3.6.3 Multi-Player Time Travel

Time-travel in the single player case closely resembles saving the game state and reloading past checkpoints later on. In contrast, in the multiplayer case things get a lot more interesting. When a number of players co-exist in the same environment it is unclear how time travel of one player should change the current time step of other players. A naive approach is to simply use the single-player option, whereby all players are "dragged through time" by the time travel decisions of any player. However, this will likely make for a confusing playing experience and also break game dynamics since any player might be incentivised to travel in time when things are not working well for them.

Instead, we suggest a new approach for multi-player TT which relies on *branching timelines*: Any player can travel back in time *independently* while all other players have the option to continue playing on their current timeline or TT independently. This leaves a crucial question: What are the characters of other players doing on branches that the player is not currently playing? We suggest to use "zombie-actors", i.e. machine learning models that predict the actions of a player in the *alternate* reality given their realised actions in the played reality. This problem is similar to the issues caused by time-delay in multi-player games, which are solved e.g. with *Rollback* which is now being improved using machine learning [1].

To reduce computational overhead and avoid pure "zombie-games", branches that have been abandoned by all players get frozen in time until a player rejoins said branch.

References

1 Anton Ehlert. *Improving input prediction in online fighting games*. 2021

3.7 Artificial Intelligence for Audiences

Antonios Liapis (University of Malta – Msida, MT), Maren Awiszus (Leibniz Universität Hannover, DE), Alex J. Champandard (creative.ai – Wien, AT), Michael Cook (Queen Mary University of London, GB), Alena Denisova (University of York, GB), Alexander Dockhorn (Leibniz Universität Hannover, DE), Tommy Thompson (AI and Games – London, GB), and Jichen Zhu (IT University of Copenhagen, DK)

Artificial Intelligence (AI) has been leveraged for assisting individual players [20, 12] and individual designers or creators [9], but the rise of for-profit content creation platforms [3], and games as a spectacle [1] opens a new and exciting opportunity for AI support. In this working group, we explore applications, algorithms, and interfaces for *AI for audiences*.

The simplest inception of an AI application in this vein would be as mediator between a content creator (e.g. a YouTuber or a Twitch streamer) and the consumers that may be enjoying this content in real-time (e.g. during a stream) or asynchronously (e.g. watching a YouTube video). Focusing on the communication between audience and content, the working group identified the following non-exhaustive list for possible AI roles:

- **AI as mediator.** For instance, the AI may inform a viewer when the content changes (e.g. a new game area is entered or the creator changes the discussion topic), or inform a live-streamer when audience engagement shifts (in tone, volume, or discussion topic).

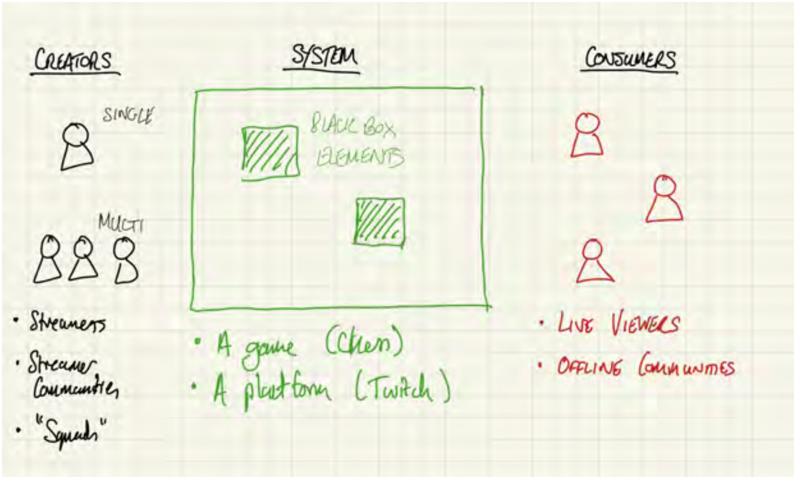

Figure 6 Envisioned AI as mediator between an audience and one or many content creators.

- **AI as entertainer.** For instance, the AI can add a (textual) commentary to a playthrough in real-time. In this role, the AI may act as an *unreliable narrator*, in which case the state of the game need not be described reliably in order to increase engagement through uncertainty and curiosity. Similar patterns are observed in e.g. e-sport competitive matches, where (human) casters give more "optimistic" predictions for a comeback of the currently losing team.
- **AI for hype.** For instance, the AI can algorithmically generate audio, visual, or text assets to promote content scheduled in the future by connecting it with past content from the same creator or a broader context. Similarly, the AI can promote existing content to the audience based on more in-depth patterns (e.g. gameplay progression) and player/viewer models than current recommender systems.
- **AI as tutor.** For instance, when requested by a viewer an AI could explain game mechanics and their interactions as relevant to the current context. The issue of personalisation is pertinent here, as modeling the viewer's expertise (based on the number of similar content they have viewed or games they have played, as well as questions they have asked the AI) could impact the level of explanation and possible examples or anchor points to scaffold the explanation.
- **AI as filter of needless data.** For instance, an on-demand AI can jump to the highlights in the video, or an always-on AI can remove uninteresting or toxic chat between audience members.

