

Tutorial @ IEEE CoG'23

TAG: Advancements in AI-Driven Tabletop Games

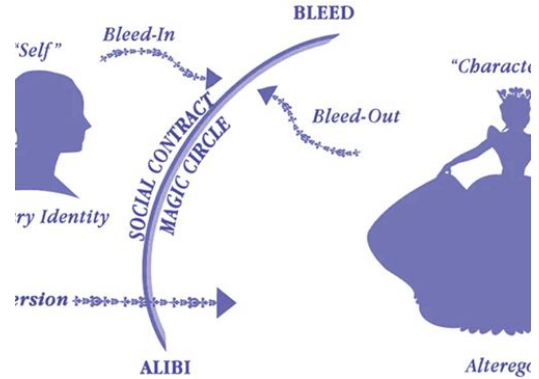

**Raluca D. Gaina, Martin Balla,
(James Goodman), Diego Perez Liebana**

AI and Tabletop Games

Tabletop Board Games



Tabletop Games and their Social Impact



...dientes Personal für therapeutische Maßnahmen siehe Code-Book.

TUBUS	LEBENSUNTER					AT-Nr.		
Ja	6	5	4	3	2	1	tot	22
Nein								

M 30 J.

ORT	MABNAHME
Sichtungsplatz	Sichtung
Schockraum	Primary Survey
NFA / eFAST	eFAST
Behandlungsber.	Röntgen
Behandlungsber.	CT
Behandlungsber.	ETC
Bariuminzie	DCS
OP	TASC
AWR	Kons. Ther.
ICU	Keine Ther.
IMC	Keine weit. Maßnahmen
NST	


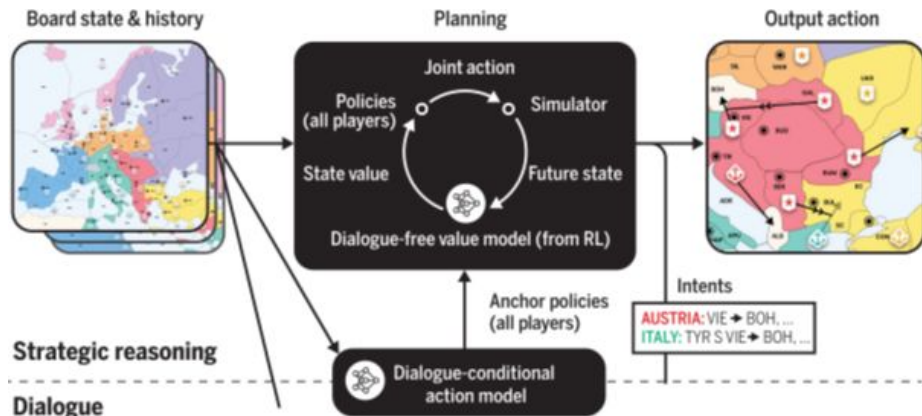
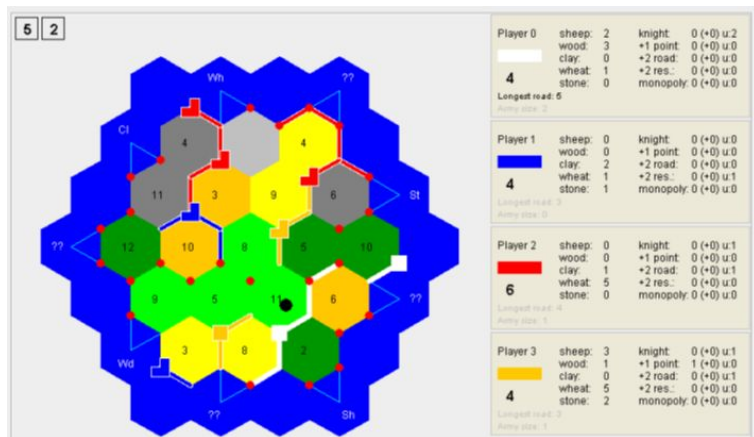
ZEIT, Fertig Stop Go 1 2 3 4 5 6

Mögliche Maßnahmen:

- Konventionelles Röntgen (Zeit +1)
- eFAST (Zeit +0)
- Primary Survey (Zeit +0)
- OP

Arzt Anästhesie
Arzt Chirurgie
Pflege ZN

Tabletop Games and AI



Tabletop Games Framework (TAG)

TAG: Tabletop Games Framework

<https://tinyurl.com/TAGframework>

Summarizing backbones of tabletop games, easy to implement / prototype games

- **Actions:** things the player does
- **Rules:** things the game does
- **Turn Order:** what order to players go in?
- **Game Phase:** time frames with different actions/rules
- **Components:** game objects/pieces, the state of which is changed by actions/rules

Easy to plug in a bunch of AI players and analysis tools

- **Game balancing:** is this game fair to all players? does it have the right difficulty?
- **Game testing:** does this game work right?
- **Information extraction:** what does this game even mean?

