

## **HANE (Hybrid Attention with self-supervised kNowledge)**

The HANE engine represents a novel approach to building a powerful Go-playing AI, engineered around two core principles: deep game intuition and context-aware strategic reasoning. Our methodology diverges from standard models by employing a sophisticated two-phase training regimen and a unique hybrid neural network architecture.

1. **Intuition First (Self-Supervised Pre-training):** The model first learns the fundamental "grammar" of the game of Go without direct supervision, building a robust internal representation of board dynamics.

2. **Strategic Reasoning (Hybrid Architecture):** The network's architecture combines the proven pattern-recognition strength of Convolutional Neural Networks (CNNs) with a custom, Go-specific attention mechanism. This allows the AI to reason about the global state of the board in a manner that is aware of key Go concepts like group connectivity and tactical urgency.

To develop a truly strong AI, the model moved beyond conventional supervised learning from the outset. Our training process is divided into two distinct, sequential phases.

### **Phase 1: Self-Supervised Pre-training with Masked Go Modeling (MGM)**

Before ever attempting to predict the next best move, the model undergoes a foundational training phase inspired by cutting-edge techniques in natural language processing (like BERT). We treat the Go board as a sentence of 361 "tokens." During each training step, we randomly "mask" (hide) a portion of the stones on the board. The model's sole objective is to reconstruct the original board state by correctly predicting the identity of the masked positions (Black stone, White stone, or Empty). This process forces the model to learn the intricate relationships between stone placements and the fundamental principles of shape, influence, and connectivity. It develops a powerful intuition about what constitutes a "sensible" or "logical" board state. This provides an exceptionally strong foundation for the next phase.

## Phase 2: Supervised Fine-tuning

Once the model has developed its core intuition, we fine-tune it on a traditional supervised task. Using the pre-trained weights from Phase 1 as a starting point, the model learns to predict the next move (the policy head) and the eventual game outcome (the value head) from a dataset of expert game records. This phase is significantly more effective and data-efficient due to the powerful foundation built during pre-training.

The "brain" of HANE is a custom-designed hybrid neural network that leverages the strengths of two different architectures. The foundation of our network is a deep stack of residual convolutional blocks (ResNet). These layers are exceptionally effective at identifying local patterns, shapes, and tactical motifs within small regions of the board—a critical skill for any Go AI. Integrated at strategic depths within the convolutional backbone are our novel **Go-Attention layers**. A standard attention mechanism allows a model to weigh the importance of all positions on the board simultaneously, overcoming the "local vision" limitation of CNNs. However, our mechanism is uniquely enhanced with Go-specific domain knowledge. When the attention mechanism calculates the relationships between every pair of points on the board, we introduce an "inductive bias" that explicitly tells the model to pay more attention to connections that are important in Go. This bias is dynamically computed for every board state and gives higher importance to points that are:

- **Part of the same stone group (group connectivity).**
- **In a state of tactical urgency (atari).**
- **Spatially adjacent and relevant to a group's life and death (liberty proximity).**

This allows the model to perform a kind of relational reasoning that mirrors human expert intuition. It doesn't just see the board as a grid of pixels; it understands it as a dynamic collection of interconnected, warring groups.

The unique design of HANE provides several key advantages:

- **Deep Game Understanding:** The MGM pre-training ensures the model has a robust understanding of Go fundamentals beyond simple move prediction.
- **Global Strategic Awareness:** The attention mechanism enables the model to understand the long-range influence of moves and evaluate the global board balance effectively.
- **Tactically-Aware Reasoning:** The Go-specific bias allows the model to dynamically focus its computational resources on the most critical and urgent areas of the board.
- **Synergistic Design:** The training philosophy and the network architecture are designed to work in concert. The deep intuition from pre-training makes the advanced reasoning of the Go-Attention layers even more effective.

The model is named **HANE** not only as an acronym (*Hybrid Attention with self-supervised kNowledge*), but also after the Go move “*hane*” (ハネ), a decisive tactic where a stone bends around the opponent’s position to seize initiative. Just as the *hane* move transforms the local fight into strategic advantage, our HANE model bends beyond conventional training by combining self-supervised intuition building with Go-specific attention, enabling it to capture both tactical urgency and global board balance with initiative.