




Efficiently Solving Games with Expected Work Search

Owen Randall, Martin Müller, Ting-Han Wei, and Ryan Hayward



Introduction

- Expected Work Search (**EWS**) is a new combinatorial search algorithm
 - Not a search engine like Google or Bing
 - Searches for solutions to problems with many possible positions
- Use **heuristics** to efficiently guide search
 - Can exponentially improve performance
- EWS can outperform existing algorithms
 - **3.8x to 291x** faster when solving Go
 - **2.0x to 3032x** faster when solving Hex





Solving

- Evaluate EWS by **solving** two player perfect information games
 - Specifically Go and Hex
 - Well defined rules
 - Scalable difficulty
- Solving games prove who wins under **optimal play**
 - Any winning move = winning position
 - All losing moves = losing position




Heuristics

- Win rate estimation
 - Monte Carlo Tree Search
 - Prioritizes strong moves
- Proof size estimation
 - Proof number search
 - Prioritizes quickly ending the game
- Both have strengths and weaknesses
- EWS combines both using Expected Work




Expected Work Search

- Build a node tree in memory
- Nodes represent positions
- Each node has a win rate, and Expected Work estimate
- Each iteration of EWS builds the tree, refines node estimates, and backs up proofs




Expected Work Calculation

- How can we estimate the work it will take to solve a node?
- Inductively assume children have expected work values
- Separate cases for winning & losing
- Use MCTS style win rates to estimate winning probability
- Base case expected work values come from random simulations



Expected Work Calculation


- A position is either winning or losing
- Losing case:



Expected Work Calculation




- A position is either winning or losing
- Losing case:




Expected Work Calculation

- A position is either winning or losing
- Losing case:




Expected Work Calculation

- A position is either winning or losing
- Losing case:



Expected Work Calculation


- A position is either winning or losing
- Losing case:
- Sum of winning children's EW



$$EW_{loss}(X) := \sum_{i=1}^n EW_{win}(C_i)$$


Expected Work Calculation

- A position is either winning or losing
- Winning case:




Expected Work Calculation

- A position is either winning or losing
- Winning case:




Expected Work Calculation

- A position is either winning or losing
- Winning case:



Expected Work Calculation

- A position is either winning or losing
- Winning case:
- **Weighted** sum of losing children's EW




$$EW_{win}(X) := \sum_{i=1}^n (EW_{loss}(C_i) \cdot \prod_{j=1}^{i-1} WR(C_j))$$

$$EW_{win}(X) := \sum_{i=0}^n (EW_{loss}(C_i) \cdot \prod_{j=0}^{i-1} WR(C_j)) \quad (2)$$

- Children are searched in order
- Search stops once a losing child is found
- The probability of searching a child = the probability none of the previous children are losing


Win Rate 0.8

Search Prob. 1.0




$$EW_{win}(X) := \sum_{i=0}^n (EW_{loss}(C_i) \cdot \prod_{j=0}^{i-1} WR(C_j)) \quad (2)$$


- Children are searched in order
- Search stops once a losing child is found
- The probability of searching a child = the probability none of the previous children are losing



$$EW_{win}(X) := \sum_{i=0}^n (EW_{loss}(C_i) \cdot \prod_{j=0}^{i-1} WR(C_j)) \quad (2)$$


- Children are searched in order
- Search stops once a losing child is found
- The probability of searching a child = the probability none of the previous children are losing





Selection

- Traverse the search tree
- Select the best child according to the move ordering
- If the selected child has been expanded, continue selection
- Otherwise, expand the leaf node



Child Ordering



- Sort children with:

$$EW_{loss}(A)/(1-WR(A)) < EW_{loss}(B)/(1-WR(B))$$


$$\iff$$

$$A < B$$

- Best first move ordering
- Move ordering of children determines EW_{win}
 - **This ordering minimizes EW_{win}**


Child Ordering

$$EW_{loss}(A)/(1-WR(A)) < EW_{loss}(B)/(1-WR(B))$$



Child Ordering


$$EW_{loss}(A)/(1-WR(A)) < EW_{loss}(B)/(1-WR(B))$$



WR	0.6	0.2	0.15
EW_{loss}	900	1000	1500
$EW/(1-WR)$	2250	1250	1765

Child Ordering

$$EW_{loss}(A)/(1-WR(A)) < EW_{loss}(B)/(1-WR(B))$$




Selection

- Traverse the search tree
- Select the best child according to the move ordering
- If the selected child has been expanded, continue selection
- Otherwise, expand the leaf node


