WallZero: Mastering the Game of WallGo with Strategic Analysis

Hsing-Yu Chen[1,2], Jérôme Arjonilla[2], I-Chen Wu[1], Ti-Rong Wu[2]

[1] National Yang Ming Chiao Tung University
[2] Academia Sinica

The Game of WallGo

- A New Board Game popularized by Netflix's *The Devil's Plan* in 2025
  - Lee Sedol participated in this show
  - Attracted attention from Go and AI communities
- Played on a 7×7 board by 2–4 players
- Estimated game-tree complexity ≈ 10^{87}

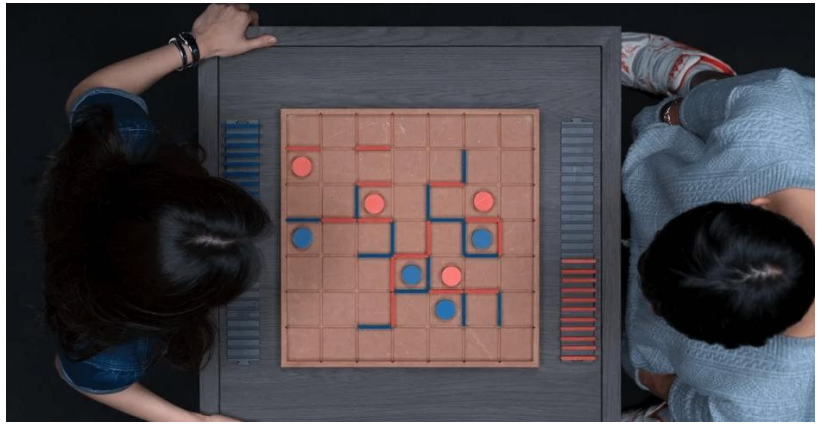

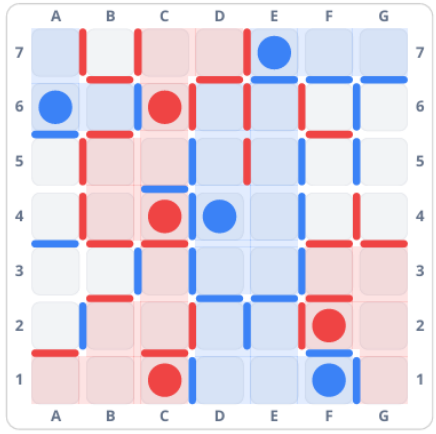

Game	Game Tree Complexity
Othello 8x8	10^{58}
Gomoku 15x15	10^{70}
WallGo 7x7	10^{87}
Hex 11x11	10^{98}
Chess 8x8	10^{123}
Go 19x19	10^{360}

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
  - There are eight stones in total