TAG: Framework Games

- **Simple:** Tic Tac Toe, Dots & Boxes
- **Social, hidden information:** Love Letter
- **Deck building:** Dominion
- **Casino:** Poker Texas Hold'em, Blackjack
- **Card:** Uno, Virus!, Exploding Kittens, Hanabi

- **Strategy:** Colt Express, Pandemic, Diamant, Terraforming Mars, Settlers of Catan, BattleLore, Stratego, Puerto Rico, 7 Wonders
- **In-progress:** Descent2e, Monopoly Deal
- **External contributions:** Sushi Go!



Modern Tabletop Games

- **Multi-player dynamics:** more than 2 players. Winning? Coalitions?

Modern

- Mult



Modern Tabletop Games

- **Multi-player dynamics:** more than 2 players. Winning? Coalitions?
- **Many strategies available:** no fixed way of winning. Dominant strategy? How many? Explainability? More interesting, or too much cognitive load?

Model

- M
- M

ma



? How

Modern Tabletop Games

- **Multi-player dynamics:** more than 2 players. Winning? Coalitions?
- **Many strategies available:** no fixed way of winning. Dominant strategy? How many? Explainability? More interesting, or too much cognitive load?
- **Large design spaces:** many parameters. Variations? Balance?
- **Different definitions of balance:** all cards equally strong? One overpowered, but counter available? Or CHAOS!

M



N

,

Modern Tabletop Games

- **Multi-player dynamics:** more than 2 players. Winning? Coalitions?
- **Many strategies available:** no fixed way of winning. Dominant strategy? How many? Explainability? More interesting, or too much cognitive load?
- **Large design spaces:** many parameters. Variations? Balance?
- **Different definitions of balance:** all cards equally strong? One overpowered, but counter available? Or CHAOS!
- **Meaningful metrics:** inform game design decisions.

AI in Tabletop Games: 3 Perspectives

1. AI for playing tabletop games

Goodman, J., Perez-Liebana, D. and Lucas, S., MultiTree MCTS in Tabletop Games. In 2022 IEEE Conference on Games (CoG). IEEE

One tree per player. Each iteration traverses and updates multiple trees.

AI in Tabletop Games: 3 Perspectives

2. AI for tabletop games evaluation and statistics extraction

Gaina, R.D., Goodman, J. and Perez-Liebana, D., 2021, October. TAG: Terraforming Mars. In Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (Vol. 17, No. 1, pp. 148-155).



AI in Tabletop Games: 3 Perspectives

2. AI for tabletop games evaluation and statistics extraction

Gaina, R.D. and Balla, M.,
2022, August. TAG:
Pandemic Competition. In
2022 IEEE Conference on
Games (CoG) (pp. 552-559).
IEEE.



AI in Tabletop Games: 3 Perspectives

2. AI for tabletop games evaluation and statistics extraction

Goodman, J.,
Perez-Liebana, D., Lucas,
S.M., 2023, August. A case
study in AI-assisted board
game design. In 2023 IEEE
Conference on Games
(CoG). IEEE.



Thursday, 14:30, ISEC 136

AI in Tabletop Games: 3 Perspectives

2. AI for tabletop games evaluation and statistics extraction

New PhD students starting Sept 2023: Joshua Kritz and Dien Ngyuen – watch out for their upcoming work!

