CS 4260 and CS 5260
Vanderbilt University

More thoughts on Game Search

Consider game tree search

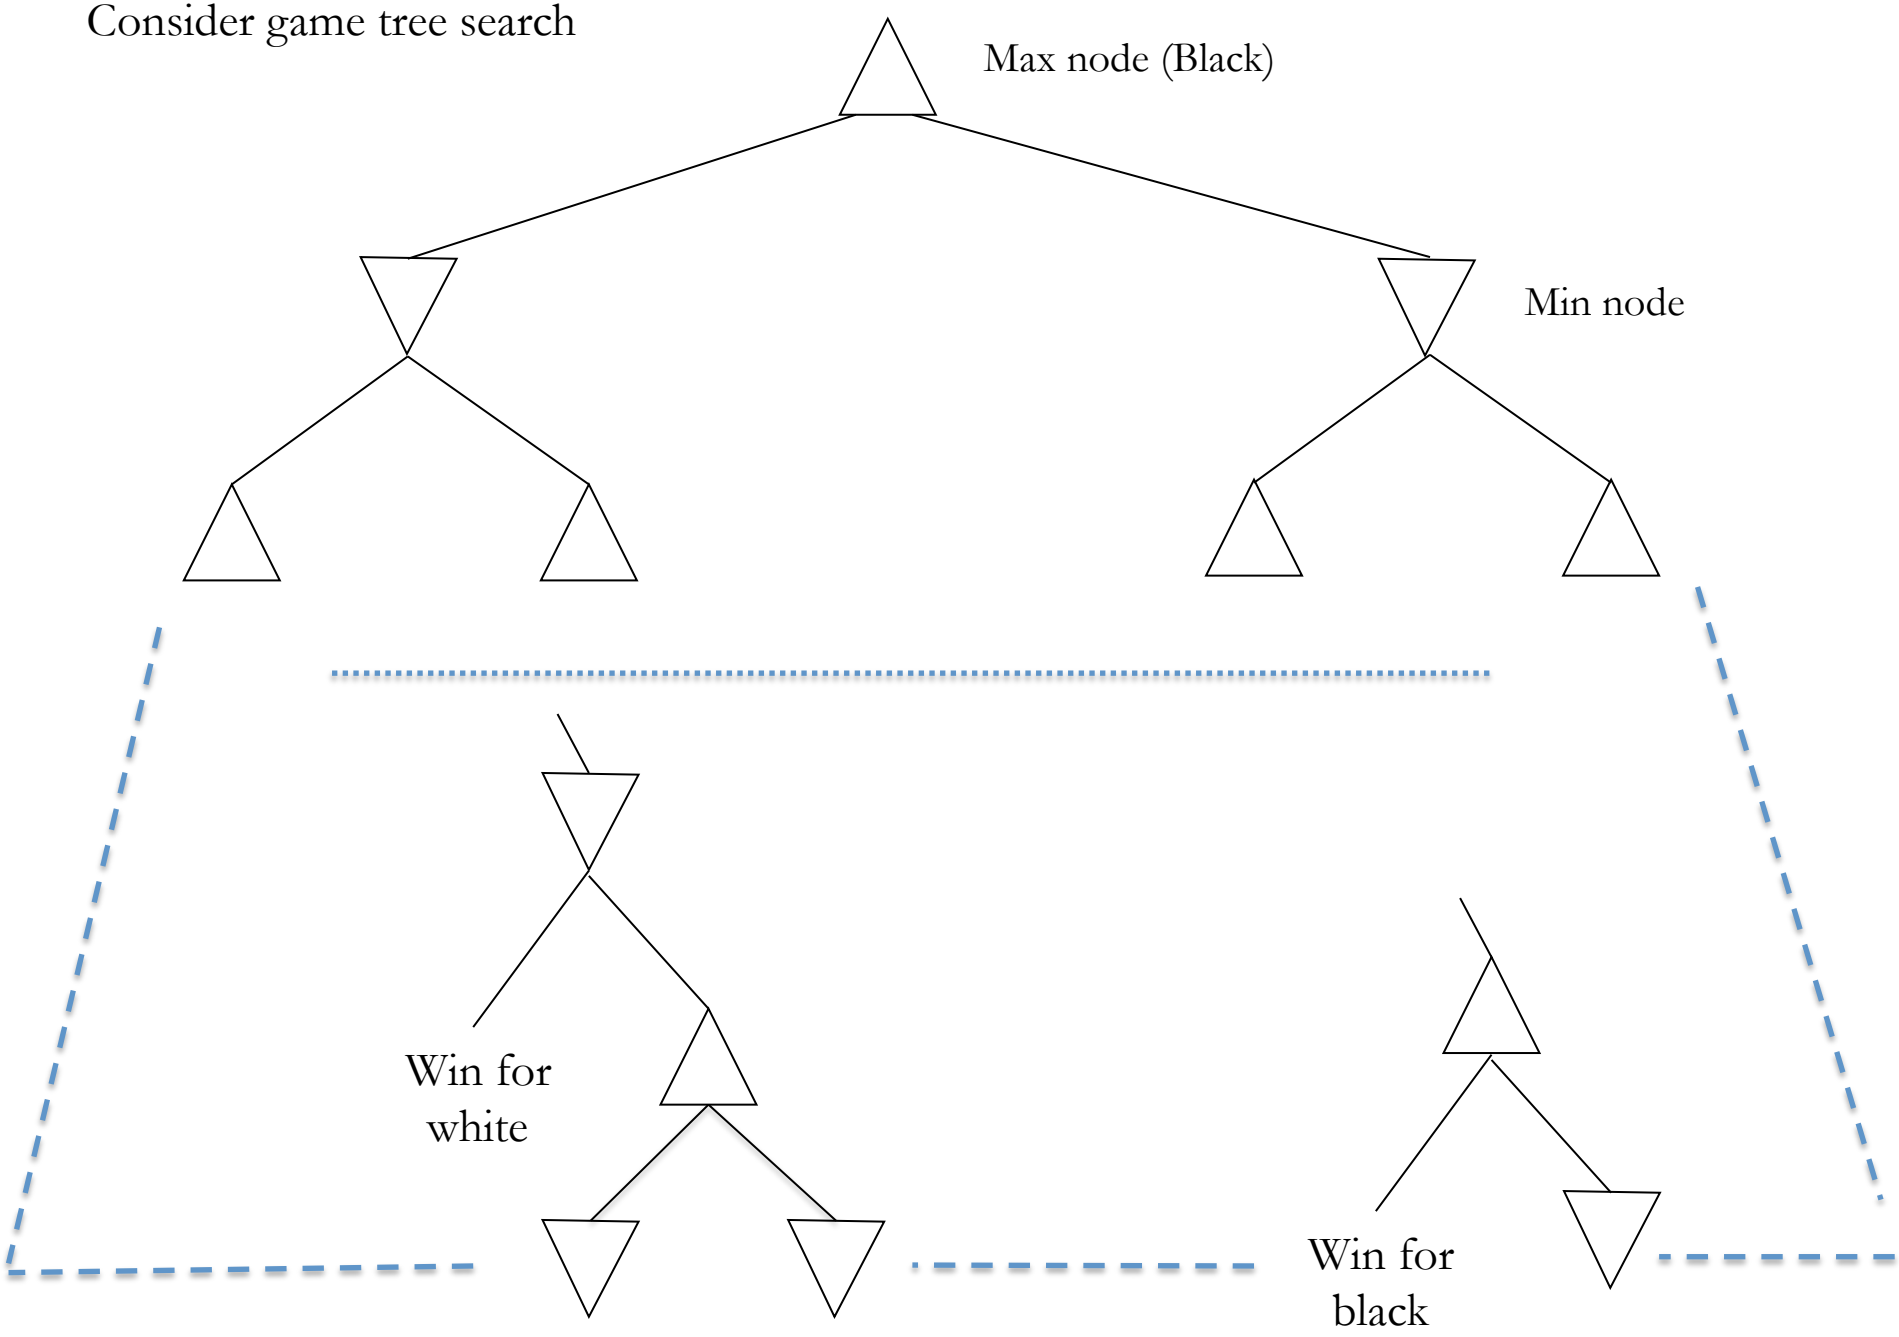

Game trees are huge for non-trivial games

≈ 10^{120} games (paths) for Chess
≈ 10^{761} for Go

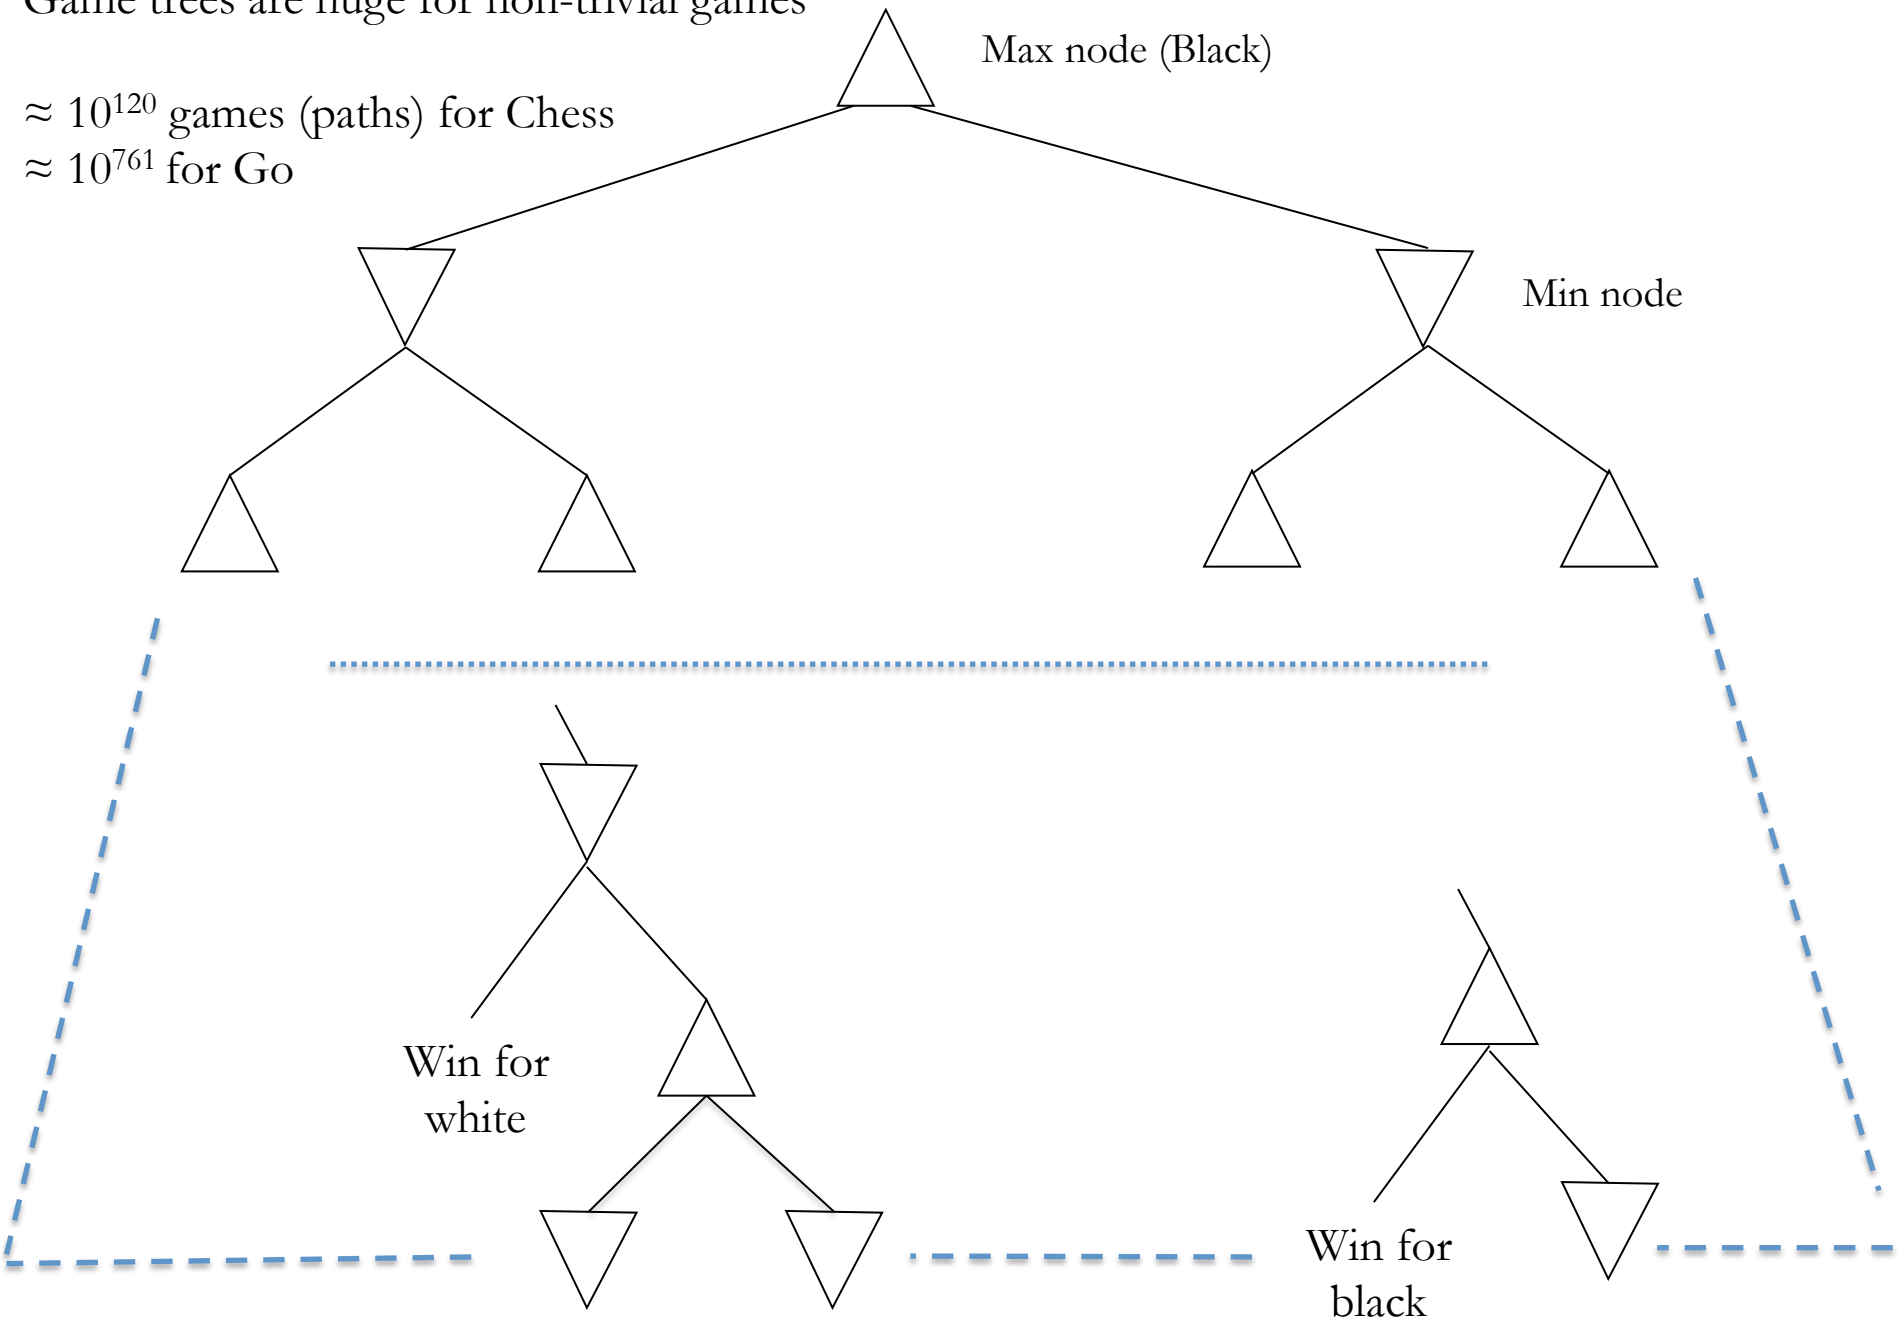

In Go, or any game

Play others or play yourself
and learn from results

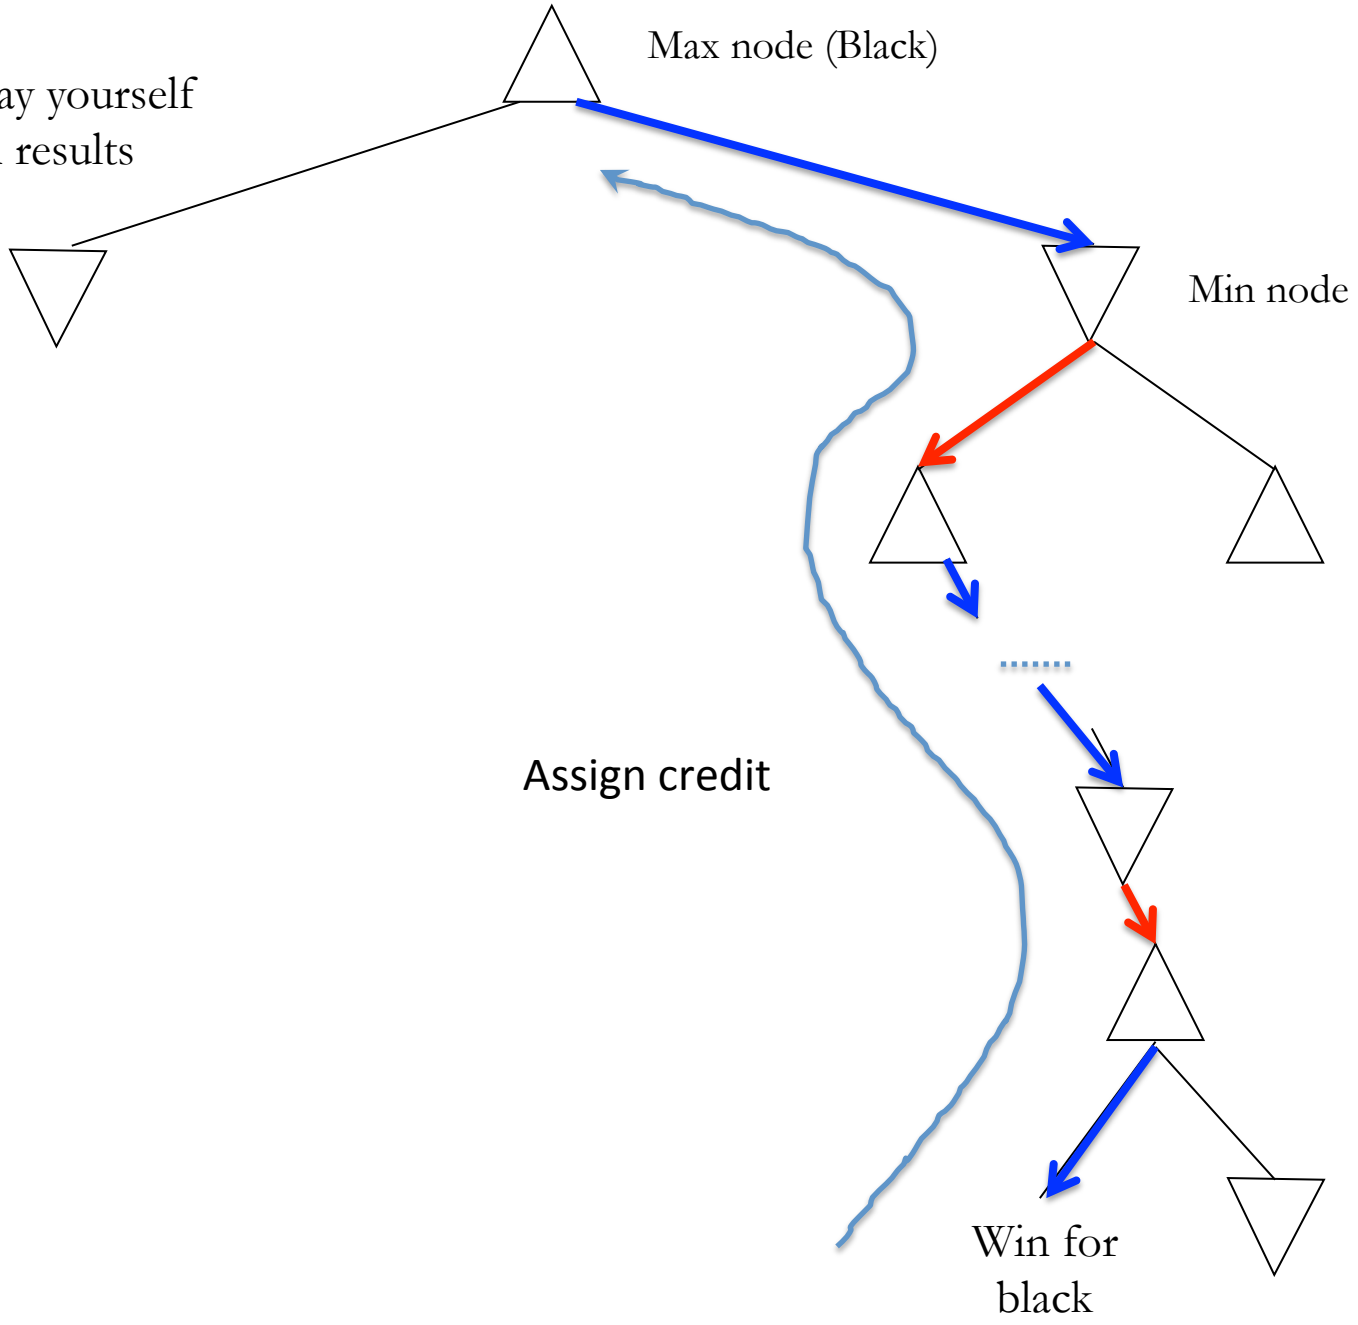

In Go, or any game

Play others or play yourself
and learn from results