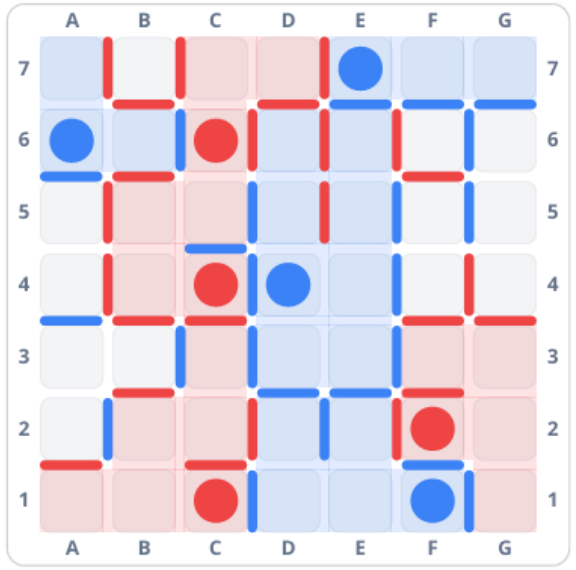

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

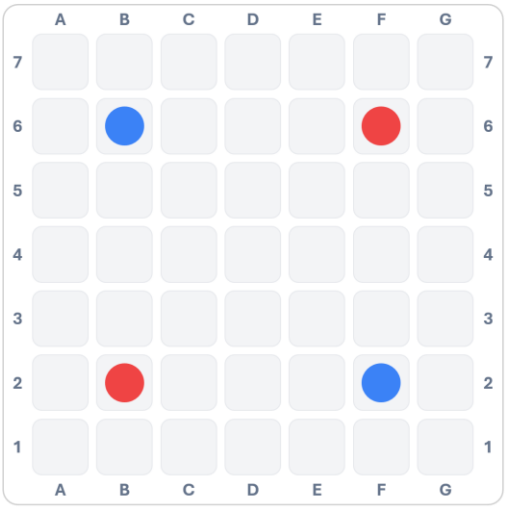

4-stone Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

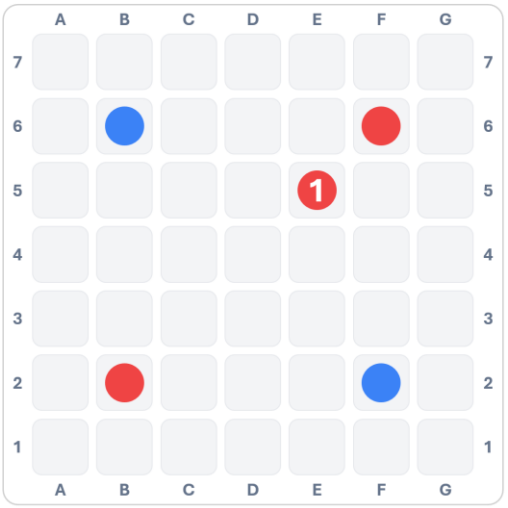

4-stone Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

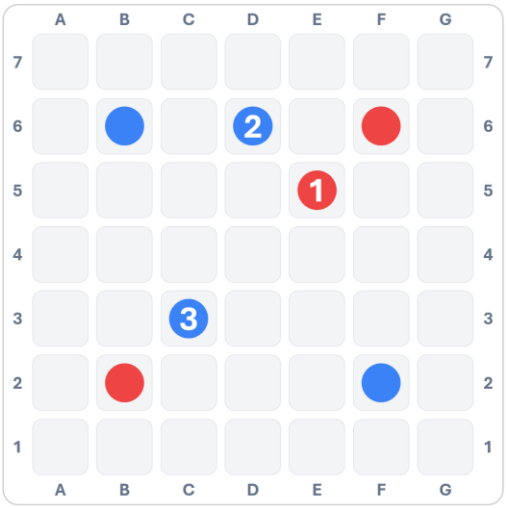

4-stone Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

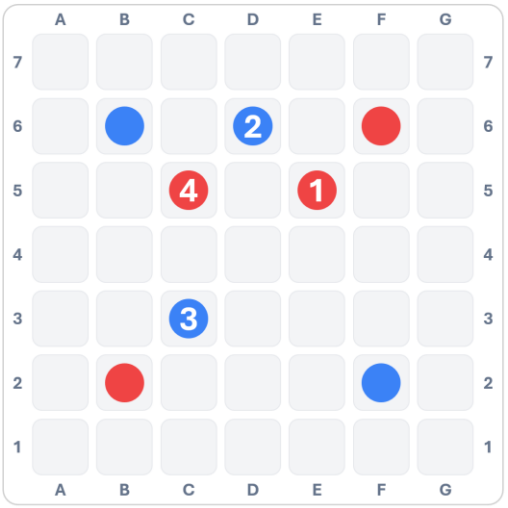

4-stone Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

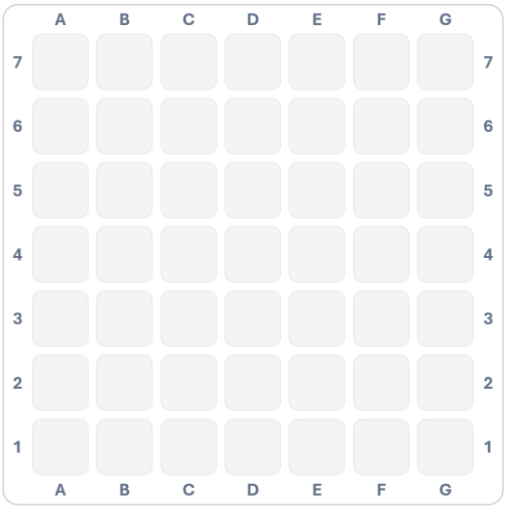

Empty Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

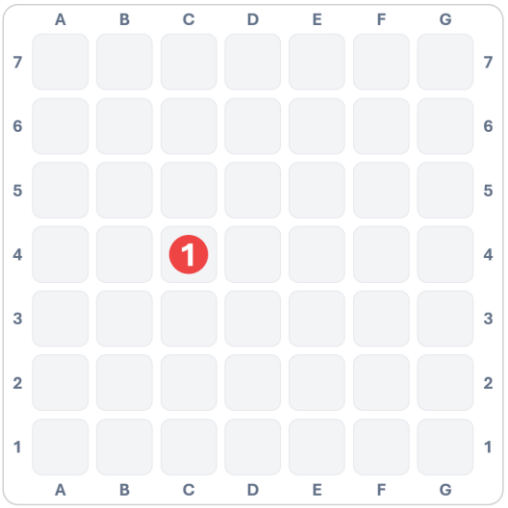

Empty Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

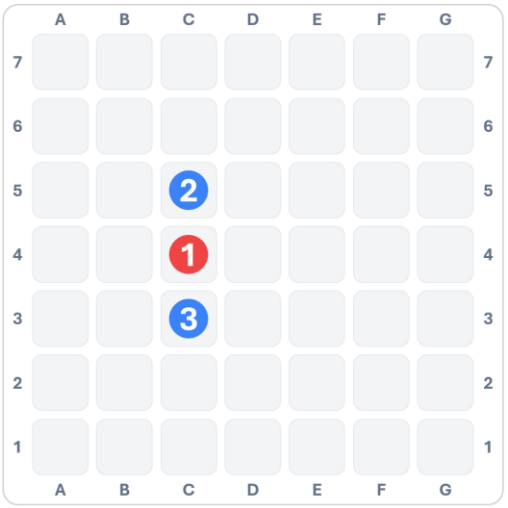

Empty Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

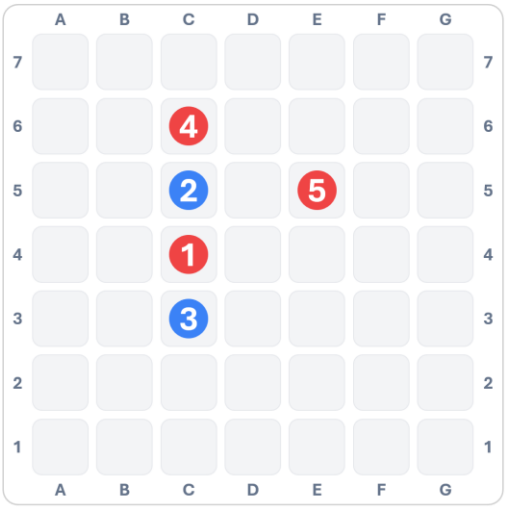

Empty Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

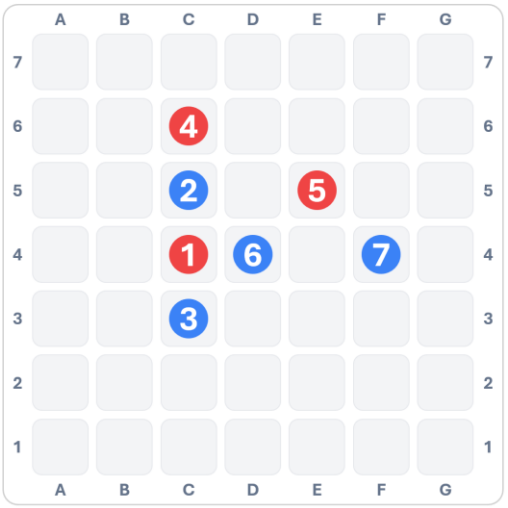

Empty Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones on the board
    - 4-stone mode (used in Netflix's The Devil's Plan)
      - Four stones are pre-positioned
      - Place order: Red→ Blue→ Blue→ Red
    - Empty mode (used on some online platforms)
      - The board starts empty
      - Place order: Red→ Blue→ Blue→ Red→ Red→ Blue→ Blue→ Red

