

The 11 Core Principles of Guiguzi (鬼谷子)

Ancient Strategy Through the Lens of Modern Optimization

1. 揣摩 (Chuǎi Mó) — Assessment & Probing

The foundational skill. Like shooting in basketball — without it, all other techniques are guesses.

- 揣 (Chuǎi): Initial assessment — weighing and measuring the situation
- 摩 (Mó): Active probing — testing your read through subtle interaction

Methodology: Probe when targets are emotionally heightened (joy or fear), as they cannot mask true intentions. Methods include: peaceful approach, righteous challenge, flattery, provocation, intimidation, coercion, moral appeal, trust-building, incentive, and feigned humility.

ML Parallel: Building the reward model in RLHF. You must accurately read what the other party values before any policy update means anything. A bad reward model makes every downstream technique backfire.

2. 捭阖 (Bǎi Hé) — Opening & Closing

The rhythm of engagement. Alternating between expansion and contraction.

- 捭 (Bǎi) — **Open:** Make things public, draw out speech, encourage expression
- 阖 (Hé) — **Close:** Conceal, stay silent, consolidate

Core Insight: All things belong to either yin or yang — soft or hard, open or closed, relaxed or tense. The sage masters the timing of transitions between these states.

ML Parallel: The RLHF cycle — generate diverse outputs (open), receive feedback signal, update policy and tighten (close), repeat. Too much exploration and you never converge; too much exploitation and you collapse. PPO's clipping mechanism is a form of controlled closing.

3. 反应 (Fǎn Yìng) — Response & Mirroring

Extract signal by reflecting. Echo the other's words to make them reveal their true position.

Core Insight: By mirroring and responding strategically, you provoke reactions that expose hidden intentions. The response itself becomes a diagnostic tool.

ML Parallel: Adversarial probing — feeding inputs designed to elicit informative outputs. Like prompt engineering: the quality of what you extract depends on how you frame the reflection.

4. 内键 (Nèi Jiàn) — Inner Connection

Gain trust so deep that your counsel is adopted. Align your values with the decision-maker's until your interests merge.

- 内 (Nèi): Getting your strategy accepted
- 键 (Jiàn): Persisting in maintaining your position

Core Insight: "以情为重" — emotion is paramount. Win people through emotion, reason, and righteousness. The goal is to become so aligned with the decision-maker that you can "come and go freely."

ML Parallel: Reward alignment — ensuring your objective function matches the principal's. When inner alignment is achieved, the agent operates with full autonomy because its goals are indistinguishable from the principal's.

5. 抵巇 (Dǐ Xī) — Filling Gaps & Exploiting Cracks

Detect vulnerabilities in systems or relationships. Then decide: repair them or exploit them.

Core Insight: All systems develop cracks over time. The strategist's role is to identify these gaps before others do, then choose whether to seal them (defensive) or leverage them (offensive).

ML Parallel: Anomaly detection. Monitoring for distribution shifts, system vulnerabilities, or misalignments. The same detection capability serves both defense (patching) and offense (exploitation).

6. 飞箝 (Fēi Qián) — Flattery & Control

Elevate, then constrain. First "fly" (praise, validate) to open the target up, then "clamp" (control, direct).

- **飞 (Fēi):** Let emotions run free, encourage open expression through praise
- **箝 (Qián):** Clamp down, restrict movement, steer toward your objective

Methodology: Assess intelligence, weigh capability, estimate momentum — then use these as pivot points. Approach, follow, clamp to harmonize, adjust to fit intention.

ML Parallel: Exploration then exploitation. The "fly" phase is gathering information about the target's preference landscape. The "clamp" phase is optimizing over that landscape. Also mirrors curriculum learning — start permissive, then tighten constraints.

7. 忤合 (Wǔ Hé) — Opposition & Alliance

To align with one side, you must oppose another. No strategy serves two masters.