Algorithm 1 SelectBackpropagate X : returns whether X is solved and if so whether X is winning

```
1:  $C := X.children[0]$  ▷ Selection
2: Make move  $C.move$ 
3: if  $C.expanded$  then
4:    $isSolved, isWinning = SelectBackpropagate(C)$ 
5: else
6:    $isSolved, isWinning = Expand(C)$ 
7: Undo move  $C.move$  ▷ Backpropagation
8: if  $isSolved$  and  $isWinning$  then
9:   Remove  $C$  from  $X.children$ 
10:  if  $X.children$  is empty then
11:    return true, false ▷ Solved loss
12:  else if  $isSolved$  and not  $isWinning$  then
13:    return true, true ▷ Solved win
14: Sort  $X.children$  with (3)
15: Update  $X.winRate$ 
16: Update  $X.expectedWorkLoss$  with (1)
17: Update  $X.expectedWorkWin$  with (2)
18: return false, false ▷ Unsolved
```



Expansion


- Create new nodes
- Check for terminal children
- Initialize WR and EW with random simulation results
- Return if the expanded node is proven



Initial Evaluation

- Random simulations
- Win rates = wins / visits
- Initialize EW as the sum of the branching factors of positions in random simulations

$$EW_0(X) := \sum_{i=0}^m b(P_i)$$




Expansion



- Create new nodes
- Check for terminal children
- Initialize WR and EW with random simulation results
- Return if the expanded node is proven


Algorithm 2 Expand X : returns whether X is solved and if so whether X is winning

```
1:  $X$ .expanded := true
2: for all legal moves  $m$  do
3:   if  $m$  is a terminal winning move for  $X$  then
4:     return true, true ▷ Solved win
5:   else if  $m$  is not a terminal losing move for  $X$  then
6:     Create node  $C$ 
7:      $C$ .expanded = false
8:      $C$ .move :=  $m$ 
9:     Evaluate  $C$ .winRate
10:    Evaluate  $C$ .expectedWorkLoss
11:    Evaluate  $C$ .expectedWorkWin
12:    Add  $C$  to  $X$ .children
13: if  $X$ .children is empty then
14:   return true, false ▷ Solved loss
15: else
16:   return false, false ▷ Unsolved
```



Backpropagation

- Back up new information along the path chosen by selection:
- Win rate
- Expected Work
- Proofs
- Proven nodes are pruned



Backpropagation

- Back up new information along the path chosen by selection:
- Win rate
- Expected Work
- Proofs
- Proven nodes are pruned


Algorithm 1 SelectBackpropagate X : returns whether X is solved and if so whether X is winning

```
1:  $C := X.children[0]$  ▷ Selection
2: Make move  $C.move$ 
3: if  $C.expanded$  then
4:    $isSolved, isWinning = SelectBackpropagate(C)$ 
5: else
6:    $isSolved, isWinning = Expand(C)$ 
7: Undo move  $C.move$  ▷ Backpropagation
8: if  $isSolved$  and  $isWinning$  then
9:   Remove  $C$  from  $X.children$ 
10:  if  $X.children$  is empty then
11:    return true, false ▷ Solved loss
12: else if  $isSolved$  and not  $isWinning$  then
13:   return true, true ▷ Solved win
14: Sort  $X.children$  with (3)
15: Update  $X.winRate$ 
16: Update  $X.expectedWorkLoss$  with (1)
17: Update  $X.expectedWorkWin$  with (2)
18: return false, false ▷ Unsolved
```



Results






- Compared against different independent solving implementations
 - GoSolver, MIGOS, Enhanced AB, Morat PNS, Morat MCTS
- Ablation study
 - Removed proof-size / win-rate information
 - Performance drops
- Evaluated a variety of Go positions
 - Strong general performance
 - New results on empty boards
- Evaluated empty $n \times n$ Hex
 - Strong results with and without knowledge



Program	Av. solve time (s)	# Faster than EWS	# Timeouts
EWS	2.32	-	-
Go-Solver	22.63	31 (5.2%)	11 (1.8%)




Program	Av. solve time (s)	# Faster than EWS	# Timeouts
EWS	2.32	-	-
Go-Solver	22.63	31 (5.2%)	11 (1.8%)
EWS-WR	19.70	96 (16.0%)	11 (1.8%)



