

CivRealm: A Learning and Reasoning Odyssey in Civilization for Decision-Making Agents

Siyuan Qi, Shuo Chen, Yexin Li, Xiangyu Kong, Junqi Wang, Bangcheng Yang, Pring Wong, Yifan Zhong, Xiaoyuan Zhang, Zhaowei Zhang, Nian Liu, Yaodong Yang, Song-Chun Zhu

@ BIGAI (Beijing Institute for General Artificial Intelligence)



| Development of Decision-Making Algorithms

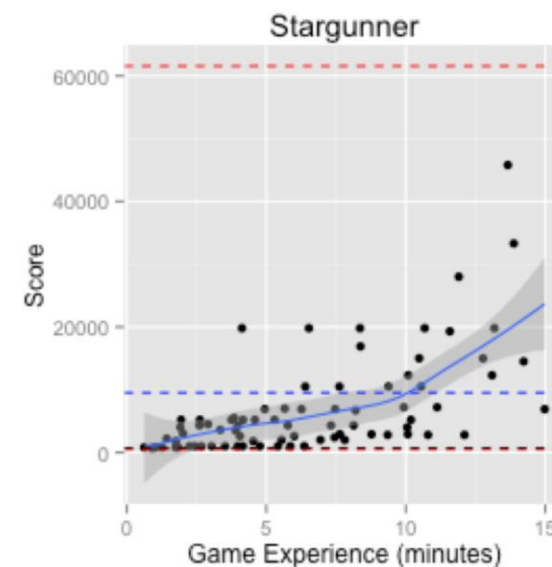
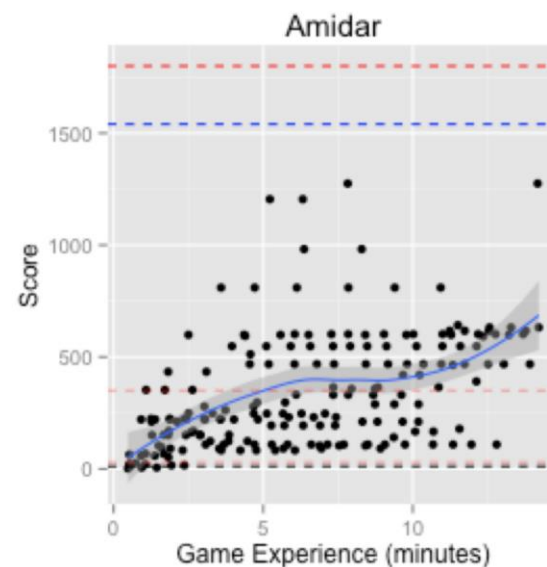

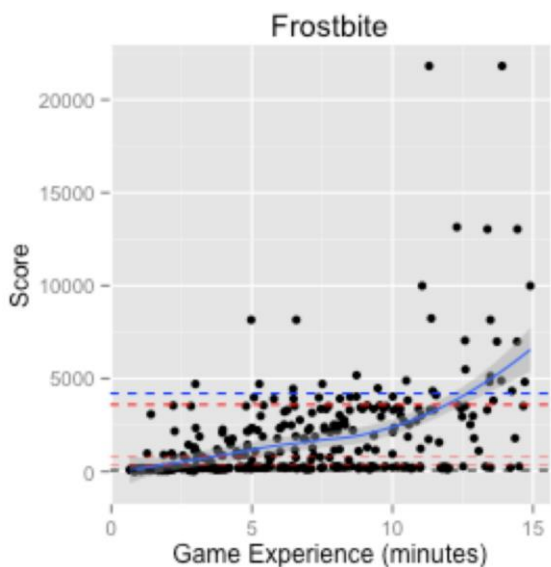


DQN, 2013



AlphaGo, 2016

Challenge I: Learning at Human-Level Speed




Humans reach “expert” performance within 15 mins, and exceed DDQN’s 10- and 25- million-frame with a few minutes.¹



