

Personality Exceptions

Elinor Rubin-McGregor

Department of Computer Science University of Kentucky
Davis Marksbury Building, 329 Rose Street
Lexington, KY 40506-0633 USA
erru227@uky.edu

Abstract

Adding character depth in narrative planners is a complex subject with many approaches. One common strategy used to add depth to characters is to imbue them with personality traits. One attribute that few personality models consider, however, is how character behavior should change when the character is exposed to reminders of powerful psychological influences. Through creating and refining a new digital construct known as *Personality Exceptions*, I will provide a tool that makes it much easier for personality systems to support such conditions. I intend to make a modular tool that researchers can attach to existing personality models to enhance the character depth of their characters and, thus, improve the overall quality of computational narratives.

Introduction

One way to increase the quality of stories generated by narrative planners is to imbue characters with a sense of personality. This helps create believable characters and can create engaging narratives. There are many existing personality models for narrative planners, and these focus on a multitude of different aspects of grasping personality. One aspect that is rarely explored is how characters' personality and behavior can change under emotionally intense situations, as many personality models either give characters a static personality (Bahamón and Young 2013), change a character's entire personality to change behavior (Evans and Short 2014), or else require large amounts of authorial-provided information and considerable amounts of work to accomplish this task such as social facts databases (McCoy 2014) and collections of information on entire (possibly fictional) societies (Kreminski et al. 2020). For my dissertation, I plan to develop what I call *Personality Exceptions* that can support explainable breaks in personality or behavior in order to further enhance the character depth of characters portrayed in narrative planners. The final goal is to make it possible for the planner to generate Personality Exceptions without any authorial input required, thus lessening the burden on the author.

Personality exceptions, at their core, hold two elements: one is a contextual-based trigger to determine when the exception should be active, and the other is a specific change

or change formula to add to the character's normal personality model. The idea for a contextual-based trigger comes from the Indexter (Cardona-Rivera et al. 2012) five-index values of Time, Space, Protagonist, Causal, and Intention indexes to measure narrative memory salience in audiences, with the values adapted to simulate internal memory salience in characters. I have created and tested a set of six indexes for use with Personality Exceptions (Rubin-McGregor, Harrison, and Siler 2023), with the Causal index being split into two to allow for more diversity in approach, one being action type for a direct cause and the other being precondition events for indirect causes.

However, it is my hope to make Personality Exceptions fully modular by testing them with multiple unique personality systems to see how well they function, and if any patterns can be identified in their usefulness. Likewise, I intend to progress this work further by enabling the creation of Personality Exceptions during the generation of a story, thus allowing events the audience witnesses to have a memorable impact on the behavior of the characters when it is relevant.

Work so Far

My work has largely been focused around the narrative planning model of Sabre (Ware and Siler 2021), an expansion of the earlier Glaive (Ware and Young 2014). The Glaive narrative planner produces stories that meet the author's intended narrative goals while honoring that characters would not act against their individual personal goals. This ensures that the narratives generated use believable characters with clear motives understandable to the audience. Sabre allows characters to have and act on their own (possibly inaccurate) beliefs about the world (Shirvani, Ware, and Farrell 2017; Shirvani, Farrell, and Ware 2018). Following the precedent established in Glaive, Sabre distinguished between the concepts of *author goals* and *character goals*, collections of propositions that the author or character "want" to be true. The planner will attempt to generate a story that meets all authorial goals, while only permitting characters to be *consenting characters* for actions that help make their character goals true (Ware and Young 2014). Characters are allowed make multi-step plans known as *character plans* at each timestep to progress their goals (Shirvani, Farrell, and Ware 2018).

To support a flexible personality model in response to

external stimuli, I created **Personality Exceptions**, a combined index holding a narrative-internal situational **trigger**, and an associated personality modification **effect** to be applied when a given character is exposed to the aforementioned trigger. My goal was to enable character depth in narrative generators by creating internal situation-based exceptions to character personality that activate based on an event in the present that is related to some highly salient past event that the character has experienced. My approach took Indexer's (Cardona-Rivera et al. 2012) concept of narrative salience, which was originally proposed to measure memorability of events to a reader, and used it instead to measure memorability of events to a character. I then simulated the influence of powerful pre-existing character memories from the character's backstory by assigning a shift to the character's personality that only applies when they "recall" these memories, as indicated by encountering an event that scores highly along one of Indexer's indices. This allowed for a slew of complex behaviors that are seen often in media: A cruel and malicious antagonist has one friend they are kind towards. Similarly, the pragmatic Batman refuses to use guns because that weapon killed his parents. For my experiments I applied a modified variant of Shirvani's OCEAN based personality model (Shirvani and Ware 2019) to the Sabre belief-intentionality intersecting state-space narrative planner. The OCEAN personality model being a commonly accepted set of five scales to measure personality in terms of Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism. My set of experiments were focused on proving that the exceptions could work to improve the range and quality of personality-implementing planners. As a proof of concept, I sought to provide examples of narratives where character depth is not supported by static personality models, but is supported when Personality Exceptions are added in.