Program	Av. solve time (s)	# Faster than EWS	# Timeouts
EWS	2.32	-	-
Go-Solver	22.63	31 (5.2%)	11 (1.8%)
EWS-WR	19.70	96 (16.0%)	11 (1.8%)
EWS-PS	19.58	171 (28.5%)	14 (2.3%)

Solving Square Go Boards


- Compare against:
 - EWS-WR (without win rate estimation)
 - EWS-PS (without proof size estimation)



Program	3×3 Time (s)	3×3 Nodes	4×4 Time (s)	4×4 Nodes	5×5 Time (h)	5×5 Nodes
EWS	0.053	161	13.626	495,494	5.252	2,605,781,360
EWS-WR	0.074	796	40.396	1,562,718	> 24	-
EWS-PS	0.070	232	18.320	744,169	> 24	-

Solving Square Go Boards


- Compare against:
 - EWS-WR (without win rate estimation)
 - EWS-PS (without proof size estimation)
 - Go-Solver (iterative deepening alpha beta)
 - MIGOS (state of the art using Japanese rules)



Program	3×3 Time (s)	3×3 Nodes	4×4 Time (s)	4×4 Nodes	5×5 Time (h)	5×5 Nodes
EWS	0.053	161	13.626	495,494	5.252	2,605,781,360
EWS-WR	0.074	796	40.396	1,562,718	> 24	-
EWS-PS	0.070	232	18.320	744,169	> 24	-
Go-Solver	1.299	1,628	51.522	799,607	> 24	-
MIGOS SSK	< 3.3	~ 25,118	~3,960	~3,162,277,660	-	-

Solving Square Hex Boards


- Comparing against:
 - An enhanced alpha beta solver we developed
 - Morat Monte Carlo Tree Search
 - Morat Proof Number Search



Program	4×4 Time (s)	4×4 Nodes	5×5 Time (s)	5×5 Nodes	6×6 Time (h)	6×6 Nodes
EWS	0.002	283	0.253	37,034	0.422	93,963,192
Enhanced AB	0.004	1,673	3.935	698,402	> 24	-
Morat MCTS	0.035	4,644	29.669	3,554,546	> 24	-
Morat PNS	0.054	8,871	758.102	154,539,591	> 24	-


Solving with Hex Knowledge

- Virtual Connections
- Strategies to guarantee connections
- Full board connections determine the winner
- Intersections of first player VCs form a must-play zone





Solving with Hex Knowledge

- Dead cells / Inferior cell analysis
- Some moves are irrelevant to the game's theoretical outcome
- Dead cells can be filled in, inferior moves can be ignored
- Found with pattern matching



Solving with Hex Knowledge


- Significantly reduces search space
- For 6x6 Hex: 93,963,192 nodes -> 26 nodes
- 8x8 can be solved
- Search is much slower, but more node efficient



Program	6×6 Time (s)	6×6 Nodes	7×7 Time (s)	7×7 Nodes	8×8 Time (s)	8×8 Nodes
EWS with knowledge	0.005	26	0.588	2,318	234.449	555,158

Conclusion

- Expected Work Search combines proof size and win rate heuristics
 - Balances these heuristics
 - Covers previous weaknesses
- Strong results on Go and Hex
 - New results solving 5x5 Go with pos. superko
 - More Hex knowledge is needed for S.O.T.A results
- Plenty of future work for EWS



Thank you for your time!

Questions?