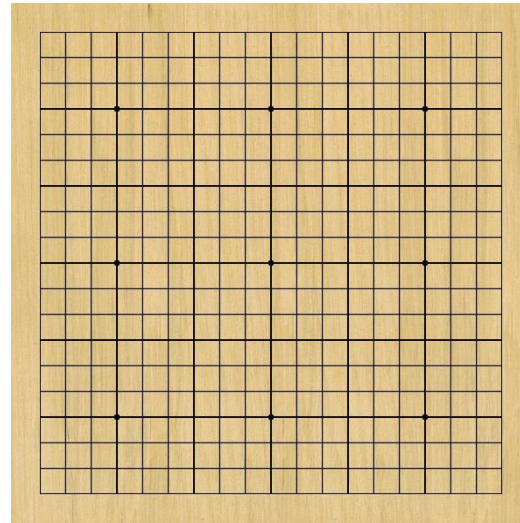
¹Tsividis, Pedro A., et al. “Human learning in Atari.”, 2017.

| Challenge II: Reasoning in an Ever-Expanding Space


- 2009, 9x9 game, Fuego beat Zhou Junxun (9 dan professional).²
- 2009, 19x19 game, MoGo won against Zhou Junxun, with 7 handicap stones.²



9×9



19×19



23×23

Stone to **Fire**: 18,000,000 million years

Steam engines to **Computers**: 247 years



² Gelly, Sylvain, et al. “The grand challenge of computer Go: Monte Carlo tree search and extensions.”, 2012.