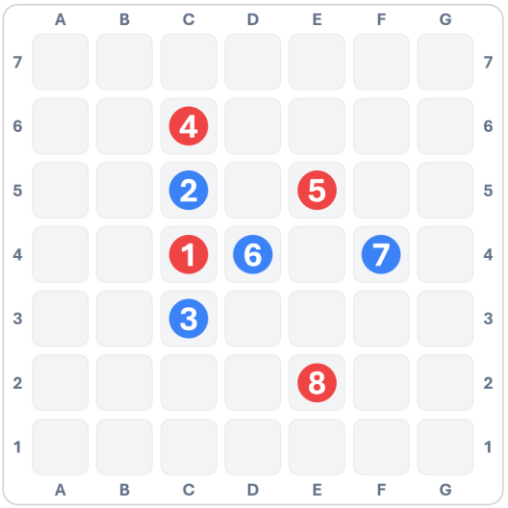

Empty Mode

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
    - Move one stone
      - By zero, one, or two orthogonal steps
      - Stones cannot move through walls or other stones
    - Build one wall
      - Adjacent to the stone's final position

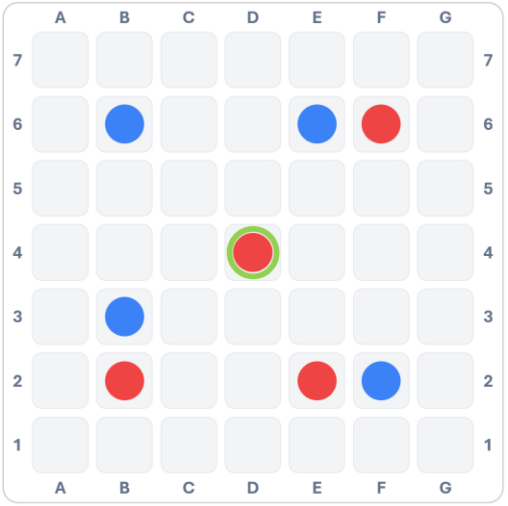

Select the stone

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
    - Move one stone
      - By zero, one, or two orthogonal steps
      - Stones cannot move through walls or other stones
    - Build one wall
      - Adjacent to the stone's final position

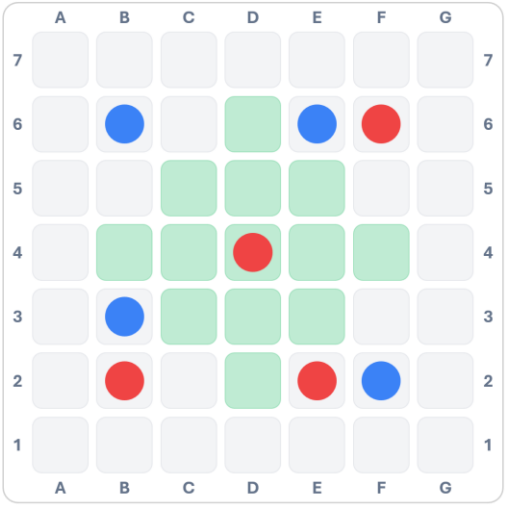

0 ~ 2 steps

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
    - Move one stone
      - By zero, one, or two orthogonal steps
      - Stones cannot move through walls or other stones
    - Build one wall
      - Adjacent to the stone's final position

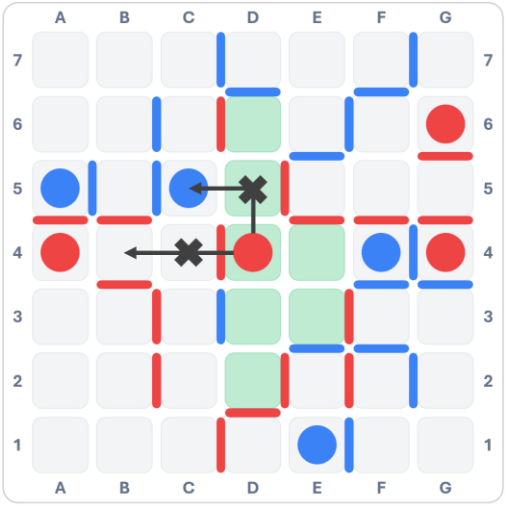

Cannot move through walls and stones

Rules of WallGo in Two-player

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
    - Move one stone
      - By zero, one, or two orthogonal steps
      - Stones cannot move through walls or other stones
    - Build one wall
      - Adjacent to the stone's final position

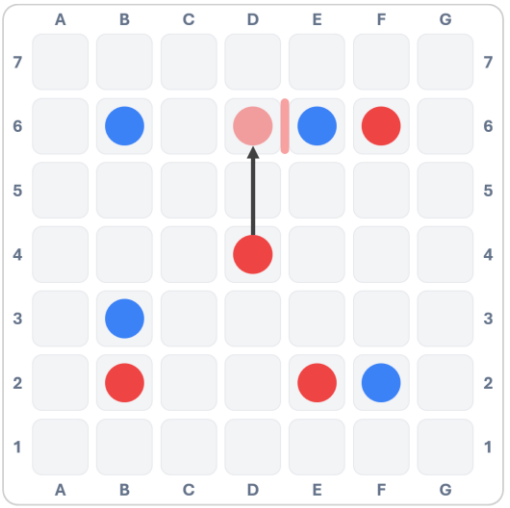

Build one wall

Rules of WallGo in Two-player

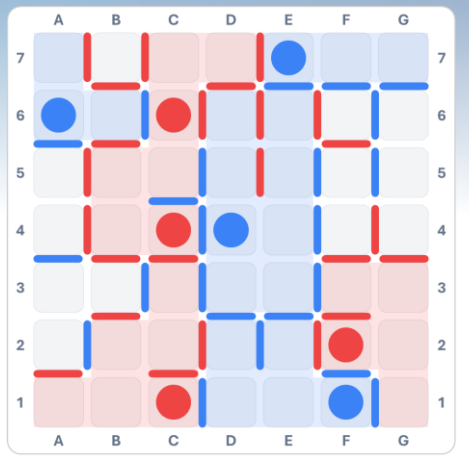

Red 18 vs. Blue 19

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
- Goal: enclose more territory
  - Game ends when all stones are enclosed into single-player regions
    - No region contains both Red and Blue stones
    - Territory is counted from these enclosed regions
  - Winning condition