TAG Demo @ UKGE'23 – more later in the talk

AI in Tabletop Games: 3 Perspectives

3. AI for content generation in tabletop games

- TAG in the future?
- **Hybrid games:** great medium!
 - Add new scenarios / adventures / stories
 - Update game balance
 - Customised / personalised content
 - Procedurally generated stories and event results



TAG for Education

- Framework already used in **university courses**. Example assignments:
 - Create AI players for a game
 - Analysis of AI play / game
 - Implementation of new games (or fill in existing skeleton with key parts of the code)
- Also in many **UG, PG, PhD projects** world-wide

- **Understand code in context** and in larger framework
- **Practical** problem solving and applications
- **Real-world connection**: students can play the physical board games to understand the problem better (and it's fun!), or they might know the games beforehand (due to use of modern games).

PyTAG



**BEST
PAPER
AWARD**
(nomination)

PyTAG

9:00-10:40



Best of CoG
Chair: Jialin Liu
Location: ISEC 102


Outline

- PyTAG overview
- Reinforcement Learning
- Challenges and solutions
- Usage
- Opportunities


PyTAG

- An interface to interact with TAG from python
- Support for as much games as possible
- Currently aimed at supporting Reinforcement Learning
- Lots of other possibilities available
 - Access to game data
 - Access to forward models
 - Access to game metrics
 - Ability to run any game with any players

Reinforcement Learning



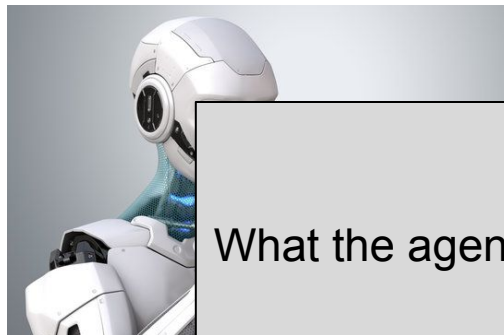
Observation,
Reward



Action



Reinforcement Learning



Observation,
Reward



What the agent sees at each step of the game



Observation




Observation



Observation - Approach 1

Vectorised observation representation



Observation - Approach 1

Vectorised observation representation

TicTacToe: encoding of the board

Diamant: Gems, score, number of players in cave..

Exploding Kittens: Card encodings + additional features (i.e: turn phase, alive players, #cards in draw pile)

Issues?

	X	O
	X	
O	X	



Observation - Approach 2

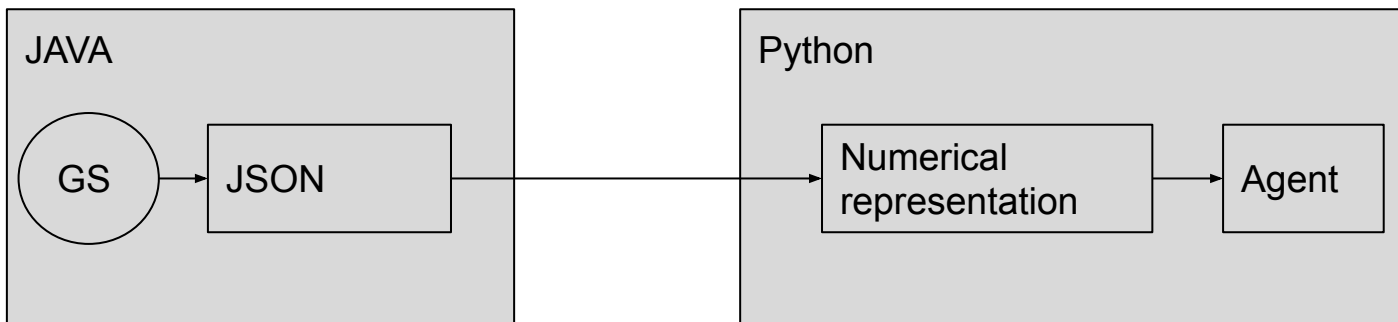
JSON observation representation

For more complex games - hard to vectorise

Idea: extract information from the GameState



- Process the information on the agent side
- More flexibility on how to process information



Observation - Approach 2

JSON observation representation

SushiGo

- cards in hand
- cards in front of player
- cards in front of opponent
- Round count
- Player scores



Observation Summary


Vectorisation

- Process on the JAVA side
- faster

JSON

- Easier to handle on the JAVA side
- just convert information from GS to JSON
- Process on the python side
- More flexibility

Actions



Observation,
Reward



Action



What the agent can do at each step of the game

Actions

Exploding Kittens

- 1, Play any card from hand
- 2, Draw a card to end turn



Actions

Exploding Kittens

- 1, Play any card from hand
- 2, Draw a card to end turn

All cards have special effects

- Draw a card from the opponent
- Which player and card?