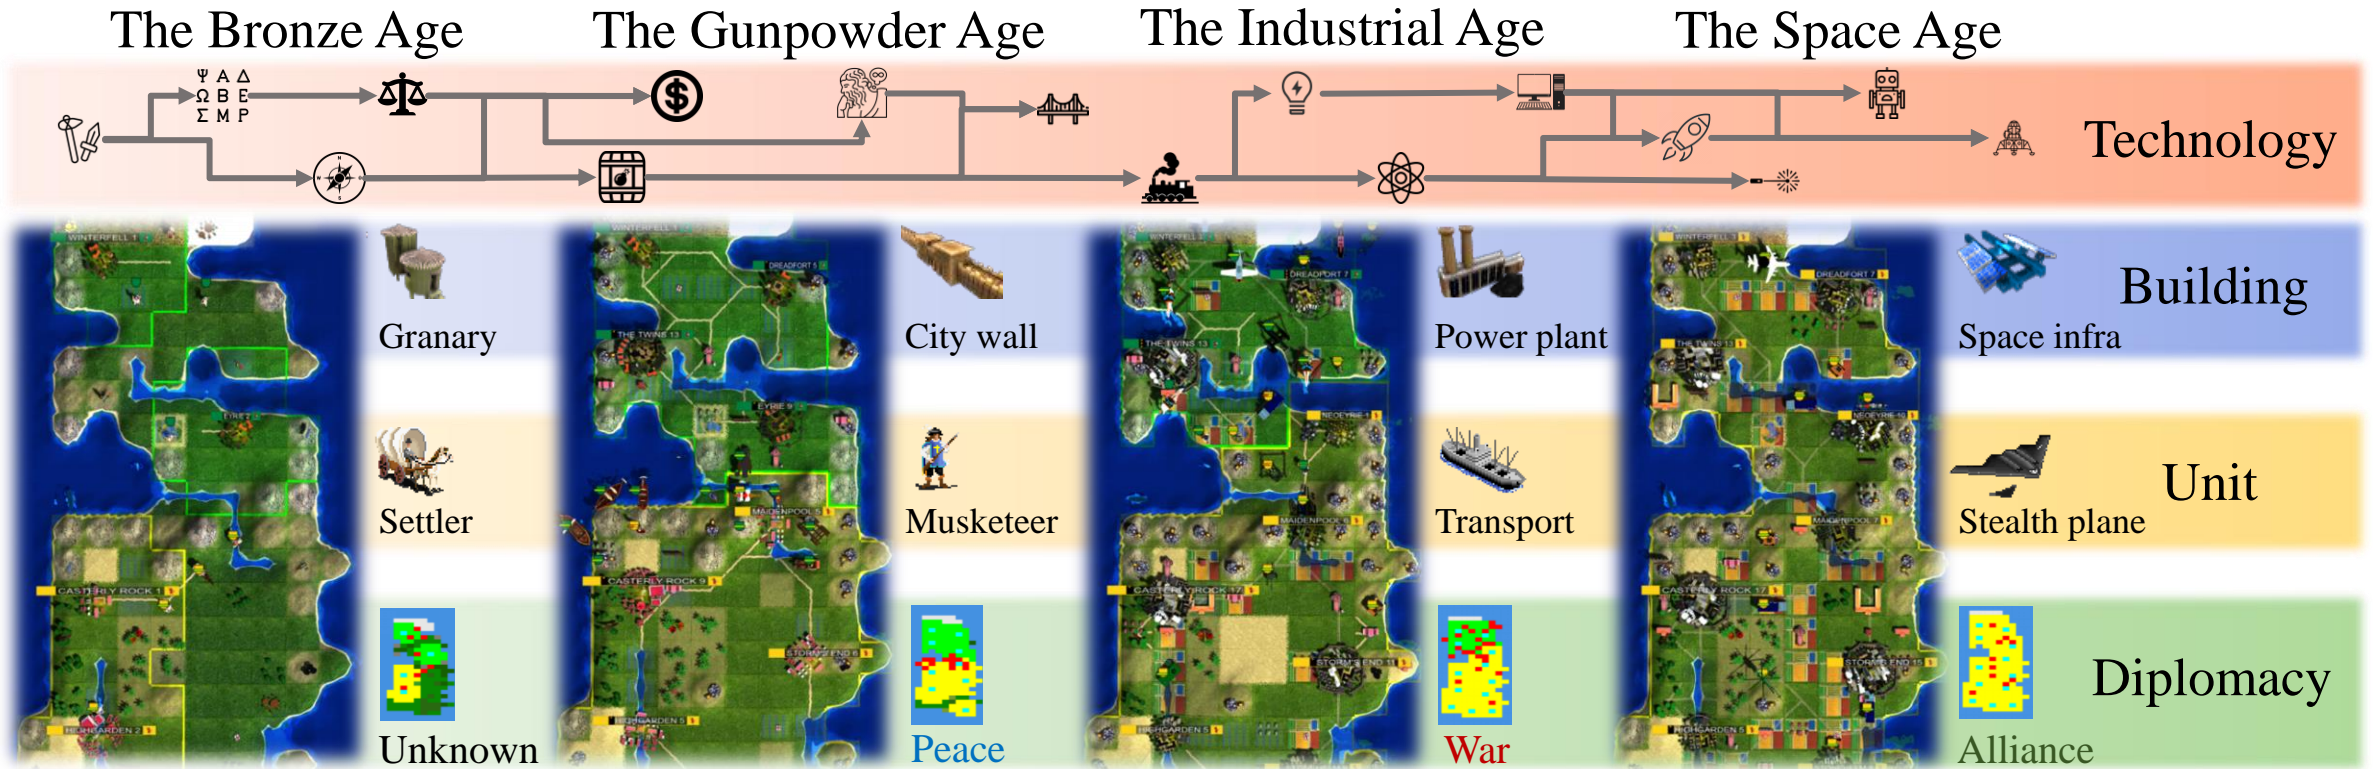
| CivRealm: A Learning and Reasoning Odyssey

✓ The *Civilization* Game:

- ❑ Multi-agent interaction, a general-sum game.
- ❑ Long-term planning, multi-task generalization.
- ❑ Interactive and open-ended decision making.



Expansion of State and Action Spaces



| Challenges of CivRealm

Environment	Imperfect info	Stochastic	Multi-goal	Dynamic space	Multi-agent	General-sum	Changing players	Comm.	Tensor & Lang.
MineDojo [18]	✓	✓	✓	✓	✗	✓	✗	✗	✓
MPE [50]	✗	✓	✗	✗	✓	✓	✗	✓	✗
Hanabi [5]	✓	✗	✓	✗	✓	✓	✗	✓	✗
Hold'em [10]	✓	✗	✗	✗	✓	✗	✗	✗	✗
Diplomacy [51]	✗	✗	✗	✗	✓	✗	✓	✓	✓
Melting pot [43]	✓	✓	✓	✗	✓	✓	✓	✗	✗
Google Football [42]	✗	✓	✗	✗	✓	✗	✗	✗	✗
Stratego [53]	✓	✗	✗	✗	✓	✗	✗	✗	✗
SMAC [58]	✓	✗	✗	✗	✓	✗	✗	✗	✗
Dota 2 [8]	✓	✓	✗	✓	✓	✗	✗	✗	✗
StarCraft II [79]	✓	✗	✗	✓	✓	✗	✗	✗	✗
<i>CivRealm</i>	✓	✓	✓	✓	✓	✓	✓	✓	✓