Rules of WallGo in Two-player

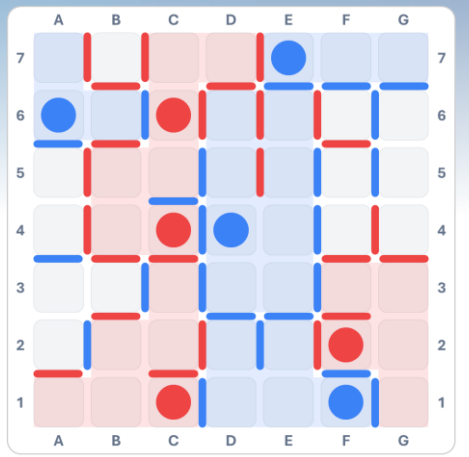

Red 18 vs. Blue 19

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
- Goal: enclose more territory
  - Game ends when all stones are enclosed into single-player regions
  - Winning condition
    - The player with the larger total territory wins
    - If total territory is tied, the player with the larger single region wins
    - If both are tied, the game is a draw

Rules of WallGo in Two-player

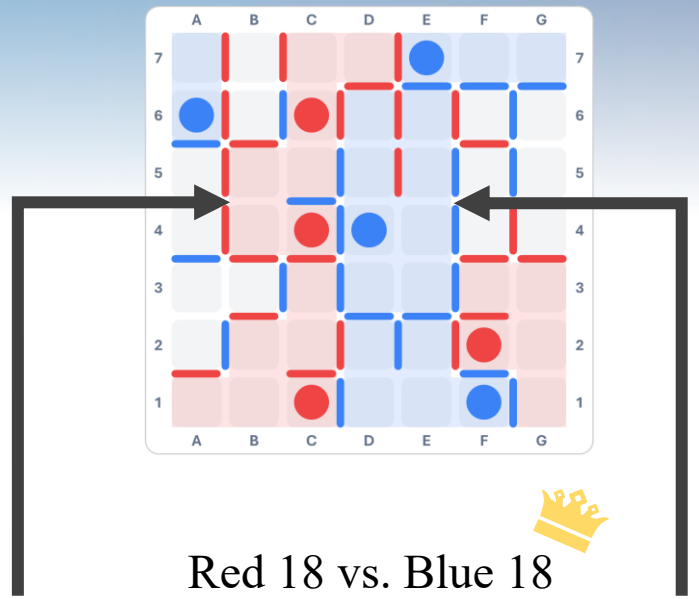

Red 18 vs. Blue 18

Red max region 7 vs. Blue max region 8

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
- Goal: enclose more territory
  - Game ends when all stones are enclosed into single-player regions
  - Winning condition
    - The player with the larger total territory wins
    - If total territory is tied, the player with the larger single region wins
    - If both are tied, the game is a draw

Rules of WallGo in Two-player

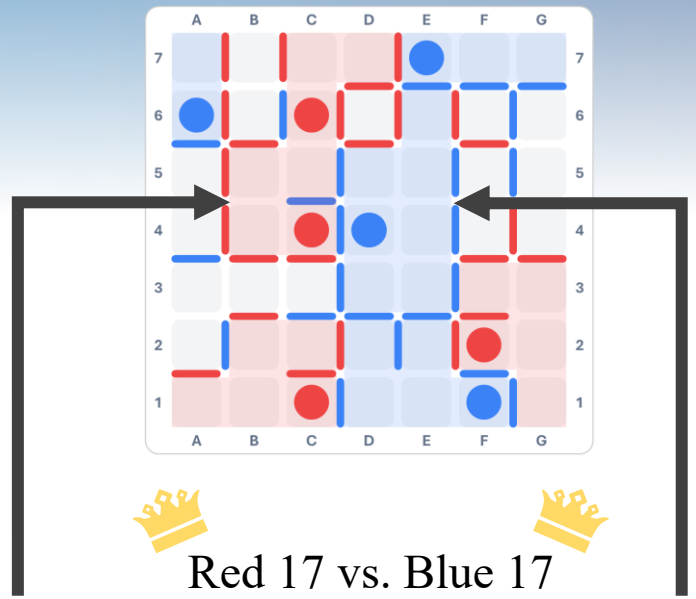

- Played on a 7×7 board
- Each player (Red and Blue) controls four stones
- Two phases:
  - Setup phase: place stones
  - Play phase: move one stone and build one wall
- Goal: enclose more territory
  - Game ends when all stones are enclosed into single-player regions
  - Winning condition
    - The player with the larger total territory wins
    - If total territory is tied, the player with the larger single region wins
    - If both are tied, the game is a draw

Motivation

- Issues:
  - Most existing online AI programs have not reached a high level of play
  - The game's strategies and balance have not been systematically analyzed

