AI and Mahjong: A New Frontier in Game Intelligence

Chi-Kwong Li
College of William & Mary

March 8, 2025

Why Mahjong?

A Complex Game

Mahjong challenges AI with:

Why Mahjong?

A Complex Game

Mahjong challenges AI with:

- **Imperfect Information**: Hidden tiles.

Why Mahjong?

A Complex Game

Mahjong challenges AI with:

- **Imperfect Information**: Hidden tiles.
- **Multiplayer**: 3–4 players compete.

Why Mahjong?

A Complex Game

Mahjong challenges AI with:

- **Imperfect Information**: Hidden tiles.
- **Multiplayer**: 3–4 players compete.
- **Randomness**: Tile draws add chance.

Why Mahjong?

A Complex Game

Mahjong challenges AI with:

- **Imperfect Information**: Hidden tiles.
- **Multiplayer**: 3–4 players compete.
- **Randomness**: Tile draws add chance.
- **Scoring**: Complex winning patterns.

Why Mahjong?

A Complex Game

Mahjong challenges AI with:

- **Imperfect Information**: Hidden tiles.
- **Multiplayer**: 3–4 players compete.
- **Randomness**: Tile draws add chance.
- **Scoring**: Complex winning patterns.

"Harder than Chess or Go" — AI Researchers

Mahjong Challenges for AI

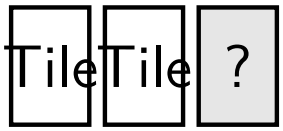

Mahjong Challenges for AI

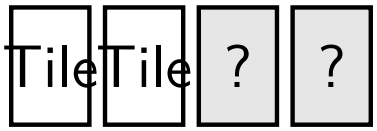

- **Uncertainty**: Only see own hand + discards.

Mahjong Challenges for AI

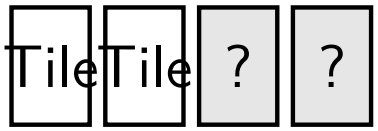

Opponents' Hidden Hands

- **Uncertainty**: Only see own hand + discards.
- **Strategy**: Model multiple agents' moves.

Mahjong Challenges for AI

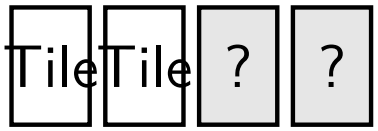

Opponents' Hidden Hands

- **Uncertainty**: Only see own hand + discards.
- **Strategy**: Model multiple agents' moves.
- **Computation**: Evaluate 10^{48} game states (Riichi Mahjong).

Mahjong Challenges for AI

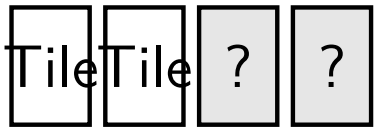

Opponents' Hidden Hands

- **Uncertainty**: Only see own hand + discards.
- **Strategy**: Model multiple agents' moves.
- **Computation**: Evaluate 10^{48} game states (Riichi Mahjong).

Goal

Build AIs that play like—or beat—top humans!

Suphx: Superhuman Mahjong AI

Microsoft Research Asia (2020)

Suphx: Superhuman Mahjong AI

Microsoft Research Asia (2020)

- **Approach**: Deep RL + 5 CNNs (discard, Riichi, etc.).

Suphx: Superhuman Mahjong AI

Microsoft Research Asia (2020)

- **Approach**: Deep RL + 5 CNNs (discard, Riichi, etc.).
- **Training**: Supervised (Tenhou data) + self-play RL.

Suphx: Superhuman Mahjong AI

Microsoft Research Asia (2020)

- **Approach**: Deep RL + 5 CNNs (discard, Riichi, etc.).
- **Training**: Supervised (Tenhou data) + self-play RL.
- **Result**: 10th dan, beats 99.99% of players.

Suphx: Superhuman Mahjong AI

Microsoft Research Asia (2020)

- **Approach**: Deep RL + 5 CNNs (discard, Riichi, etc.).
- **Training**: Supervised (Tenhou data) + self-play RL.
- **Result**: 10th dan, beats 99.99% of players.
- **Link**: arXiv:2003.13590

Suphx: Superhuman Mahjong AI

Microsoft Research Asia (2020)

- **Approach**: Deep RL + 5 CNNs (discard, Riichi, etc.).
- **Training**: Supervised (Tenhou data) + self-play RL.
- **Result**: 10th dan, beats 99.99% of players.
- **Link**: arXiv:2003.13590