- **忤 (Wǔ):** Opposition, resistance, going against

- 合 (Hé): Alignment, joining, coming together

Core Insight: "计谋不两忠，必有反忤" — No plan can be loyal to two opposing sides. You must choose. The art is in knowing which side to align with and when to switch.伊尹 (Yi Yin) served both Shang Tang and Xia Jie five times each before committing — testing alignment through opposition.

ML Parallel: Multi-objective optimization with competing loss functions. You cannot minimize contradictory objectives simultaneously. The strategist must choose which objective to prioritize, understanding the tradeoff.

8. 权 (Quán) — Weighing & Authority

Measure before you act. Weigh all factors — capability, resources, timing, relationships — to determine the right course.

Core Insight: Authority comes from accurate assessment. The one who weighs correctly earns the right to decide. Rash action without weighing leads to failure.

ML Parallel: The evaluation phase in any decision pipeline. Hyperparameter tuning, A/B testing, cross-validation — all are forms of weighing before committing resources.

9. 谋 (Móu) — Planning & Strategy

Deep thought before action. "计谋者，存亡枢机" — Planning is the pivot between survival and destruction.

Three Prerequisites: 1. 先知 (Xiān Zhī): Know yourself and your opponent — gather intelligence 2. 周密安排: Arrange every detail — any single failure affects the whole 3. 环境优势: Leverage environmental advantages — fight on familiar ground

Core Insight: "谋定而后动" — Plan thoroughly, then act. Those who act before planning fail. Those who plan before acting succeed.

ML Parallel: The design phase — architecture selection, data pipeline planning, experiment design. Corresponds to the planning phase in model-based RL, where the agent simulates outcomes before committing to actions.

10. 转圆 (Zhuǎn Yuán) — Turning the Circle

Formless adaptation at speed. Like a rolling sphere — no fixed shape, infinite configurations, unstoppable momentum.

Core Insight: The sage's wisdom turns like a never-stopping sphere. Round strategies allow language and plans to rotate freely; square strategies stabilize and anchor. Master both: fluidity for maneuvering, structure for execution.

Key Principle: "以速取胜" — Victory through speed. Capture opportunities instantly, decide without hesitation, act before the opponent can react.

ML Parallel: Online learning and real-time adaptation. The model that adapts fastest to distribution shifts wins. Also mirrors the explore-exploit tradeoff in multi-armed bandits — the "circle" never stops turning.

11. 符言 (Fú Yán) — Keeping One's Word


Credibility is infrastructure. The leader's code: 安 (calm), 徐 (measured), 正 (upright), 静 (still).

Nine Governance Principles: 1. **主位:** Maintain calm authority — observe the world from a position of stability 2. **主明:** See clearly — remove all obstacles to perception 3. **主听:** Listen broadly — "以天下之耳听者，则无不闻" 4. **主赏:** Reward with integrity — credibility makes incentives work 5. **主问:** Investigate thoroughly — examine all factors before deciding 6. **主因:** Follow natural principles — align rewards with actual needs 7. **主周:** Know comprehensively — understand both internal and external dynamics 8. **主参:** Cross-reference — use multiple sources to verify truth 9. **主名:** Name accurately — ensure labels match reality ("名实相符")

Core Insight: "赏罚必以诚信征服天下人" — Reward and punishment must be grounded in credibility to command the world. Without trust, no technique works.

ML Parallel: The training infrastructure itself. Model reliability, reproducibility, consistent evaluation metrics. A model that produces unreliable outputs — no matter how sophisticated its architecture — cannot be trusted or deployed. Credibility is the foundation that makes all optimization meaningful.

The System Architecture



Key Insight: Guiguzi vs. Sun Tzu

Sun Tzu operates at the **policy optimization** level — given the battlefield, given the enemy, find the best response. Reactive.

Guiguzi operates at the **environment design** level — shape the game itself before the opponent even begins to optimize. Proactive.

Sun Tzu: "The enemy is here. What now?" Guiguzi: "I decided where 'here' is."

One is the best general on the battlefield. The other chose which battlefield, when, and whether there even needs to be one.

The framework is neutral — it describes optimization dynamics that are mathematical in nature. What you optimize for is the human choice.