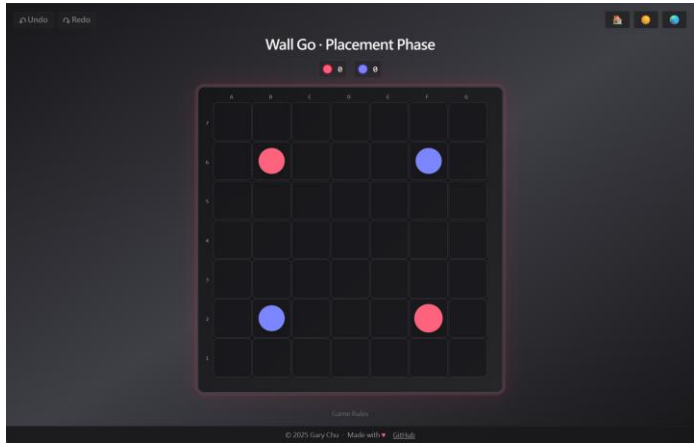

https://schaoss.github.io/wall-go/

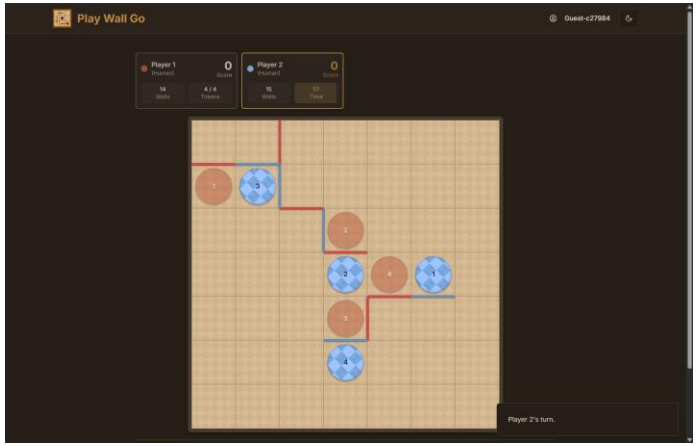

https://www.playwallgo.com/

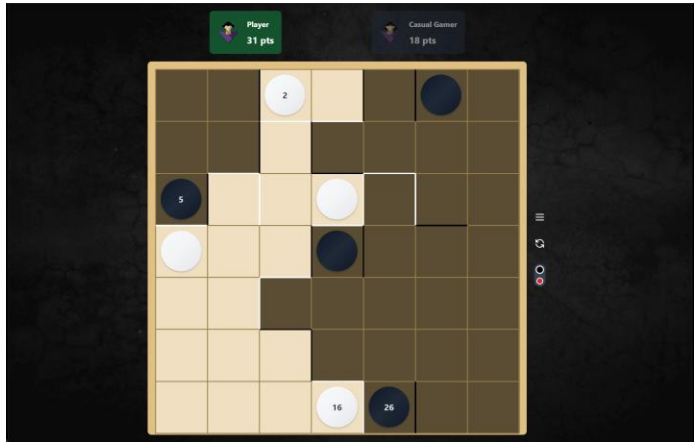

https://www.wallbad.uk/

Why Feature Design Matter?

- State representation is key to AlphaZero-based agents
  - Learning efficiency[13]
  - Final playing performance[4]
- Game-specific features have improved prior systems
  - Faster training in Go[13]
  - Stronger play in Chess[4]
- Motivation: design WallGo-specific features for WallZero

[4] Czech, J., Blüml, J., Kersting, K., Steingrimsson, H.: Representation Matters for Mastering Chess: Improved Feature Representation in AlphaZero Outper forms Switching to Transformers. In: ECAI 2024, pp. 2378–2385. IOS Press, https://ebooks.iospress.nl/doi/10.3233/FAIA240763
[13] Wu, D.J.: Accelerating Self-Play Learning in Go. In: Proceedings of the AAAI Workshop on Reinforcement Learning in Games. https://arxiv.org/pdf/1902.10565.pdf

WallZero Overview

- Develop *WallZero*, an AlphaZero-based agent
- Learn WallGo through self-play
- Use:
  - MCTS for search
  - Neural network for policy and value prediction
- Two WallGo-specific designs:
  - Action design
  - Feature design

WallZero Action Designs

- Setup phase: select a board position (49)
- Play phase: select stone (49), destination (13), and wall direction (4)
- Invalid actions are filtered by an action mask

$$|\mathcal{A}| = \underbrace{49}_{\text{setup phase}} + \underbrace{49 \times 13 \times 4}_{\text{play phase}} = 2597.$$

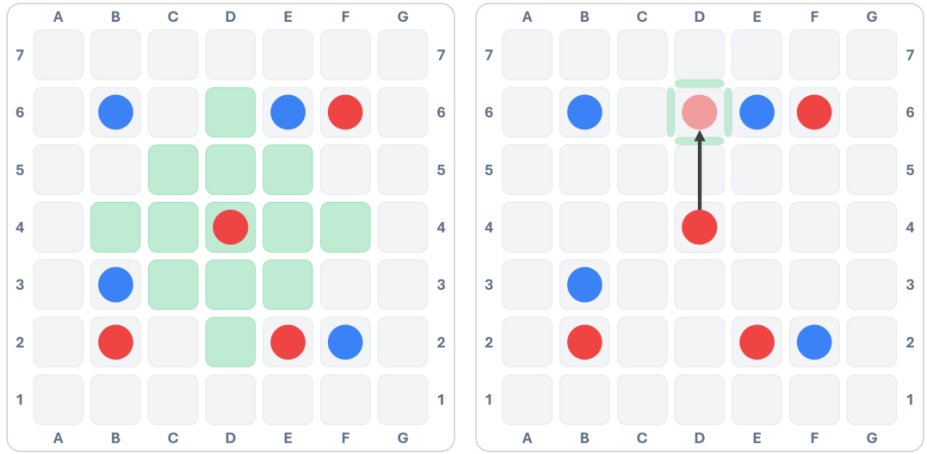

WallZero Feature Designs

- Feature design

Feature	# of planes	Description
Stone	2	Red / Blue stone
Horizontal Wall	2	Red / Blue horizontal wall
Vertical Wall	2	Red / Blue vertical wall
Player Turn	2	Indicates the player to move
Territory (T)	3	Red / Blue / Neutral territory
Reachability (R)	2	Positions reachable within one turn for each player
History (H)	36	Four-step history of stones, walls, and territory (9 per step)

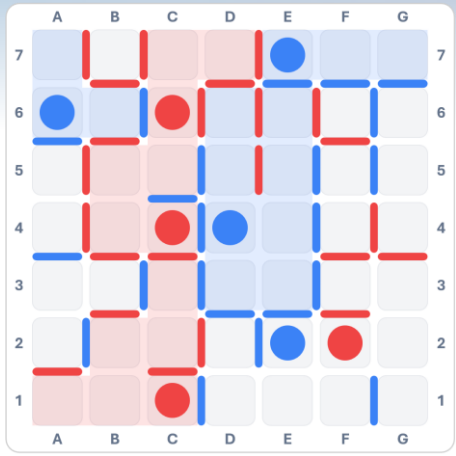

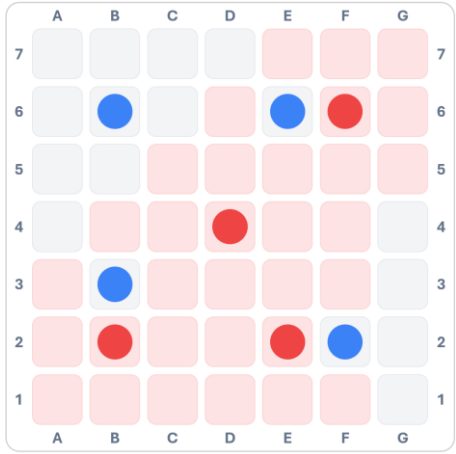

Experimental Setup

- Trained *WallZero* in two modes:
  - Empty mode
  - 4-stone mode
- Training framework: MiniZero[14]
  - 1-block model
  - 200 MCTS simulations
  - 500 training iterations
  - 2000 self-play games per iteration
  - 500 optimization steps per iteration