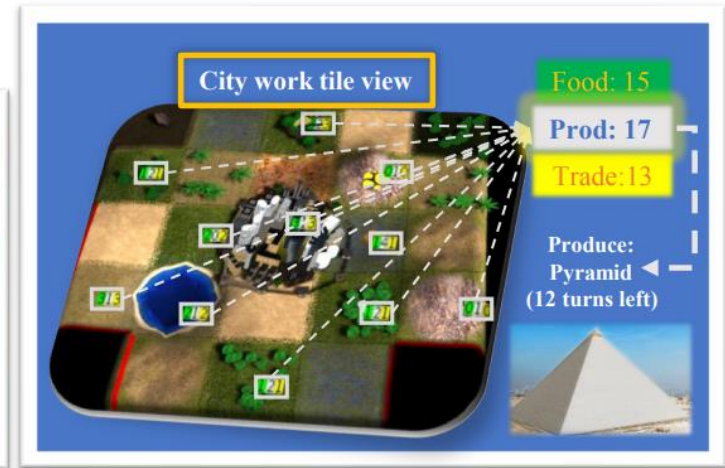
Mini-Games: Development, Battle, Diplomacy



Development: TransBuildCity



Development: WorkerBuildInfra



Development: CityTileWonder



Battle: LandBattleModern