The issue of synchronous versus asynchronous engagement can heavily impact the affordances and constraints for both the AI algorithms and the user interfaces. Beyond the obvious fast-response and low-latency requirements, the issue is pertinent because synchronous viewing may foster shorter but more direct interactions between content creator and audience and between members of the audience (e.g. chat). Synchronous viewing opens additional opportunities for AI assistance, such as a personalized recap of the stream so far in case a viewer joins late, or a recap of events while the user was away in case they leave and rejoin. On the other hand, asynchronous viewing allows for more thoughtful discussions to emerge in comments; at the same time interaction with the content is more granular and controlled as viewers can choose which parts of the video to view, rewind, etc.

(a) Lichess turn-by-turn replays with predicted wins and suggested moves.

(b) DotA 2 real-time match progression with gold, experience, deaths, and predicted win chances.

(c) YouTube viewership analytics, including the a “most replayed” label for popular video sections.

Figure 7 Current examples of visualizations, analytics, and predictions intended for audiences.

Note, that the data format of the content that is made available to the AI should ideally not be simply the end-product (e.g. a video) but additional meta-data regarding game actions, context, and potentially even game-specific AI game players. An example of such rich data is provided in *lichess*[1] where viewers (or players after the game is completed) can watch replays of chess matches along with AI-based predictions of win versus loss after every move, as well as suggested moves instead of the one played. Beyond chess, having access to such granular game data could allow for highlight detection (e.g. at points where the predictions shift dramatically between players), summarization (e.g. grouping similar moves together and focusing on highlights), or tutoring (e.g. showing the causal links between early choices and later outcomes). To maximize the potential of such an approach, however, the game developers would need to provide not only game state and action events but also ideally some game-specific AI that could provide nuanced context-specific metrics such as predicted win probability or chosen next moves. Such meta-data and AI-predicted game metrics are already made available for certain games that embrace the game as spectacle philosophy, especially e-sports such as *Dota 2* (Valve, 2013).

However, AI for audiences need not rely on the assumption of a *one-to-many* interaction, or the implicit assumption that the audience consists of passive consumers with no agency over the content or how they interact with it. AI for audiences can be used to promote and support *augmented communities*, where some or all of the audience members can take more proactive roles (indicatively, live commentators with AI visualization assistance or cinematographers by creating custom camera positions in live or replay game data). Audience interactions

[1] `https://lichess.org/`

with the AI itself can also lead to improved computational models, including player models [26, 19] that can provide personalized tutoring (based on detected expertise level) but also for matchmaking between audience members (especially those with proactive roles). Similarly, the AI can operate on a *many-to-many* assumption and find similar content with similar game-states from other streamers to propose to viewers, but also for matchmaking between content creators. The simplest form of AI for content creators could suggest scheduling clashes with popular content creators in the same genre (or followed by the same audience) or niche topics that have not been explored by other content creators. A more proactive AI could also act as a matchmaker between content creators, suggesting ideas on how and on what topic this collaboration could be built on. Algorithms and interfaces for this type of AI assistance can have broader ramifications, as similar many-to-many relationships can be found in crowdfunding platforms (e.g. Kickstarter), virtual crowd working platforms (e.g. Fiverr or creative.ai), and service providers more broadly (e.g. Uber, Wolt).

Several existing algorithmic advancements can be leveraged towards the goals laid out above, including recommender systems [19, 12], text summarisation [13, 21], personalisation [10] and personas [6], highlight detection [12], video indexing and matching [22], viewership analytics [7], coordination and scheduling [2], monetisation and churn prediction [8], expressive range analysis [16] and quality-diversity search [4], AI directors [11, 17], and more. However, novel AI research will be warranted in this vein tailored to the format (video, speech, and game meta-data) and user requirements of such applications. Example directions for AI research include question-answering systems (including natural language processing), text summarisation of real-time expanding datasets (of comments or gameplay), context-aware detection of video segments (e.g. based on text mentions in the comments), or causal models [14] based on audio, visual, video, gameplay, and comment/chat data.