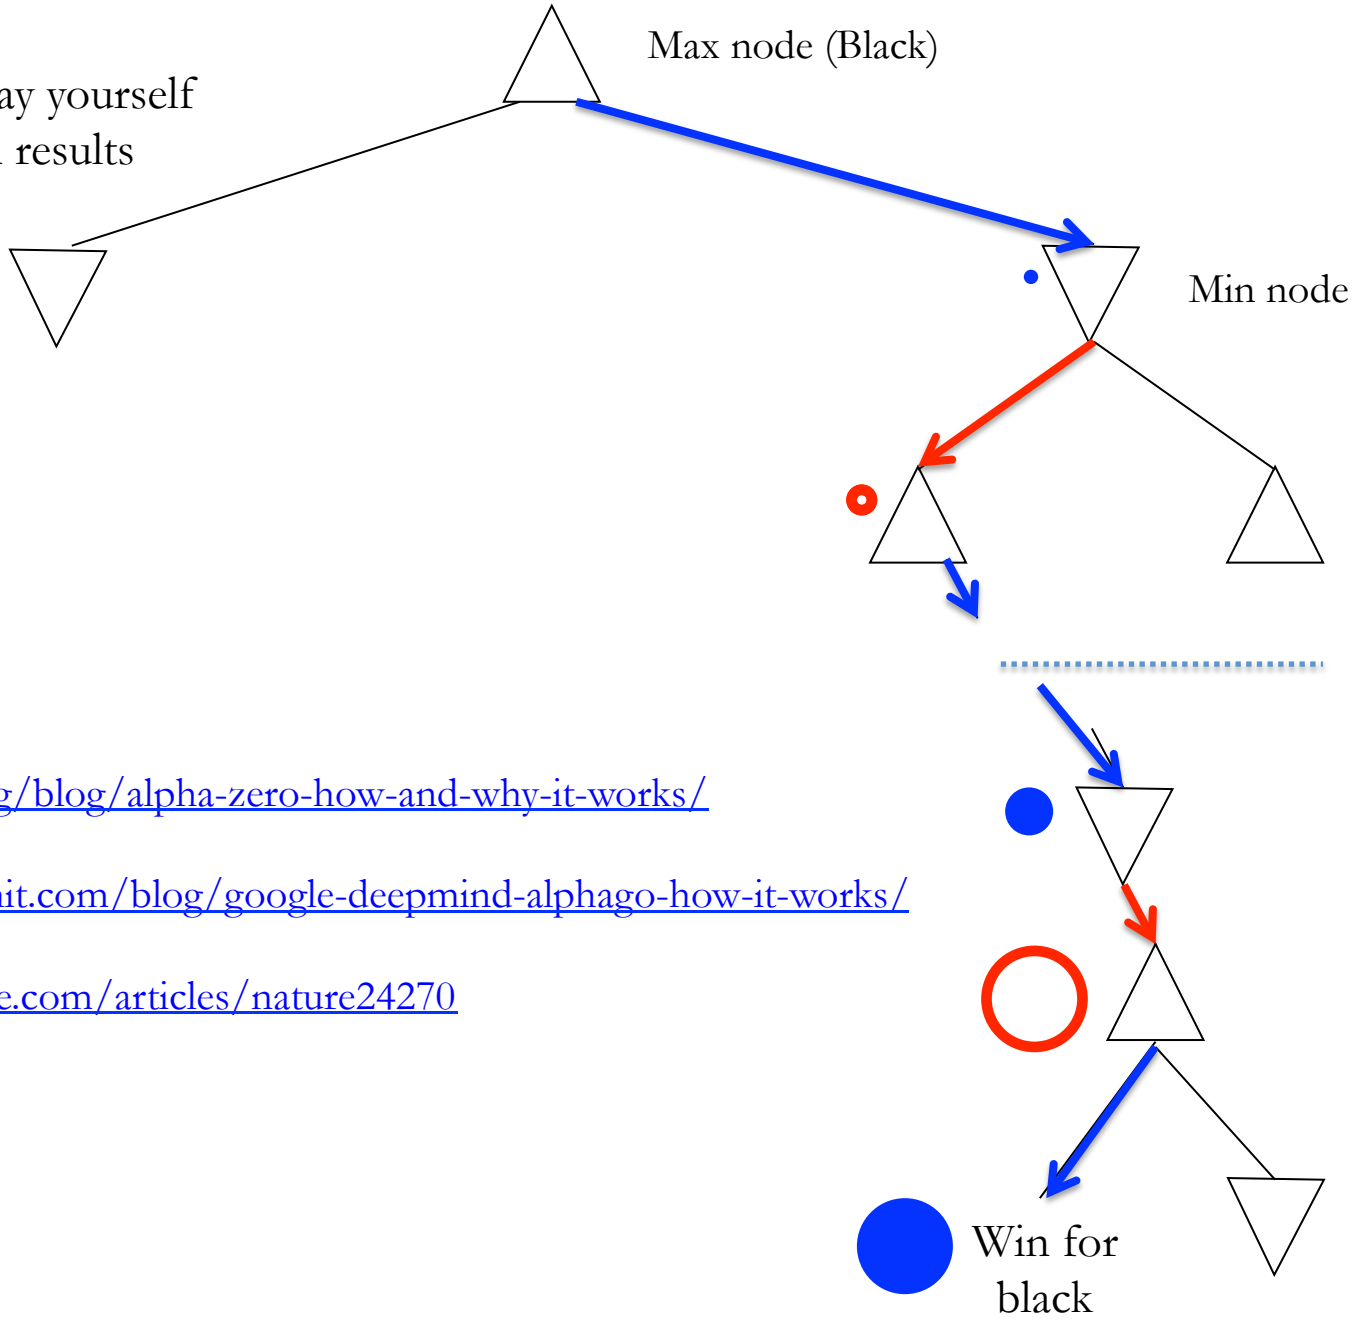

http://tim.hibal.org/blog/alpha-zero-how-and-why-it-works/

https://www.tastehit.com/blog/google-deepmind-alphago-how-it-works/

https://www.nature.com/articles/nature24270

In Go, or any game

Play others or play yourself and learn from results

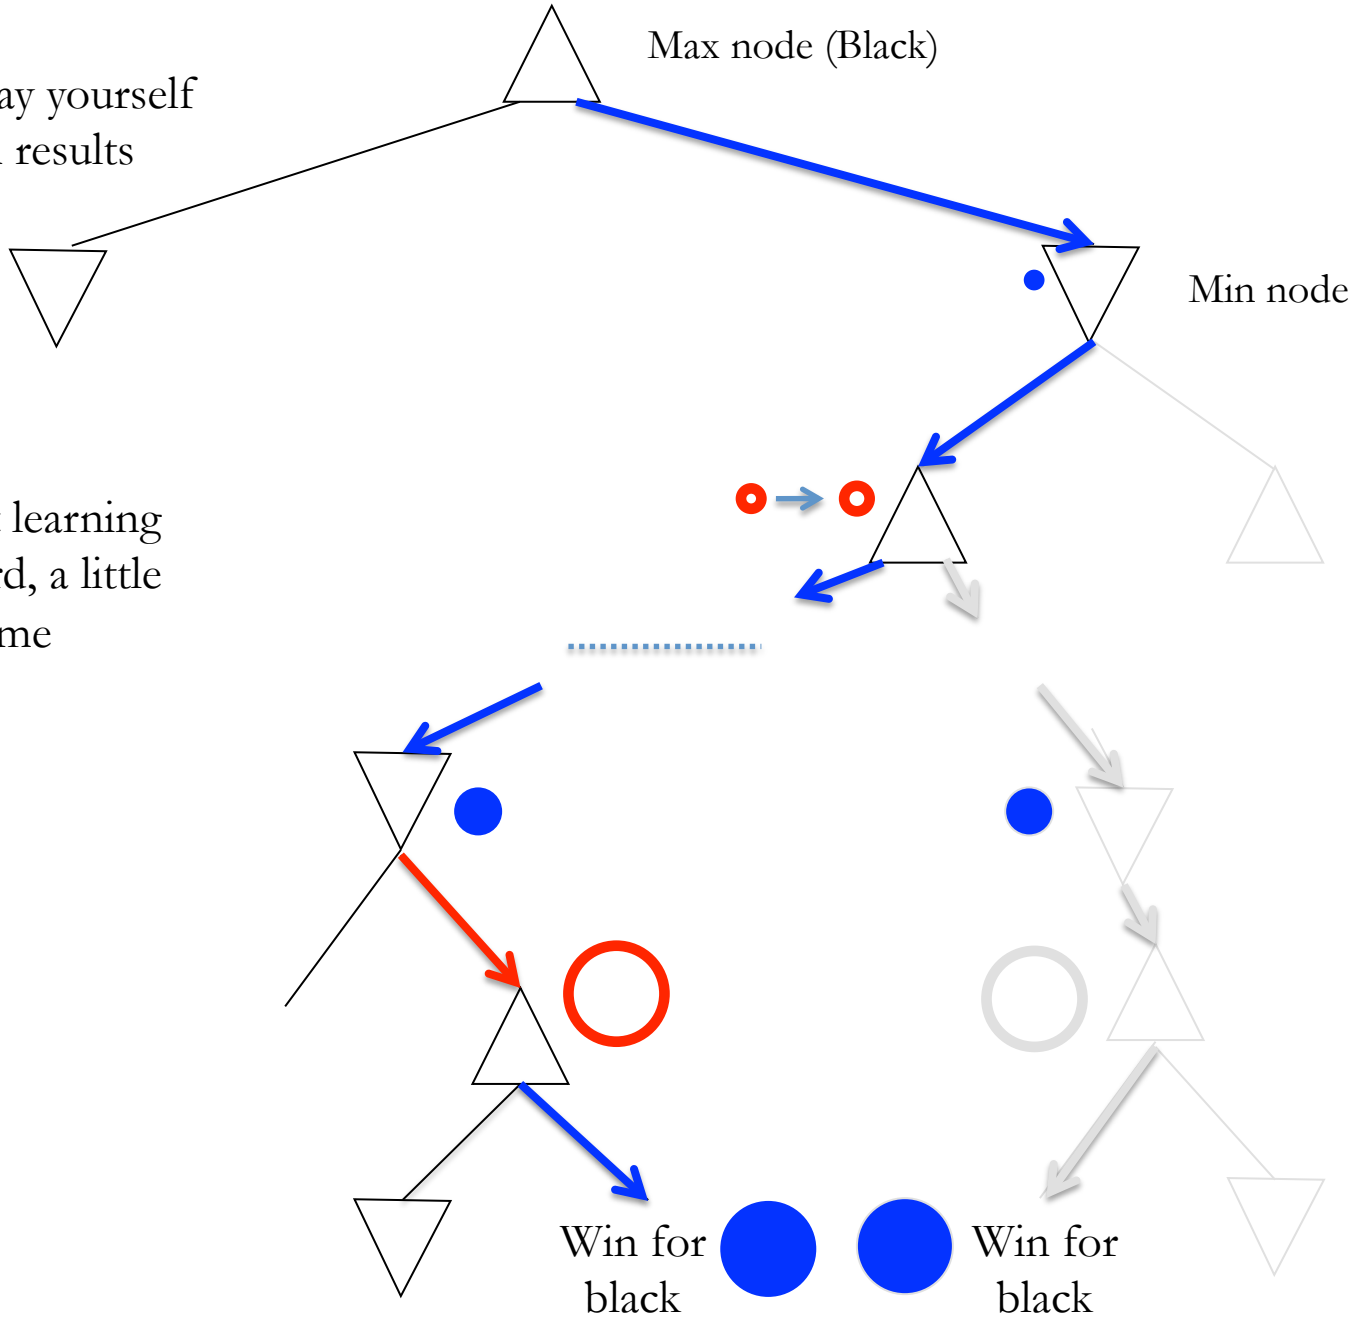

Reinforcement learning backs up reward, a little bit at a time

Arthur Samuel

From Wikipedia, the free encyclopedia

For the British politician, see Arthur Samuel, 1st Baron Mancroft.