Battle: NavalBattleModern

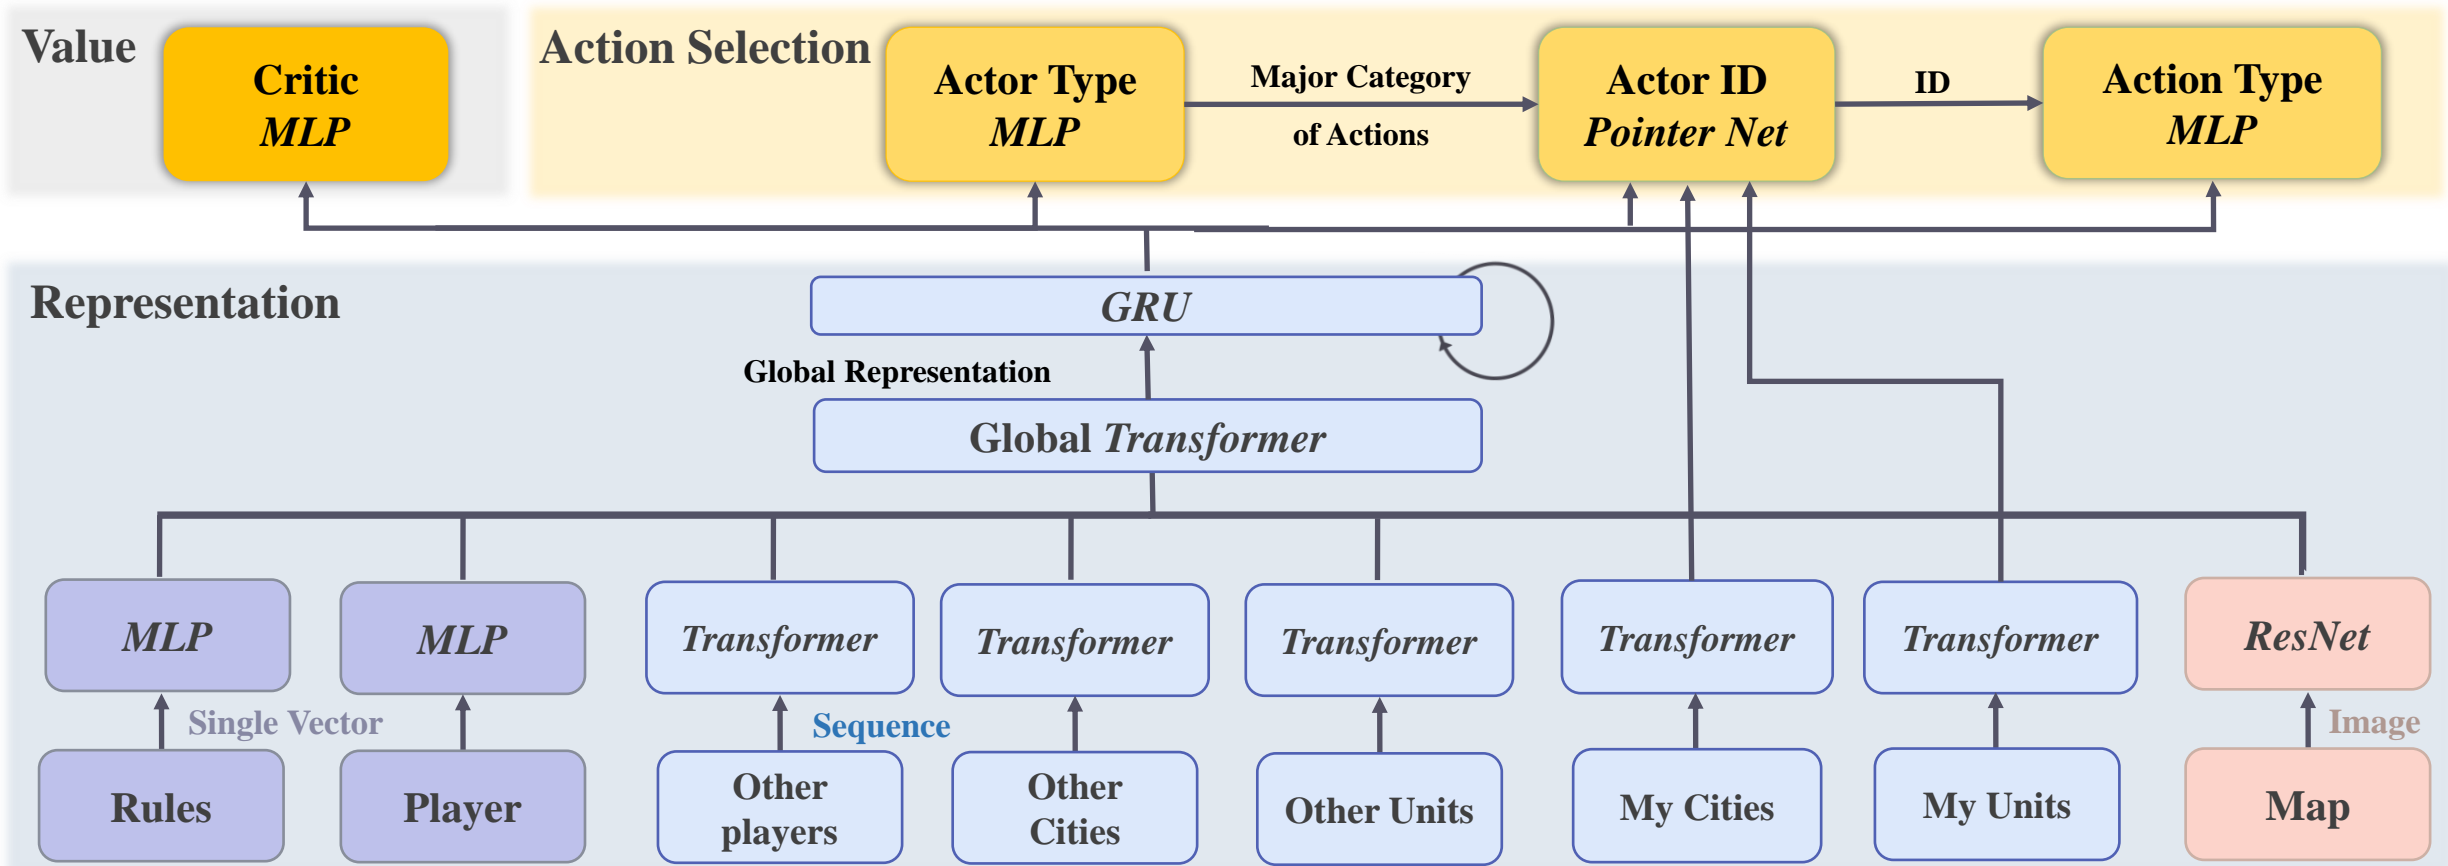


Diplomacy: TradeTechs




Tensor-based Reinforcement Learning Agent

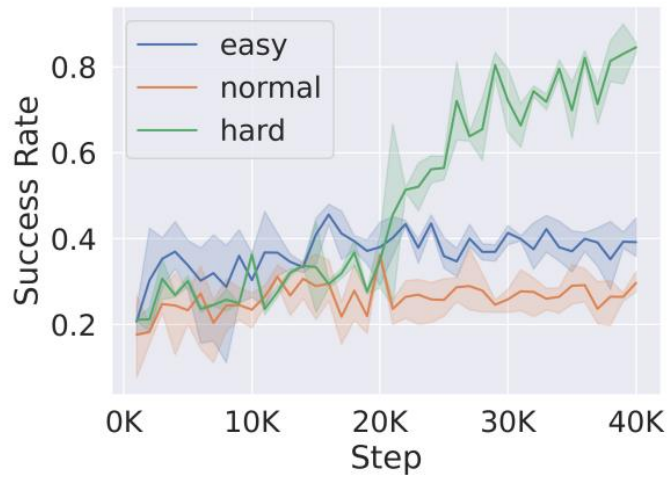
- Static state and action space
- Data-hungry