References

- [1] ALLIS, L., HERIK, H., AND HUNTJENS, M. Go-Moku and Threat-Space Search. *Computational Intelligence* 12 (1994), pp. 7–23.
- [2] ALLIS, L., VAN DER MEULEN, M., AND VAN DEN HERIK, H. Proof-Number Search. *Artificial Intelligence* 66, 1 (1994), pp. 91–124.
- [3] ANSHELEVICH, V. V. The Game of Hex: An Automatic Theorem Proving Approach to Game Programming. In *AAAI/IAAI* 17 (2000), pp. 189–294.
- [4] ANSHELEVICH, V. V. A Hierarchical Approach to Computer Hex. *Artificial Intelligence* 134, 1 (2002), pp. 101–120.
- [5] ARNESON, B., HAYWARD, R., AND HENDERSON, P. Solving Hex: Beyond Humans. In *Computers and Games* (2010), pp. 1–10.
- [6] BENSON, D. Life in the Game of Go. *Information Sciences* 10 (1976), pp. 17–29.
- [7] BROWNE, C. B., POWLEY, E., WHITEHOUSE, D., LUCAS, S. M., COWLING, P. I., ROHLFSHAGEN, P., TAVENER, S., PEREZ, D., SAMOTHRAKIS, S., AND COLTON, S. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games* 4, 1 (2012), pp. 1–43.
- [8] CAMPBELL, M. The Graph-History Interaction: On Ignoring Position History. In *ACM Annual Conference on the Range of Computing* (1985), pp. 278–280.
- [9] DOE, E., WINANDS, M. H. M., KOWALSKI, J., SOEMERS, D. J. N. J., GÓRSKI, D., AND BROWNE, C. Proof Number Based Monte-Carlo Tree Search. In *IEEE Conference on Games* (2022), pp. 206–212.
- [10] DU, H., WEI, T. H., AND MÜLLER, M. Solving NoGo on Small Rectangular Boards. In *Advances in Computer Games* (2024), Springer Nature Switzerland, pp. 39–49.
- [11] ENZENBERGER, M., MÜLLER, M., ARNESON, B., AND SEGAL, R. Fuego—An Open-Source Framework for Board Games and Go Engine Based on Monte Carlo Tree Search. *IEEE Transactions on Computational Intelligence and AI in Games* 2 (2011), pp. 259–270.
- [12] EWALDS, T. V. Playing and Solving Havannah, 2012. *University of Alberta Master's Thesis* (2012), <https://github.com/tewalds/morat>.
- [13] HAYWARD, R., BJÖRNSSON, Y., JOHANSON, M., KAN, M., PO, N., AND VAN RIJSWIJCK, J. Solving 7×7 Hex: Virtual Connections and Game-State Reduction. *Advances in Computer Games: Many Games, Many Challenges* (2004), pp. 261–278.
- [14] HENDERSON, P., ARNESON, B., AND HAYWARD, R. Solving 8x8 Hex. *IJCAI International Joint Conference on Artificial Intelligence* (2009), pp. 505–510.
- [15] KISHIMOTO, A., AND MÜLLER, M. A General Solution to the Graph History Interaction Problem. In *Proceedings of the 19th National Conference on Artificial Intelligence* (2004), pp. 644–649.
- [16] KISHIMOTO, A., WINANDS, M., MÜLLER, M., AND SAITO, J. T. Game-Tree Search Using Proof Numbers: The First Twenty Years. *ICGA Journal* 35 (2012), pp. 131–156.
- [17] KNUTH, D. E., AND MOORE, R. W. An Analysis of Alpha-Beta Pruning. *Artificial Intelligence* 6, 4 (1975), pp. 293–326.
- [18] KOCSIS, L., AND SZEPEŠVÁRI, C. Bandit Based Monte-Carlo Planning. In *Machine Learning: ECML* (2006), pp. 282–293.
- [19] MÜLLER, M. Playing it Safe: Recognizing Secure Territories in Computer Go by Using Static Rules and Search. In *Game Programming Workshop* (1997), pp. 80–86.
- [20] NAGAI, A. DF-PN Algorithm for Searching AND/OR Trees and its Applications. *University of Tokyo Ph.D. Thesis* (2002).
- [21] NETO, J. P., AND TAYLOR, W. Game Mutators for Restricting Play. *Game & Puzzle Design* (2015), pp. 64–65.
- [22] NIU, X., KISHIMOTO, A., AND MÜLLER, M. Recognizing Seki in Computer Go. In *Advances in Computer Games. 11th International Conference* (2006), vol. 4250 of *Lecture Notes in Computer Science*, Springer, pp. 88–103.
- [23] NIU, X., AND MÜLLER, M. An Improved Safety Solver for Computer Go. In *Computers and Games: 4th International Conference* (2006), vol. 3846 of *Lecture Notes in Computer Science*, Springer, pp. 97–112.
- [24] NIU, X., AND MÜLLER, M. An Open Boundary Safety-of-Territory Solver for the game of Go. In *Computers and Games. 5th International Conference* (2007), vol. 4630 of *Lecture Notes in Computer Science*, Springer, pp. 37–49.
- [25] NIU, X., AND MÜLLER, M. An Improved Safety Solver in Go using Partial Regions. In *Computers and Games. 6th International Conference* (2008), vol. 5131 of *Lecture Notes in Computer Science*, Springer, pp. 102–112.