References

1 Dave Boling. How The International became a global "Super Bowl for nerds". `https://www.espn.com/esports/story/_/id/20343989/super-bowl-nerds-dota-2-fans-globe-lured-spectacle-camaraderie-international-7`, 2017. Accessed 27 July, 2022.

2 Elisabeth Crawford and Manuela M. Veloso. Learning to select negotiation strategies in multi-agent meeting scheduling. In *Proceedings of the Portuguese Conference on Artificial Intelligence*, 2005.

3 Cecilia D'Anastasio. Amazon's Twitch seeks to revamp creator pay with focus on profit. `https://www.bloomberg.com/news/articles/2022-04-27/amazon-s-twitch-seeks-to-revamp-creator-pay-with-focus-on-profit`, 2022. Accessed 27 July, 2022.

4 Daniele Gravina, Ahmed Khalifa, Antonios Liapis, Julian Togelius, and Georgios N. Yannakakis. Procedural content generation through quality-diversity. In *Proceedings of the IEEE Conference on Games*, 2019.

5 Fabian Hadiji, Rafet Sifa, Anders Drachen, Christian Thurau, Kristian Kersting, and Christian Bauckhage. Predicting player churn in the wild. In *Proceedings of the IEEE Conference on Computational Intelligence in Games*, 2014.

6 Christoffer Holmgård, Michael Cerny Green, Antonios Liapis, and Julian Togelius. Automated playtesting with procedural personas through MCTS with evolved heuristics. *IEEE Transactions on Games*, 11(4):352–362, 2019.

7 Andrew Hutchinson. Youtube rolls out activity graph to all videos, ups the maximum price of channel memberships. `https://www.socialmediatoday.com/news/youtube-rolls-out-activity-graph-to-all-videos-ups-the-maximum-price-of-ch/624036/`, 2022. Accessed 27 July, 2022.

8 Erik Johnson. A deep dive into Steam's Discovery Queue 2. `https://www.gamedeveloper.com/business/a-deep-dive-into-steam-s-discovery-queue`, 2019. Accessed 6 July, 2022.

9 Antonios Liapis, Gillian Smith, and Noor Shaker. Mixed-initiative content creation. In Noor Shaker, Julian Togelius, and Mark J. Nelson, editors, *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*, pages 195–214. Springer, 2016.
10 Santiago Ontanon and Jichen Zhu. The personalization paradox: The conflict between accurate user models and personalized adaptive systems. In *Companion Proceedings of the International Conference on Intelligent User Interfaces*, page 64–66, 2021.
11 Mark O. Riedl, H. Chad Lane, Randall Hill, and William Swartout. Automated story direction and intelligent tutoring: Towards a unifying architecture. In *Proceedings of the AIED Workshop on Narrative Learning Environments*, 2005.
12 Charlie Ringer and Mihalis A. Nicolaou. Deep unsupervised multi-view detection of video game stream highlights. In *Proceedings of the International Conference on the Foundations of Digital Games*, 2018.
13 Alexander M. Rush, Sumit Chopra, and Jason Weston. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2015.
14 Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
15 Adam M. Smith, Chris Lewis, Kenneth Hullet, Gillian Smith, and Anne Sullivan. An inclusive taxonomy of player modeling. Technical Report UCSC-SOE-11-13, 2011, University California Santa Cruz, 2011.
16 Gillian Smith and Jim Whitehead. Analyzing the expressive range of a level generator. In *Proceedings of the FDG workshop on Procedural Content Generation in Games*, 2010.
17 Tommy Thompson. In the directors chair: The AI of Left 4 Dead. `https://medium.com/@t2thompson/in-the-directors-chair-the-ai-of-left-4-dead-78f0d4fbf86a`, 2014. Accessed 27 July, 2022.
18 Tommy Thompson. How Forza's Drivatar actually works. `https://www.gamedeveloper.com/design/how-forza-s-drivatar-actually-works`, 2021. Accessed 6 July, 2022.
19 Hao Wang, Naiyan Wang, and Dit-Yan Yeung. Collaborative deep learning for recommender systems. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015.
20 Georgios N. Yannakakis, Pieter Spronck, Daniele Loiacono, and Elisabeth André. Player modeling. In Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius, editors, *Artificial and Computational Intelligence in Games*, volume 6 of *Dagstuhl Follow-Ups*, pages 45–59. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.
21 Kevin Yauris and Masayu Leylia Khodra. Aspect-based summarization for game review using double propagation. In *Proceedings of the International Conference on Advanced Informatics, Concepts, Theory, and Applications*, 2017.
22 Xiaoxuan Zhang, Zeping Zhan, Misha Holtz, and Adam M. Smith. Crawling, indexing, and retrieving moments in videogames. In *Proceedings of the Foundations of Digital Games Conference*, 2018.