Tensor-based Reinforcement Learning Agent




(a) CityTileWonder



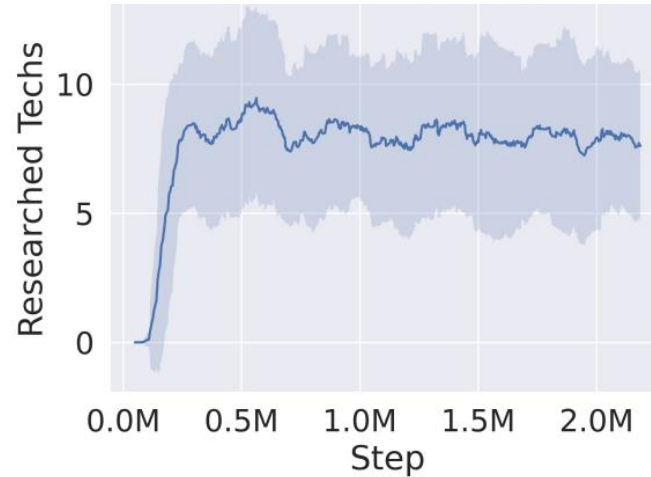
(b) SettlerBuildCity




(c) LandBattleDefendCity



(a) Game scores



(b) technologies researched

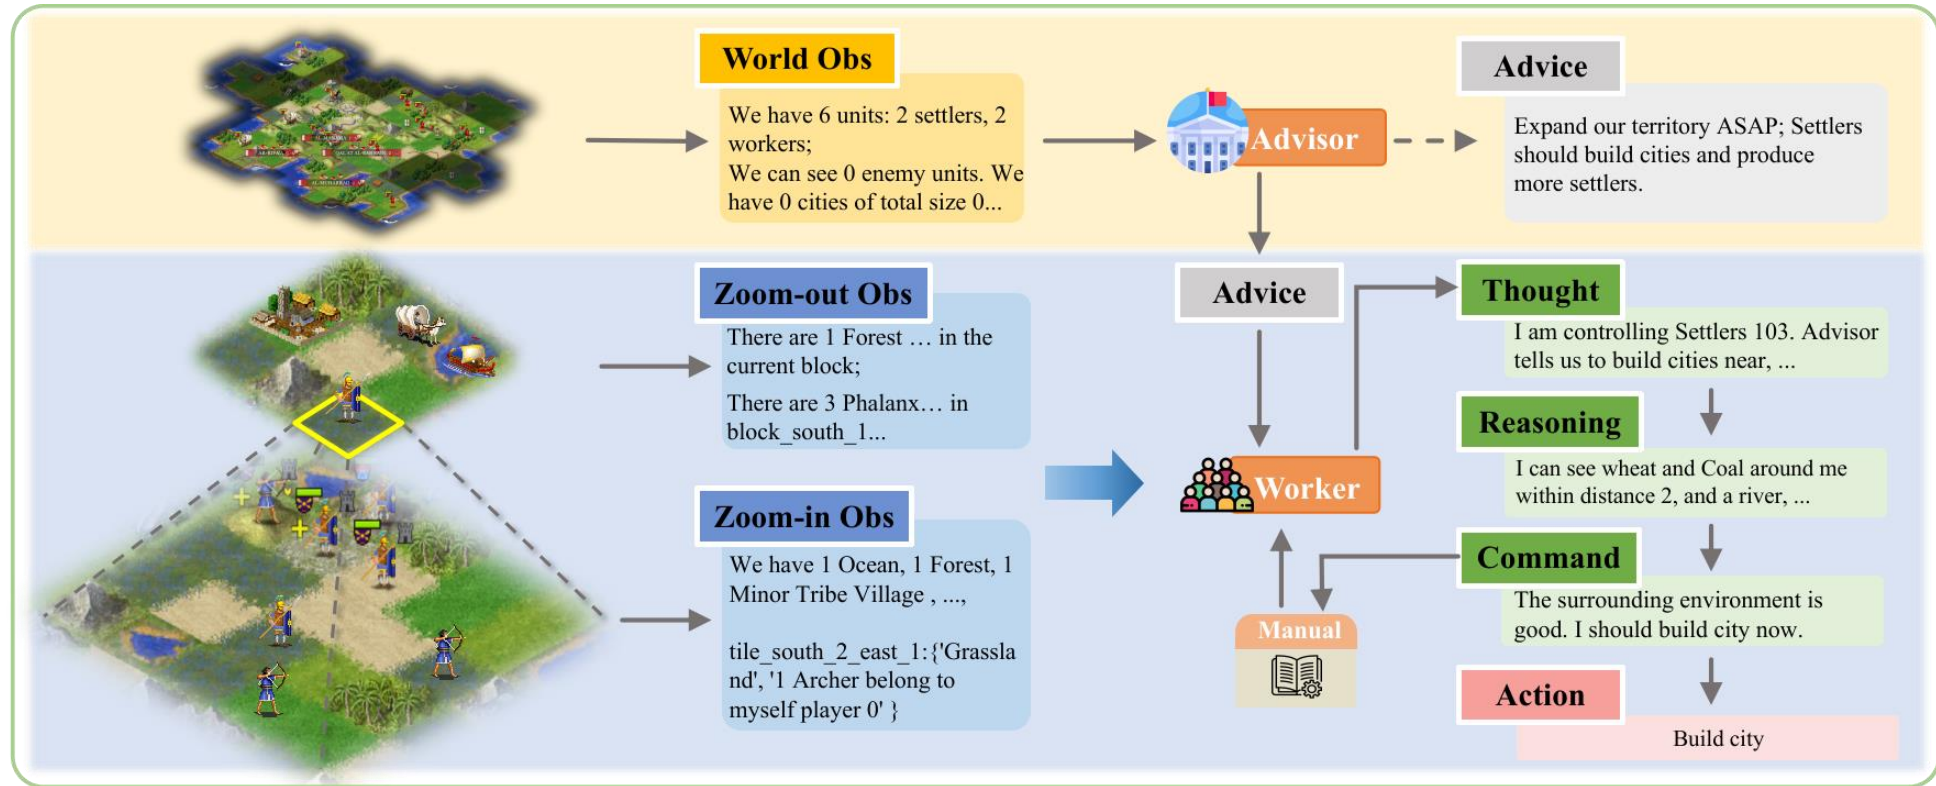
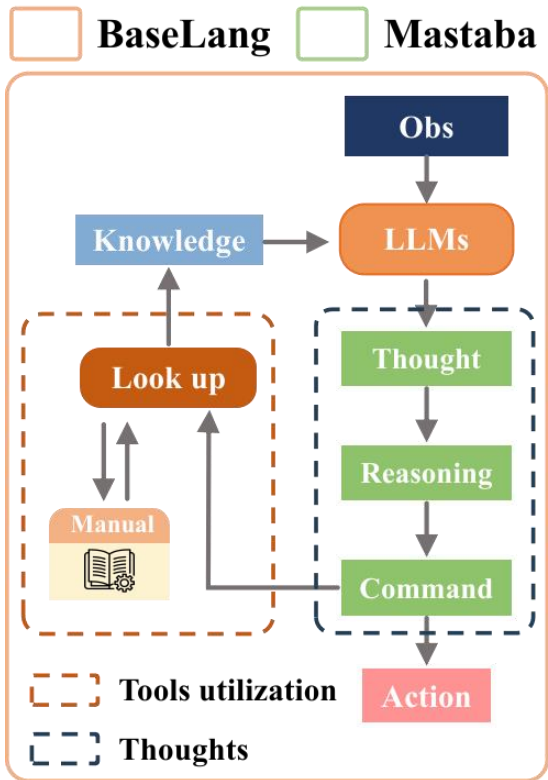