- [26] PAWLEWICZ, J., AND HAYWARD, R. B. Scalable Parallel DFPN Search. In *Computers and Games* (2013), pp. 138–150.
- [27] PAWLEWICZ, J., AND LEW, L. Improving Depth-First PN-Search: 1 + Epsilon Trick. In *Proceedings of the 5th International Conference on Computers and Games* (2006), pp. 160–171.
- [28] RANDALL, O., WEI, T.-H., HAYWARD, R., AND MÜLLER, M. Improving Search in Go Using Bounded Static Safety. In *Computers and Games* (2022), pp. 14–23.
- [29] ROSIN, C. Multi-armed Bandits with Episode Context. *Annals of Mathematics and Artificial Intelligence* 61 (2010), pp. 203–230.
- [30] SAFFIDINE, A., AND CAZENAVE, T. Developments on Product Propagation. In *Computers and Games* (2013), pp. 100–109.
- [31] SCHAEFFER, J., BURCH, N., BJÖRNSSON, Y., KISHIMOTO, A., MÜLLER, M., LAKE, R., LU, P., AND SUTPHEN, S. Checkers is Solved. *Science* 317, 5844 (2007), pp. 1518–1522.
- [32] SCHAEFFER, J., AND PLAAT, A. New Advances in Alpha-Beta Searching. In *Proceedings of the 1996 ACM 24th annual Conference on Computer Science* (1996), pp. 124–130.
- [33] SHIH, C.-C., WU, T.-R., WEI, T. H., AND WU, I.-C. A Novel Approach to Solving Goal-Achieving Problems for Board Games. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence* (2021), pp. 10362–10369.
- [34] SILVER, D., HUANG, A., MADDISON, C., GUEZ, A., SIFRE, L., DRIESCHKE, G., SCHRITTWIESER, J., ANTONOGLOU, I., PANNEERSHELVAM, V., LANCTOT, M., DIELEMAN, S., GREWE, D., NHAM, J., KALCHBRENNER, N., SUTSKEVER, I., LILLICRAP, T., LEACH, M., KAVUKCUOGLU, K., GRAEPEL, T., AND HASSABIS, D. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* 529 (2016), pp. 484–489.
- [35] SILVER, D., HUBERT, T., SCHRITTWIESER, J., ANTONOGLOU, I., LAI, M., GUEZ, A., LANCTOT, M., SIFRE, L., KUMARAN, D., GRAEPEL, T., LILLICRAP, T., SIMONYAN, K., AND HASSABIS, D. A General Reinforcement Learning Algorithm that Masters Chess, Shogi, and Go through Self-Play. *Science* 362, 6419 (2018), pp. 1140–1144.
- [36] SLATE, D., AND ATKIN, L. Chess Skill in Man and Machine. In *Springer-Verlag* (1977), pp. 82–118.
- [37] SONG, J., AND MÜLLER, M. An Enhanced Solver for the Game of Amazons. *IEEE Transactions on Computational Intelligence and AI in Games* 7 (2015), pp. 16–27.
- [38] SUTTON, R. S., AND BARTO, A. G. Reinforcement Learning: An Introduction. *MIT Press* (1998), pp. 39–40.
- [39] VAN DER WERF, E., HERIK, H., AND UITERWIJK, J. Solving Go on Small Boards. *ICGA Journal* 26 (2003), pp. 92–107.
- [40] VAN DER WERF, E., AND WINANDS, M. Solving Go for Rectangular Boards. *ICGA Journal* 32, 2 (2009), pp. 77–88.
- [41] WAGNER, J., AND VIRÁG, I. Solving Renju. *Joint International Computer Games Association* 24 (2001), pp. 30–35.
- [42] WINANDS, M. H. M., BJÖRNSSON, Y., AND SAITO, J.-T. Monte-Carlo Tree Search Solver. In *Computers and Games* (2008), pp. 25–36.
- [43] WU, T.-R., SHIH, C.-C., WEI, T.-H., TSAI, M.-Y., HSU, W.-Y., AND WU, I.-C. AlphaZero-based Proof Cost Network to Aid Game Solving. In *International Conference on Learning Representations* (2022).
- [44] ZOBRIST, A. A New Hashing Method with Application for Game Playing. *ICGA Journal* 13 (1990), pp. 69–73.