Actions

Exploding Kittens

- 1, Play any card from hand
- 2, Draw a card to end turn

All cards have special effects

- Draw a card from the opponent
- Which player and card?

Reactions?

- Do I want to use my nope card?



Actions

What are the actions?

- Play any card in hand
- Draw a card
- Reactions?

Dynamic

Combinatorial




Actions - Action Masking


TicTacToe

9 possible actions

mask out the not available actions



Action logits
[0.7, 0.3, 0.1, 0.5, ...]



Actions - Action Masking

TicTacToe

9 possible actions

mask out the not available actions




Action logits

[0.7, 0.3, 0.1, 0.5, ...]

Action mask

[0, 0, 0, 1, ...]



Actions - Action Masking

TicTacToe

9 possible actions

mask out the not available actions



Action logits


[0.7, 0.3, 0.1, 0.5, ...]

Action mask

[0, 0, 0, 1, ...]

Masked action logits

[1e-8, 1e-8, 1e-8, 0.5, ...]





Actions - Action Trees


Break down the action space further

Idea: choose action type and then arguments

Choose **row** first and then **column**



Reward



Observation,
Reward

Action



Reward

Score

- Settlers of Catan - get 10 points to win,
- Diamant - player with most points win

Win/loss

- TicTacToe - get 3 in a row
- Exploding Kittens - don't explode
- Pandemic - cure all diseases

Reward

Win/loss often works well

- Some games have long reward horizons
- Stratego may take 1000s of steps until an outcome is reached

Reward shaping

- Can implement a heuristics to alter the scoring
- May lead to suboptimal behaviours (high score \neq winning)

What do I need to do RL for a new Game in TAG?

- Observation space
 - vectoriser interface
 - JSON interface
- Action space
 - Fixed size with masking
- (optional) Reward function

RL challenges

So far

- Observation spaces
- Action spaces
- Rewards

In addition

- Multi-agent dynamics
- Coop vs Competition
- Communication
- Natural Language

PyTAG summary

Python interface to interact with TAG

Many challenges: observations, actions, rewards...

Multi-agent dynamics

Lower entry barrier for TAG - can avoid JAVA programming

Further use-cases?

Come talk to us about how this could be used in your research!

PyTAG summary

9:00-10:40



Best of CoG
Chair: Jialin Liu
Location: ISEC 102

Presented Thursday 9AM, ISEC 102

Github: <https://github.com/martinballa/PyTAG>




PyTAG: Challenges and Opportunities for Reinforcement Learning in Tabletop Games

Martin Balla, George E.M. Long, Dominik Jeurissen, James Goodman, Raluca D. Gaina, Diego Perez-Liebana
Queen Mary University of London


{m.balla, g.e.m.long, d.jeurissen, james.goodman, r.d.gaina, diego.perez}@qmul.ac.uk

Tabletop R&D: AI for Game Analysis

Bringing AI to the table: the Process



How do game designers make board games?



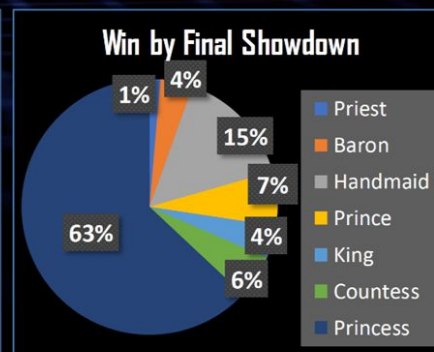
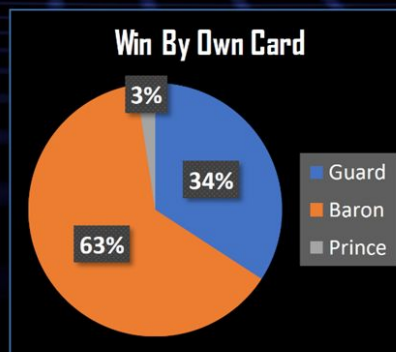