[14] Wu, T.R., Guei, H., Peng, P.C., Huang, P.W., Wei, T.H., Shih, C.C., Tsai, Y.J.: MiniZero: Comparative Analysis of AlphaZero and MuZero on Go, Othello, and Atari Games 17(1), 125–137, https://ieeexplore.ieee.org/document/10510513

Feature Design Evaluation

- Compare four feature settings:
  - B: basic board state
  - BT: + territory
  - BTR: + reachability
  - BTRH(WallZero): + history
- Reachability gives the largest improvement
- Full feature design achieves the best performance

Feature	# of planes	Description
Stone	2	Red / Blue stone
Horizontal Wall	2	Red / Blue horizontal wall
Vertical Wall	2	Red / Blue vertical wall
Player Turn	2	Indicates the player to move
Territory (T)	3	Red / Blue / Neutral territory
Reachability (R)	2	Positions reachable within one turn for each player
History (H)	36	Four-step history of stones, walls, and territory (9 per step)

+52.38 ± 2.15

	WallZero-B	WallZero-BT	WallZero-BTR	WallZero
Empty Mode (Avg.)	10.65±1.10%	27.05±1.59%	79.43±1.45%	82.87±1.35%
4-Stone Mode (Avg.)	10.45±1.09%	28.65±1.62%	78.88±1.46%	82.02±1.37%

+50.23 ± 2.18

Human-AI Evaluation

- Extend 1-block ➔ 10-block (longer training)
- *WallZero* played against two Taiwanese professional Go players
  - Chun-Hsun Chou(9-dan) and Wei Huang(3-dan)
  - They played many games before matches
  - Matches conducted in both modes and both colors
- Won all 8 games
- WallZero gains 1.98× as much territory as human players on average

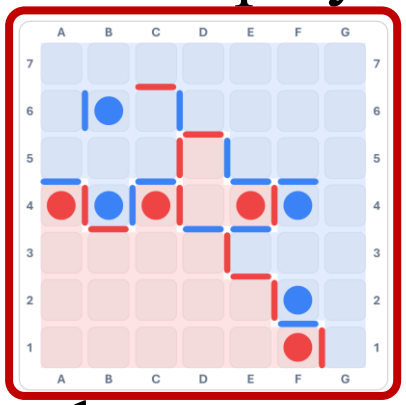

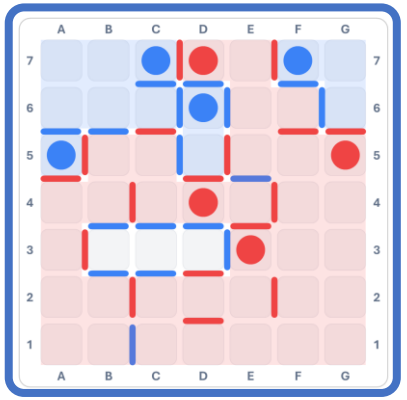

	Empty Mode		4-Stone Mode	
	Red	Blue	Red	Blue
3-dan	16–33 (2.06×)	20–29 (1.45×)	14–32 (2.29×)	19–30 (1.58×)
9-dan	12–37 (3.08×)	17–30 (1.76×)	20–29 (1.45×)	12–34 (2.83×)

Opening Strategy & Game Balance Analysis

- Empty mode:
  - openings are diverse, and stones tend to concentrate near the center
  - Red win rate = 53.57% ± 1.38% (5000 WallZero self-play games)
- 4-stone mode:
  - only a few openings, and players focus on controlling opponent reachability
  - Red win rate = 50.37% ± 1.38% (5000 WallZero self-play games)

Empty Mode

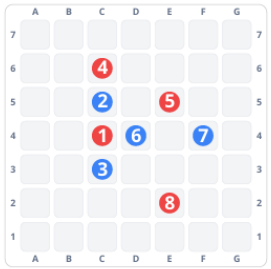

freq. = 5.8%

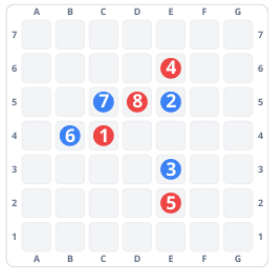

freq. = 5.0%

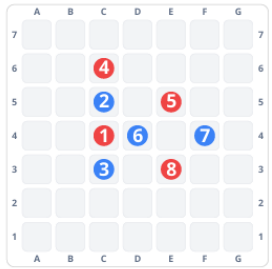

freq. = 4.5%

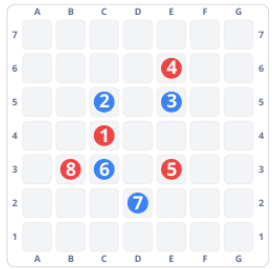

freq. = 4.1%

4-Stone Mode

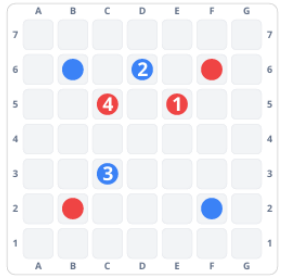

freq. = 49.0%

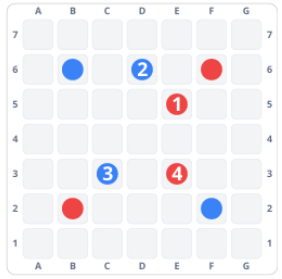

freq. = 21.0%

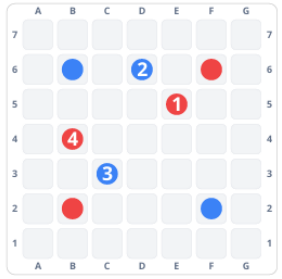

freq. = 15.1%

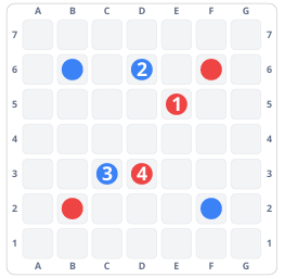

freq. = 1.8%

Strategy Analysis

- Analyze *WallZero's* 4-stone self-play games
- Incorporate expert insights from professional Go players
- Identify strategic patterns in WallGo
  - Reachability control
  - Passing strategy

Strategy 1: Reachability Control

- Reachability refers to the positions a player can access in future turns
- Example:
  - Red moves from E5 to E3, different wall placements lead to distinct future reachability
  - Percentages denote the win rate from Red's perspective