Hypotheses to Test

There are several aspects to consider before Personality Exceptions can be considered truly useful in narrative generation. I have identified several specific goals I would need to accomplish in order to complete my work.

Other Personality Models

Although Sabre and OCEAN worked very well for my tests, I attempted to build Personality Exceptions with as few specific limitations as possible so that Personality Exceptions could be implemented upon other narrative planners and to multiple personality models. I would like to attempt to adapt Personality Exceptions to some other personality models in order to see what, if any, features of them must be adapted to fit other models, and if so, what rules I can identify to properly re-define each type of Personality Exception I designed. Some models I would like to work with include:

- A more recent OCEAN-based personality model with emotions (Shirvani 2021)
- Versu (Evans and Short 2014)
- Comme il Faut (McCoy 2014)
- CB-POCL (Bahamón and Young 2013)

- Why are we Like This (Kreminski et al. 2020)

It is my hope that I can adapt the algorithms I created for Personality Exceptions directly onto multiple models, and in the event of unexpected problems I hope to identify causes of these issues that can help us to either create adaptations of Personality Exceptions or provide strict definitions of when Personality Exceptions could be applicable.

Generating Personality Exceptions

When I created Personality Exceptions, I had to build story domains where personality exceptions already existed. It was only "backstory" events that could cause lingering influences upon characters. For a truly useful narrative tool, Personality Exceptions should be able to be generated during the story in response to events that players observe directly in runtime. To this end, I have identified two different tasks for this goal: creating an algorithm set to determine when a personality exception should be generated, and creating an algorithm set to determine which triggers and what personality impact a new Personality Exception should be given upon generation.

When to Generate As a basic precaution, we should not generate personality exceptions if and when the base personality system changes an character's personality. If an event is impactful enough to change the character's personality permanently, then its influence should most likely not include a Personality Exception as well.

I have several proposals as to how to decide when to generate new personality exceptions. First is a simple threshold value, wherein a combination of goal stressors and, for systems with emotion, emotional impact will trigger the creation of a personality exception. In other words, if the situation is intense enough, the character will remember it. The second idea is based on the statements made by Jaffe, Segal, and Dumke about how emotional and psychological trauma is formed (Jaffe et al. 2005). Specifically, they list three key points behind what makes a specific event traumatic:

1. It was unexpected
2. The person was unprepared
3. There was nothing the person could do to stop the event from happening (Jaffe et al. 2005)

This works well for trauma-specific exceptions; however, I want to expand my reach into allowing strong positive memories to generate exceptions as well. Thus, I believe I might be able to modify these points as such:

1. It was unexpected
2. The person was unprepared
3. The person was made to feel either that they had no control, or that they had complete control (this can be in the context of the person feeling that if they could control things they would do things exactly the same way)

If I can evaluate these conditions as metrics, then in scenarios where all three come into play I can initiate the creation of a personality exception. By creating a set of three metrics, I hope I can analyze when an event generated by

the narrative planner is impactful enough to cause a Personality Exception to be generated. I have already come up with ideas for a few metrics I intend to refine and test in order to measure these points during story generation:

1. The probability the character gave for a given event to occur in the timestep before it occurs
2. A combined measurement of two values: how many actions the character can possibly take before and after a given timestep, and how much their potential actions would effect the world state
3. In a goal-based personality model, I would use a metric of how much a character's goals are either challenged or assisted by a given event, while in an emotion-based personality model I would use a metric of the strength and desirability of the emotions caused by a given event

While I want to perform some set of human tests to ensure these metrics work out, I will admit I am not certain how best to test that these elements are working correctly.

How to Generate As for how to create a new Personality Exception, I intend to do a fair amount of preliminary research before beginning to explore this field. My current plan for the next step is to “generate” pre-made potential triggers (having triggers generated by the planner is the step after this one) from every type based on the moment that caused the exception to generate and compare them with three evaluation metrics:

1. If they relate to the character's affected goals
2. How they relate to the other triggers
3. How closely they relate to the cause or causes that specifically sparked the exception's creation

Another idea is to attempt to apply the emotional knowledge base EmotiNet (Balahur, Hermida, and Montoyo 2012) by looking at the strongest connections between the central emotion that is tied to the creation of the Trigger Event. I could also attempt to use Indexter as an audience-measuring metric to pick the triggers the audience is most likely to recognize; thus choosing triggers based more on narrative experience than on direct psychology. This is a subject I intend to explore in much greater detail before committing to a plan.