(c) units built



Language-based Agents


- Hierarchical observation and decisions
- Communication between agents
- Grounding problem



| Language-based Agents

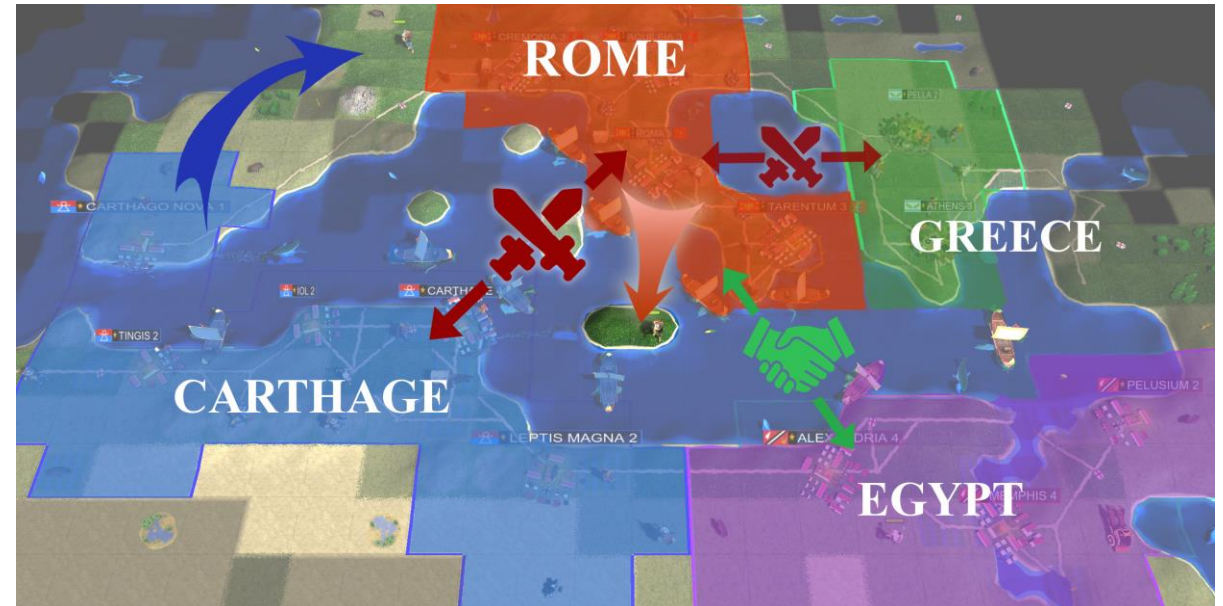
Evolution of the Civilization:

Inception, establishment, expansion, invasion, and collapse.



| Take-away

- CivRealm: an interactive environment inspired by Civilization
 - Learning with prior knowledge
 - Reasoning over an expanding space
- Two kinds of APIs: **tensor-based** and **language-based**
- Two flavors of baselines provided and analyzed



CivRealm: A Learning and Reasoning Odyssey in Civilization for Decision-Making Agents

Siyuan Qi, Shuo Chen, Yexin Li, Xiangyu Kong, Junqi Wang, Bangcheng Yang, Pring Wong, Yifan Zhong, Xiaoyuan Zhang, Zhaowei Zhang, Nian Liu, Yaodong Yang, Song-Chun Zhu

@ BIGAI (Beijing Institute for General Artificial Intelligence)