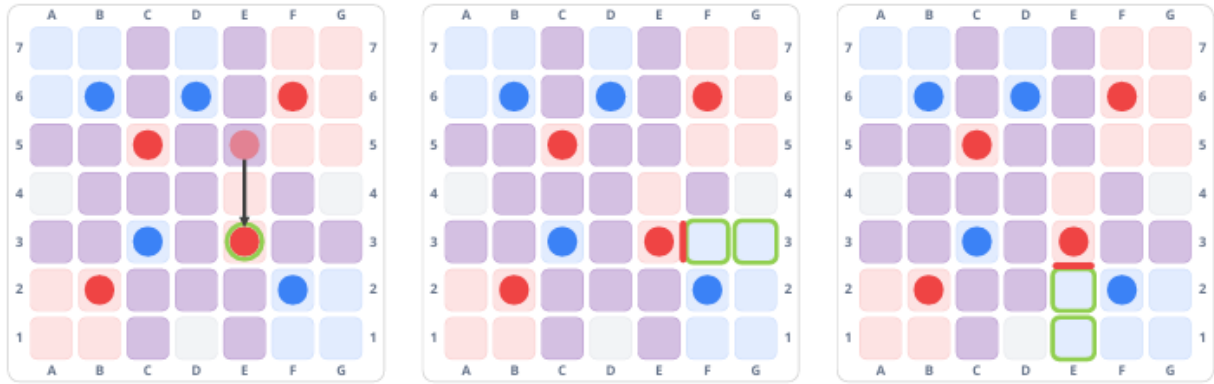

(a) (b) 46% (c) 47%

Strategy 1: Reachability Control

- Reachability refers to the positions a player can access in future turns
- Example:
  - Red moves from E5 to E3, different wall placements lead to distinct future reachability
  - Percentages denote the win rate from Red's perspective

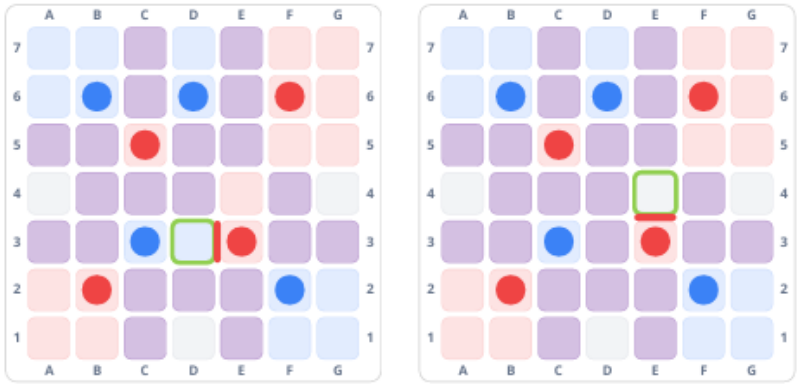

(d) 49.5%

(e) 52.5%

Strategy 1: Reachability Control

- *WallZero* often prioritizes future mobility over immediate territory
- Example:
  - Blue sacrifices one point of immediate territory to maintain access to the central region
  - Percentages denote the win rate from Blue's perspective

(a) (b) 17.5% (c) 55%

Strategy 1: Reachability Control

- *WallZero* often prioritizes future mobility over immediate territory
- Example:
  - Blue sacrifices one point of immediate territory to maintain access to the central region
  - Percentages denote the win rate from Blue's perspective

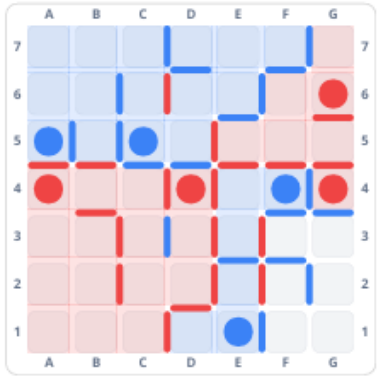

(d) Red Win (R:22, B:21)

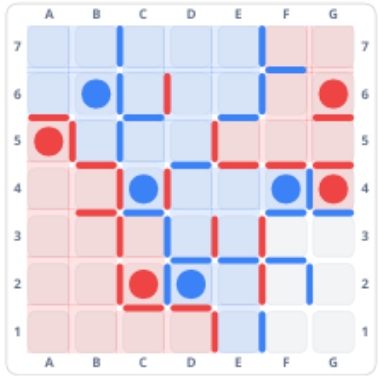

(e) Blue Win (R:21, B:22)

Strategy 2: Implicit Passing

- Problem:
  - Wall construction often restrict own's reachability
  - WallGo does not allow players to explicitly pass
- *WallZero's* solution simulates passing by:
  - Building a low-impact wall
  - Forcing the opponent to move first in a critical region

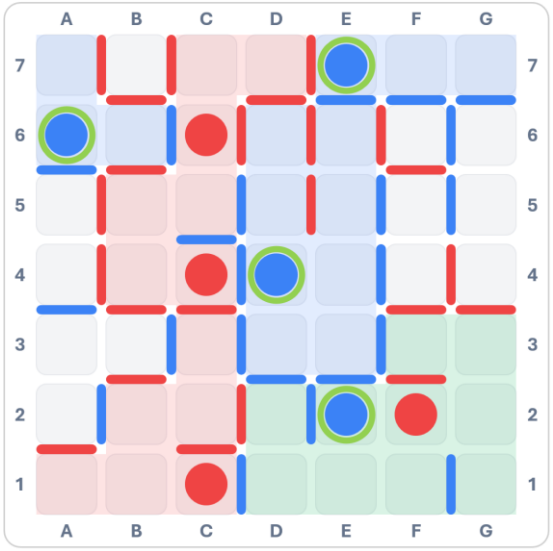

Red 13 vs. Blue 14

Strategy 2: Implicit Passing

- Problem:
  - Wall construction often restrict own's reachability
  - WallGo does not allow players to explicitly pass
- *WallZero's* solution simulates passing by:
  - Building a low-impact wall
  - Forcing the opponent to move first in a critical region