Human Experiments

Of course, I would need to ensure my work produces the results I think it will, and designing working tests is just as important as producing the hypotheses. For this, I am anticipating running around three sets of tests.

Multiple Personality Models For multiple personality models, I am largely interested in providing proof of concept again, but there is an advantage to using some existing systems. If possible, I would like to compare the outcomes of stories generated by the original personality models with the outcomes of stories generated by personality models that utilize my Personality Exceptions. By running the same inputs in the same story domains for two systems, one with Personality Exceptions and one without, I could gather stories that human subjects could then be asked to evaluate after being given the relevant information and the backstory.

Alternatively, for personality models that already included complete human interaction aspects to them, a much longer set of tests would involve having human subjects interact with both the original system and one with my personality exceptions included, and ask the subjects to compare their experiences by rating their two sets of interactions based on originality, engagement, and character depth.

Personality Exception Generation–When For this, I could have my program(s) generate several stories and report at which point(s) in the story and for what characters it expected a Personality Exception to generate based on my metrics. I could then show these stories to human subjects and ask them to select a moment in the story which they believe a selected character is likely to find most meaningful.

Personality Exception Generation–How For this, I would be tying all of my results together and attempting to give my subjects a fully interactive narrative that utilizes Personality Exceptions. I would likely want to pick systems and stories to introduce carefully, to increase the chance that my system would produce a Personality Exception for a given character during the narrative, and would need some additional measures taken to ensure that the narrative then provides a chance for the generated Personality Exception to be utilized later in the story. If subjects perceive the characters from these narratives to have consistent personalities displayed, and can identify the causes of actions influenced by a character's Personality Exception, then I will have found success.

Conclusion

Even without a personality model built for them, Personality Exceptions could still be a useful tool when they can be applied to other personality models. In creating stories about personal growth, this tool could help support changes in character behavior with relative ease. By providing precise algorithms, my work could be replicated by anyone who sees a use in portraying this aspect of human behavior.

Modeling personality in narrative planners is likely to be a task still worked on for years to come, and it is certainly one worthy of the attention it gets. Human-like behavior is critical to interesting stories and yet the complexities of “common sense” are so detailed that capturing it in an algorithm has arguably never been achieved yet. But we still reach this goal, all so that we can tell the best stories we can.

References

- Bahamón, J. C.; and Young, R. M. 2013. CB-POCL: A Choice-Based Algorithm for Character Personality in Planning-based Narrative Generation. In *Workshop on Computational Models of Narrative*.
- Balahur, A.; Hermida, J.; and Montoyo, A. 2012. Building and Exploiting EmotiNet, a Knowledge Base for Emotion Detection Based on the Appraisal Theory Model. *Affective Computing, IEEE Transactions on*, 3: 1 – 1.
- Cardona-Rivera, R. E.; Cassell, B. A.; Ware, S. G.; and Young, R. M. 2012. Indexter: a computational model of

the Event-Indexing Situation Model for characterizing narratives. In *Workshop on Computational Models of Narrative*, 34–43.

Evans, R.; and Short, E. 2014. Versu—A Simulationist Storytelling System. *IEEE Transactions on Computational Intelligence and AI in Games*, 6(2): 113–130.

Jaffe, J.; Segal, J.; Dumke, L. F.; et al. 2005. Emotional and psychological trauma: Causes, symptoms, effects, and treatment. *Mental emotional trauma*.

Kreminski, M.; Dickinson, M.; Mateas, M.; and Wardrip-Fruin, N. 2020. Why Are We Like This?: Exploring Writing Mechanics for an AI-Augmented Storytelling Game. In *International Conference on the Foundations of Digital Games (FDG '20)*.

McCoy, e. a., Joshua. 2014. Social story worlds with Comme il Faut. *IEEE Transactions on Computational intelligence and AI in Games 6.2*, 6(2): 97–112.

Rubin-McGregor, E.; Harrison, B.; and Siler, C. 2023. Enhancing Character Depth through Personality Exceptions for Narrative Planners. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 19(1): 136–144.

Shirvani, A. 2021. *Personality and Emotion for Virtual Characters in Strong-Story Narrative Planning*. Ph.D. thesis, University of Kentucky.

Shirvani, A.; Farrell, R.; and Ware, S. G. 2018. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 222–228.

Shirvani, A.; and Ware, S. G. 2019. A plan-based personality model for story characters. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 188–194.

Shirvani, A.; Ware, S. G.; and Farrell, R. 2017. A possible worlds model of belief for state-space narrative planning. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 101–107.

Ware, S. G.; and Siler, C. 2021. Sabre: A narrative planner supporting intention and deep theory of mind. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 17, 99–106.

Ware, S. G.; and Young, R. M. 2014. Glaive: a state-space narrative planner supporting intentionality and conflict. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 80–86.