Arthur Lee Samuel (December 5, 1901 – July 29, 1990)[3] was an American pioneer in the field of computer gaming and artificial intelligence.[1] He coined the term "machine learning" in 1959.[4] The Samuel Checkers-playing Program was among the world's first successful self-learning programs, and as such a very early demonstration of the fundamental concept of artificial intelligence (AI).[5] He was also a senior member in the TeX community who devoted much time giving personal attention to the needs of users and wrote an early TeX manual in 1983.[6]

Contents [hide]

Arthur Lee Samuel

Samuel also designed various mechanisms by which his program could become better. In what he called rote learning, the program remembered every position it had already seen, along with the terminal value of the reward function. This technique effectively extended the search depth at each of these positions. Samuel's later programs reevaluated the reward function based on input from professional games. He also had it play thousands of games against itself as another way of learning. With all of this work, Samuel's program reached a respectable amateur status, and was the first to play any board game at this high a level. He continued to work on checkers until the mid-1970s, at which point his program achieved sufficient skill to challenge a respectable amateur.[13]

11.2 Samuel's Checkers Player

An important precursor to Tesauro's TD-Gammon was the seminal work of Arthur Samuel (1959, 1967) in constructing programs for learning to play checkers. Samuel was one of the first to make effective use of heuristic search methods and of what we would now call temporal-difference learning. His checkers players are instructive case studies in addition to being of historical interest. We emphasize the relationship of Samuel's methods to modern reinforcement learning methods and try to convey some of Samuel's motivation for using them.