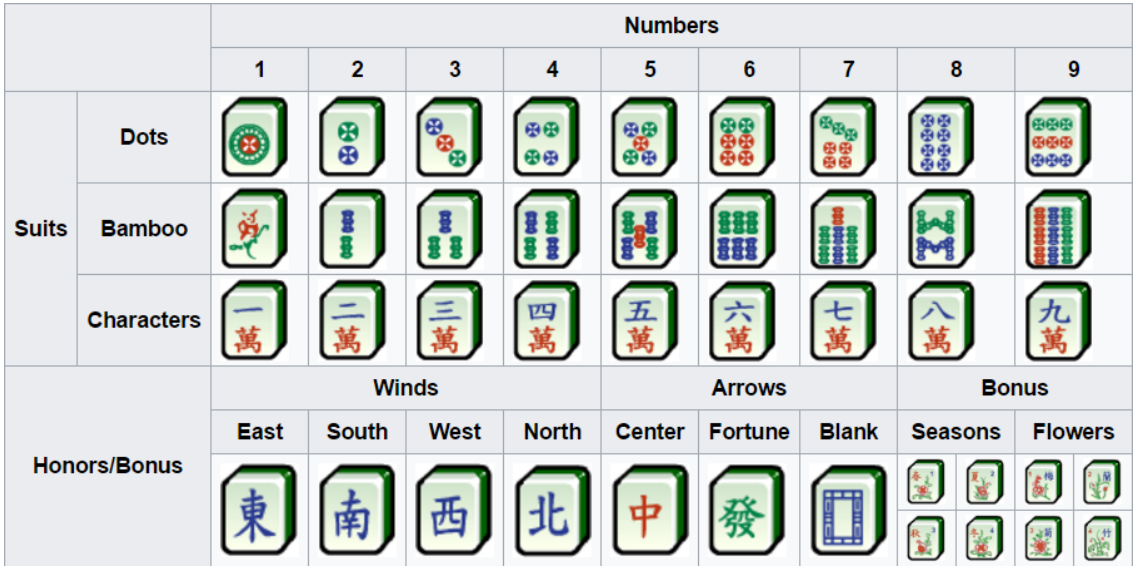

Suphx excels in defense and style.

Beyond Suphx

Meowjong (2022)

Beyond Suphx

Meowjong (2022)

- 3-player Sanma Mahjong.

Beyond Suphx

Meowjong (2022)

- 3-player Sanma Mahjong.
- RL + Monte Carlo policy gradient.

Beyond Suphx

Meowjong (2022)

- 3-player Sanma Mahjong.
- RL + Monte Carlo policy gradient.
- arXiv:2202.12847

Beyond Suphx

Meowjong (2022)

- 3-player Sanma Mahjong.
- RL + Monte Carlo policy gradient.
- arXiv:2202.12847

Tjong (2024)

Beyond Suphx

Meowjong (2022)

- 3-player Sanma Mahjong.
- RL + Monte Carlo policy gradient.
- arXiv:2202.12847

Tjong (2024)

- Transformer-based, hierarchical decisions.

Beyond Suphx

Meowjong (2022)

- 3-player Sanma Mahjong.
- RL + Monte Carlo policy gradient.
- arXiv:2202.12847

Tjong (2024)

- Transformer-based, hierarchical decisions.
- Efficient for sparse rewards.

Beyond Suphx

Meowjong (2022)

- 3-player Sanma Mahjong.
- RL + Monte Carlo policy gradient.
- arXiv:2202.12847

Tjong (2024)

- Transformer-based, hierarchical decisions.
- Efficient for sparse rewards.
- Published in CAAI Transactions (https://doi.org/10.1049/cit2.12298).

Future of AI in Mahjong

- **New Variants**: Chinese, International Mahjong.

Future of AI in Mahjong

- **New Variants**: Chinese, International Mahjong.
- **AI Assistants**: Teach players (e.g., Tencent Yuanbao, 2025 X posts).

Future of AI in Mahjong

- **New Variants**: Chinese, International Mahjong.
- **AI Assistants**: Teach players (e.g., Tencent Yuanbao, 2025 X posts).
- **Real-World Impact**: Multi-agent systems, decision-making under uncertainty.

Future of AI in Mahjong

- **New Variants**: Chinese, International Mahjong.
- **AI Assistants**: Teach players (e.g., Tencent Yuanbao, 2025 X posts).
- **Real-World Impact**: Multi-agent systems, decision-making under uncertainty.

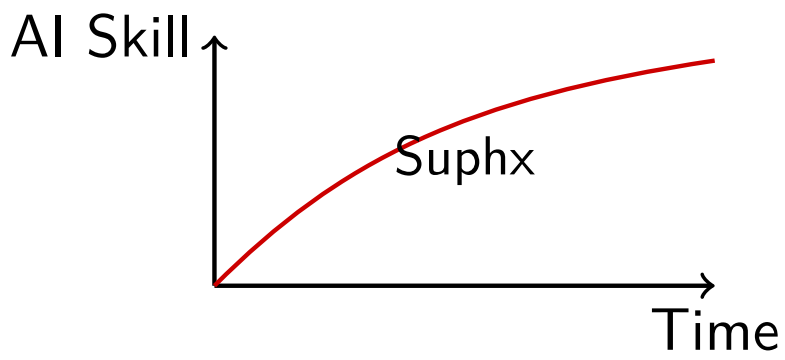

Future of AI in Mahjong

- **New Variants**: Chinese, International Mahjong.
- **AI Assistants**: Teach players (e.g., Tencent Yuanbao, 2025 X posts).
- **Real-World Impact**: Multi-agent systems, decision-making under uncertainty.

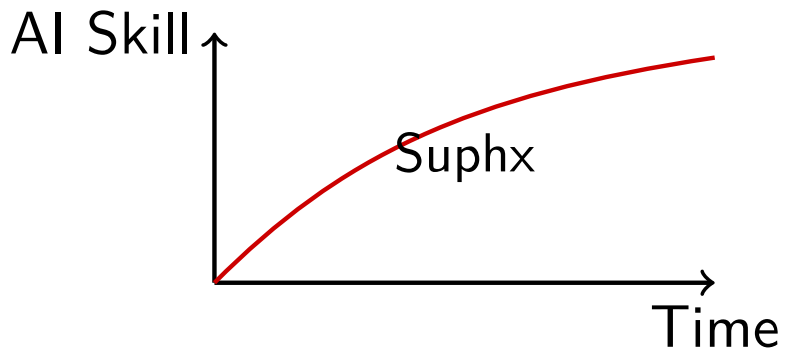

"Mahjong AI is just beginning!"

Conclusion

AI meets Mahjong:

A testbed for intelligence, strategy, and innovation.

Thank you!