	Event
	000	000	Initial state
1	001	000	Set Even 3
2	001	001	Set Odd 3
3	011	000	Set Even 2 (reset Odd 3)
4	010	010	Set Odd 2 (reset Even 3)
5	011	010	Set Even 3
6	011	011	Set Odd 3
7	111	000	Set Even 1 (reset Odd 2, 3)
8	100	100	Set Odd 1 (reset Even 2, 3)
9	101	100	Set Even 3
10	101	101	Set Odd 3
11	111	100	Set Even 2 (reset Odd 3)
12	110	110	Set Odd 2 (reset Even 3)
13	111	110	Set Even 3
14	111	111	Set Odd 3

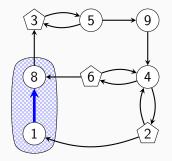
Recursive tangle learning (two-player variation)

- Same as tangle learning, and...
- ...partition every (open) region recursively:
 - "We now want to avoid the top vertex instead"
 - Remove the opponent attractor (inside the region) to the top vertex
 - Partition the remainder of the region

Recursive tangle learning (one-player variation)

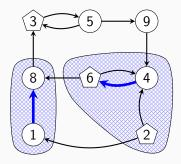
• Only consider the even priority vertices as attractor targets (or only odd priority vertices)





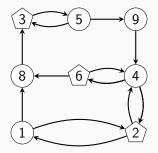
Regions:

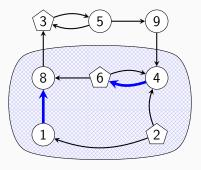
• 8 1 (open) recursive: -



Regions:

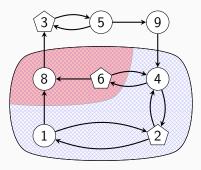
- 8 1 (open) recursive: -
- 6 4 2 (closed)





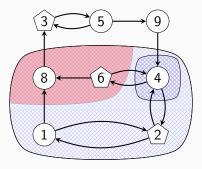
Regions:

• 8 6 4 1 2 (open) recursive: 4 2 1



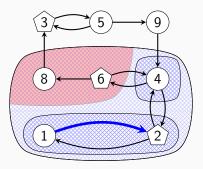
Regions:

• 8 6 4 1 2 (open) recursive: 4 2 1



Regions:

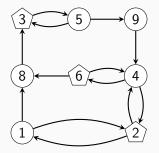
- 8 6 4 1 2 (open) recursive: 4 2 1
- 4 (open) recursive: -



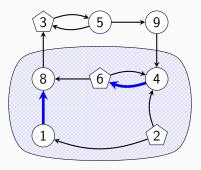
Regions:

- 8 6 4 1 2 (open) recursive: 4 2 1
- 4 (open) recursive: -
- 2 1 (closed)

Learned tangles: {6,4 \rightarrow 6}, {2,1 \rightarrow 2}



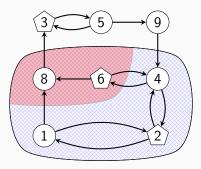
Learned tangles: $\{6, 4 \rightarrow 6\}$, $\{2, 1 \rightarrow 2\}$



Regions:

• 8 6 4 2 1 (open) recursive: 4 2 1

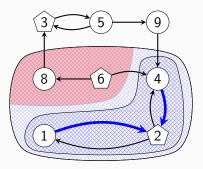
Learned tangles: {6,4 \rightarrow 6}, {2,1 \rightarrow 2}



Regions:

• 8 6 4 2 1 (open) recursive: 4 2 1

Learned tangles: {6,4 \rightarrow 6}, {2,1 \rightarrow 2}



Regions:

8 6 4 2 1 (open) recursive: 4 2 1
4 2 1 (closed)

Learned tangles: {6,4 \rightarrow 6}, {2,1 \rightarrow 2}, {4 \rightarrow 2,2,1 \rightarrow 2}

- Vertex 8 was a distraction
- Two-player tangle learning simply attracts 8 via $\{3,5\}$ to 9
- Recursive tangle learning "avoids" 8 in the recursion

- Vertex 8 was a distraction
- Two-player tangle learning simply attracts 8 via $\{3,5\}$ to 9
- Recursive tangle learning "avoids" 8 in the recursion
- Recursive tangle learning fixes TBC!
- But recursive tangle learning is easily defeated too (see paper)

- TL and RTL have worst-case exponential behavior
- Ongoing work: combining TL with universal tree values (PMTL)
- Ongoing work: heuristics for "currently distractions" (DF*TL)
- Ideas to prove that TL cannot solve in polynomial time
 - Learning tangles moves vertices along the universal tree
 - Every algorithm orders vertices with a universal tree
 - Maybe: the maximal "knowledge" (tangles and path values) per iteration can be characterized, such that the number of steps in the worst case is quasi-polynomial?
 - Could be generalized to classes of algorithms?