https://webdocs.cs.ualberta.ca/~jonathan/publications/ai_publications/samuel.pdf

In Go, or any game

Play others or play yourself
and learn from results

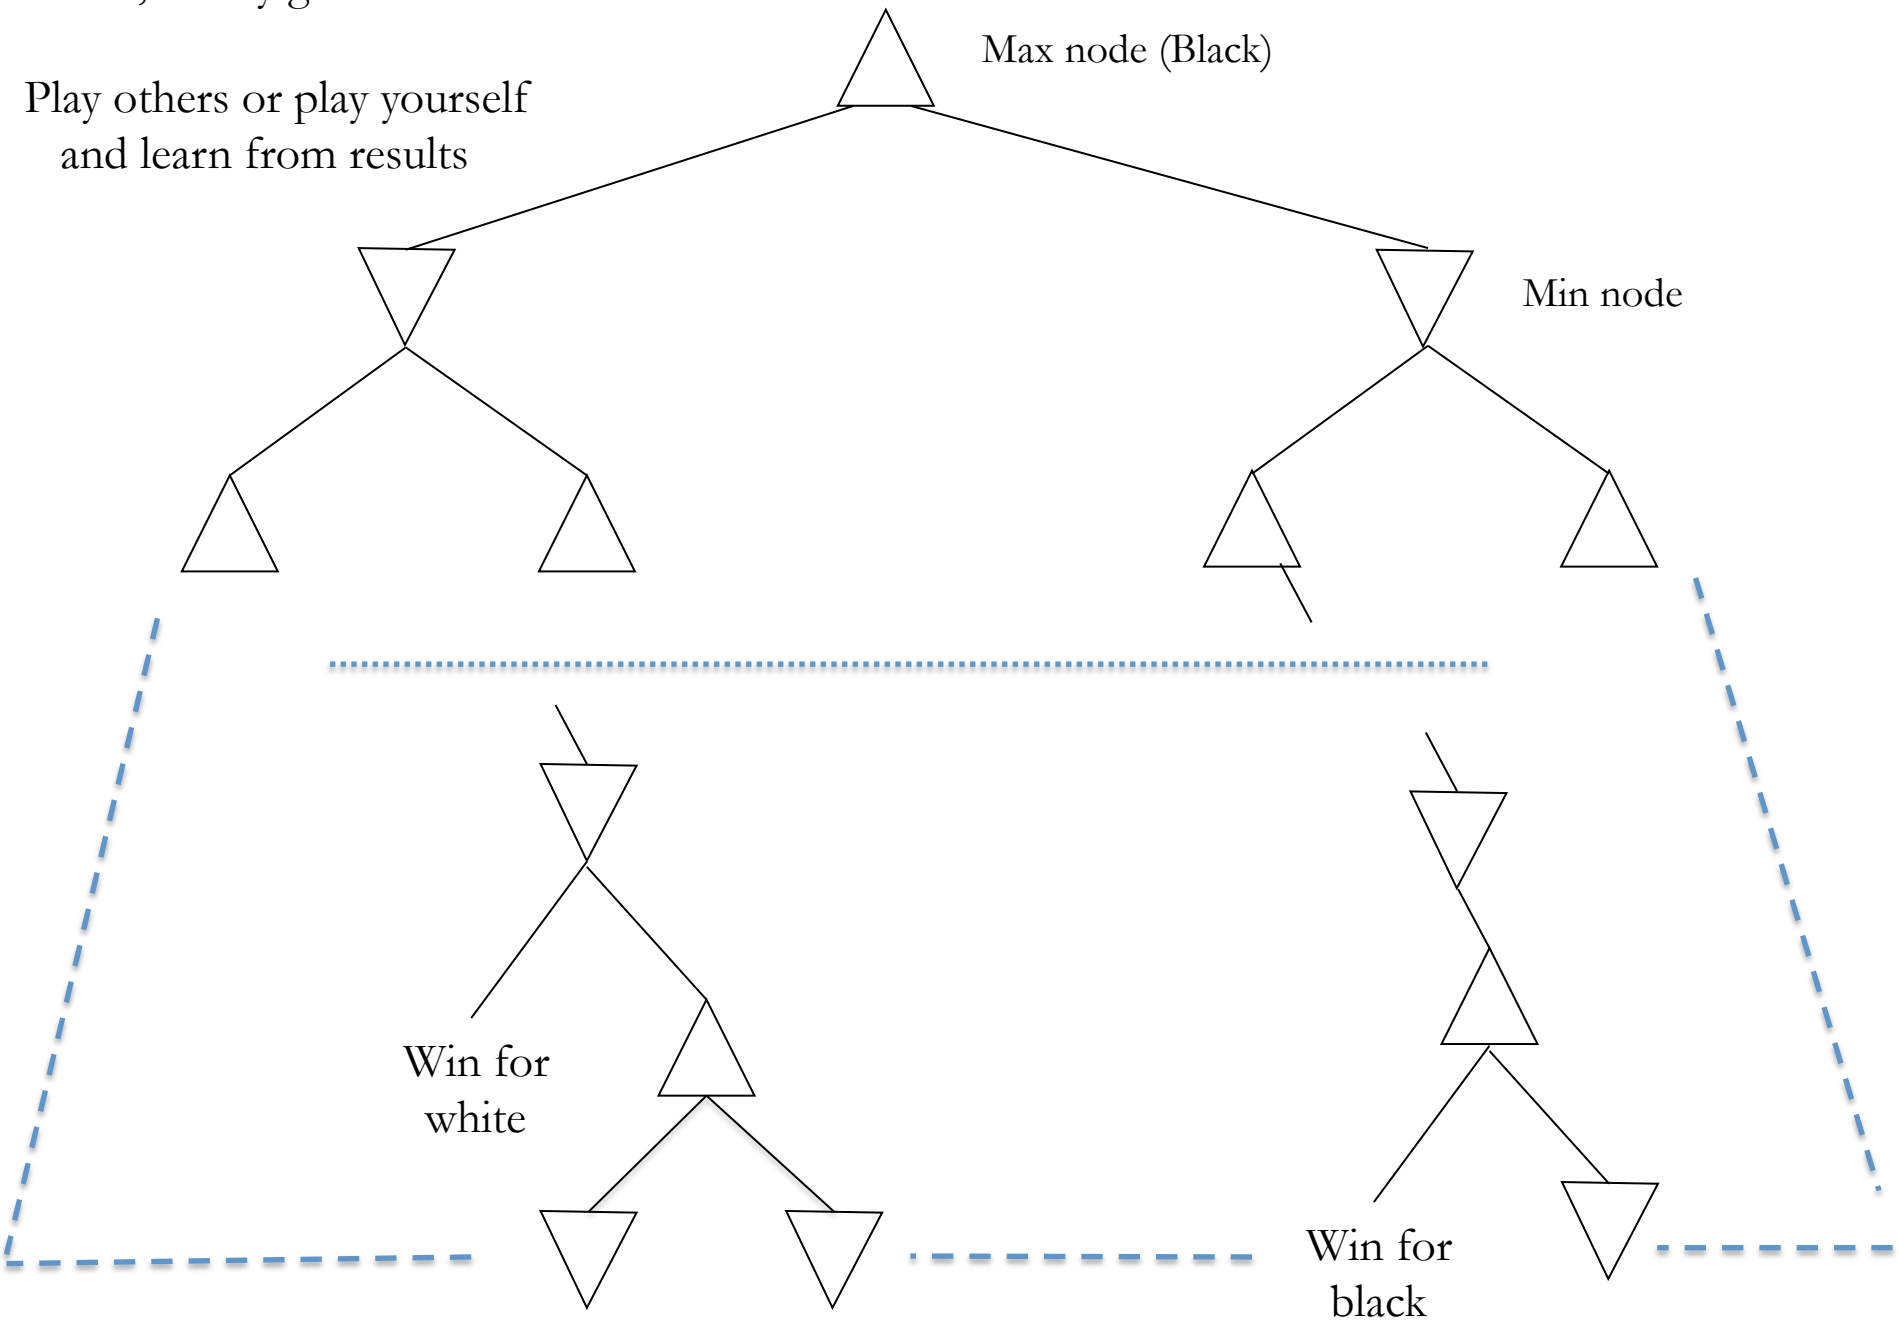

Radiolab - The Rules Can Set You Free [Brian Christian and Alison Gopnik]

https://www.youtube.com/watch?v=nwkxxs-xJHs