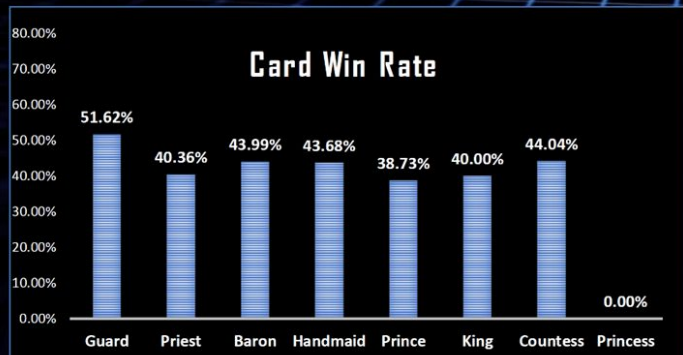
... and Play-testing

Example 1: Love Letter

Speeding up playtesting with Artificial Intelligence




1000 Love Letter games played automatically by AIs in 2 minutes







Example 2: Terraforming Mars

Speeding up playtesting with Artificial Intelligence



1000 Terraforming Mars games played automatically by AIs in 10 hours



Example: Sirius Smugglers (UKGE'23 Demo)

In development, working with designer.

How long do games last, and why do they end?

It is just about possible for a 4-player game to finish in **Round 4**; but on average 3/4-player games last for **8 rounds**, and 2-player games for **10 rounds**. Games end either because of **deck exhaustion**; or because of the **corruption track**. The **Ammonia** and **Contraband** victory conditions trigger very rarely (and even more so once you move beyond 2 players).



	Deck	Ammonia	Contraband	Corruption
2 Players	58	1	7	43
3 Players	44	0	0	64
4 Players	33	0	0	70


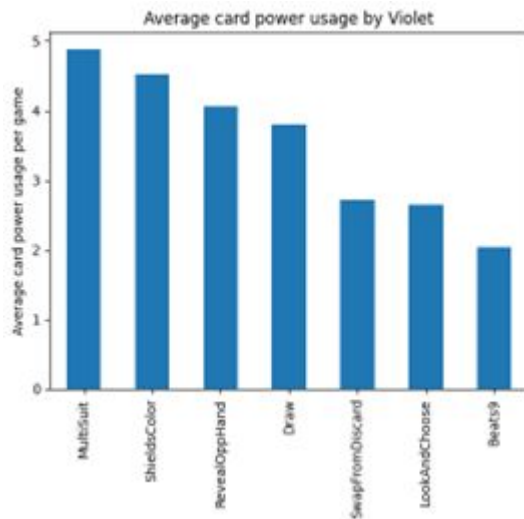
Example: Violet and the Grumpy Nisse (UKGE'23 Demo)

Preparation for second edition, working with designer.

Which card power gets used the most by Violet?

Violet players prefer the **Multi-Suit** power most, which has a **good chance of winning the hand**, but also of **not losing by a large margin**, as it sits right in the middle of card values at 5.



The opposite is true for **Beats-9**, which, although able to **beat the highest card in the game**, it will **lose** to all others with its own value of 1.



Example: Puerto Rico (UKGE'23 Demo)

Is there a clearly optimal choice of actions in the first turn?

Builder first! Or maybe you want some doubloons from the **Prospector**. Then, call on the **Mayor** to bring colonists to your buildings and plantations, and produce with the **Craftsman**. Ship your goods off with the **Captain** to end the round.



Example: Puerto Rico (UKGE'23 Demo)

Is there a difference in expected outcome if you sit next to a bad player?

Clear advantage in sitting to the left of a bad player (taking the next turn)!



Tabletop R&D

DIGITAL TWINS

fast implementation of tabletop games

AUTOMATIC PLAY

thousands of games in seconds

PLAY-STYLES

varied customised strategies

INSIGHTS

in-depth statistical analysis

BALANCE

test game designs

QMUL Spin-out

<https://www.tabletoprnd.co.uk/>



Our Team

<https://www.tabletoprnd.co.uk/>



Raluca D. Gaina

Founder, Director, Technical Manager




Diego Perez Liebana

Founder, Director, Project Manager



James Goodman

Research Associate



Martin Balla

Research Associate



Tabletop R&D

Q&A + Discussion

Tutorial archive: <https://www.tabletoprpd.co.uk/cog23>

Discussion Prompts

- What are you interested to see more of in TAG / how would you use it? AI, games, PCG
- Games industry sees AI as evil. How do we change this perspective?
- Try it out! <https://www.tabletoprnd.co.uk/cog23>
- Join Discord!

<https://discord.gg/qmTTnyMcFW>