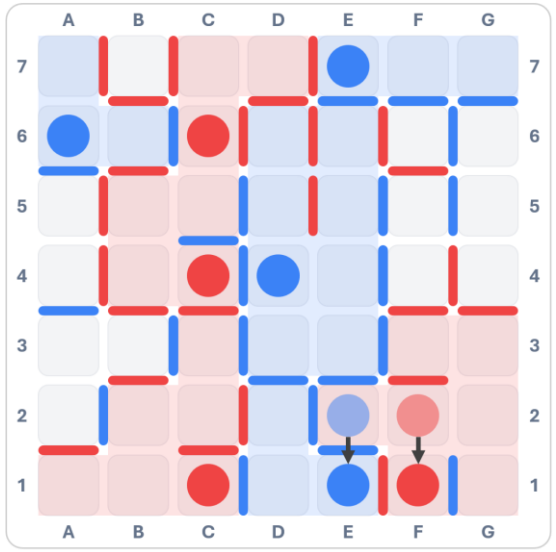

If Blue moves first,
Red will win

Strategy 2: Implicit Passing

- Problem:
  - Wall construction often restrict own's reachability
  - WallGo does not allow players to explicitly pass
- *WallZero's* solution simulates passing by:
  - Building a low-impact wall
  - Forcing the opponent to move first in a critical region

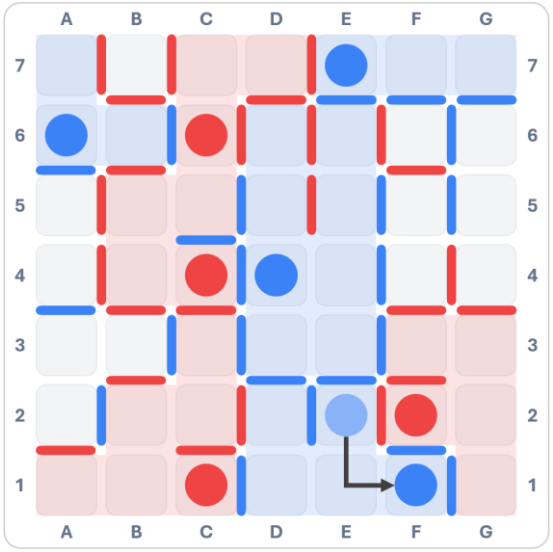

If Red moves first,
Blue will win

Strategy 2: Implicit Passing

- Problem:
  - Wall construction often restrict own's reachability
  - WallGo does not allow players to explicitly pass
- *WallZero's* solution simulates passing by:
  - Building a low-impact wall
  - Forcing the opponent to move first in a critical region

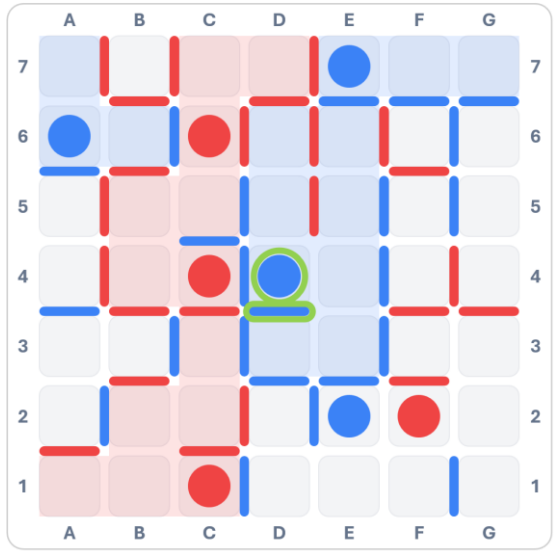

Blue loses no territory

Strategy 2: Implicit Passing

- Problem:
  - Wall construction often restrict own's reachability
  - WallGo does not allow players to explicitly pass
- *WallZero's* solution simulates passing by:
  - Building a low-impact wall
  - Forcing the opponent to move first in a critical region

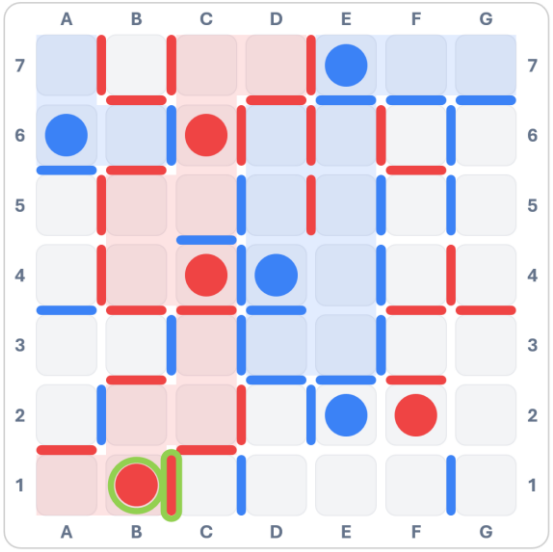

Red loses only one territory

Strategy 2: Implicit Passing

- Problem:
  - Wall construction often restrict own's reachability
  - WallGo does not allow players to explicitly pass
- *WallZero's* solution simulates passing by:
  - Building a low-impact wall
  - Forcing the opponent to move first in a critical region

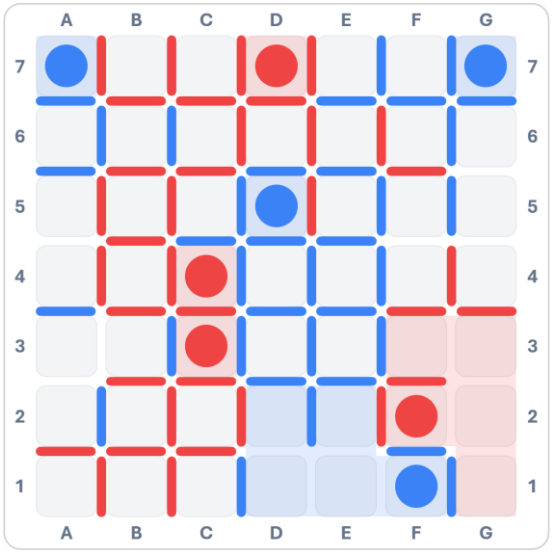

Summary

- *WallZero* achieves strong performance in WallGo through AlphaZero-style self-play.
- Both feature design evaluation and strategic analysis show that reachability is the key concept for playing WallGo.
- The 4-stone mode provides better game balance than the empty mode.

Thank You for Your Attention

Our code and data are available at
https://rlg.iis.sinica.edu.tw/papers/wallzero

Play with *WallZero*!