Proceedings of the Fifteenth AAAI Conference on Artificial
Intelligence and Interactive Digital Entertainment (AIIDE-19)

A Plan-Based Personality Model for Story Characters

Alireza Shirvani, Stephen G. Ware
Narrative Intelligence Lab, Department of Computer Science
University of Kentucky, Lexington, KY 40506
{ashirvani, sgware}@cs.uky.edu

Abstract

Believable virtual characters enhance the experience of interactive virtual worlds. BDI-based state-space narrative planners make their characters more compelling by granting them specific goals to pursue and different (possibly wrong) beliefs about the state of the world. However, such systems do not account for character individuality. We extend work on narrative planning by proposing a personality model based on the Big Five. This paper discusses how each factor of the Big Five is represented by narrative planning features. We also present a study to evaluate whether human readers can perceive the modeled personality traits through simple action patterns in short stories.

Introduction

Interactive narratives are prevalent in virtual environments for entertainment, education, and training. While extensive research has been done to make their behavior realistic, many virtual worlds will also benefit from characters who have noticeable individual personalities. Since hand-crafting individual characters in a coherent, consistent way is a time-consuming task, many procedural narrative systems can benefit from a formal, generative model of personality. In this paper, we propose a BDI-based narrative planner with a domain independent personality model based on the Big Five (DeYoung, Quilty, and Peterson 2007).

The rest of this paper is organized as follows. In section 2, we present the research on the use of personality models in virtual characters. Section 3 discusses selected narrative planning features to represent the Big Five. The proposed model is evaluated in section 4 and finally, section 5 presents conclusions and future works.

Related Works

BDI Model

Modeling individual beliefs, desires, and intentions is a popular approach to believable behavior . Assuming desires are specified by the domain author, agents choose from several planss to achieve them, and the chosen plan becomes their intention.

Planning algorithms have been adapted to ensure that agents only act in ways that they expect to contribute to achieving their goals (Riedl and Young 2010; Teutenberg and Porteous 2015; Shirvani, Ware, and Farrell 2017; Samuel et al. 2018; Ware et al. 2019). They rarely distinguish between different plans to achieve the same goal, even when they have different effects on other characters (Bahamón, Barot, and Young 2015). In this paper, we focus on the use of personality models to address this issue.

Personality

"A personality trait is an enduring personal characteristic that reveals itself in a particular pattern of behavior in different situations" (Poznanski and Thagard 2005). Personality traits can be categorized into hierarches; broad categories (e.g. extrovert) at the highest layer and specific statements about an agent's behavior (e.g. "Likes to talk to others") at the lowest.

The Big Five is a widely-studied taxonomic personality model derived from a factor analysis of a large number of self and peer reports on personality-relevant adjectives (DeYoung, Quilty, and Peterson 2007). Much of the research on The Big Five defines a two-level hierarchy, with the five broad categories, called *factors,* at the top subsuming more specific traits, called *facets*. The five factors are openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism.

There is a debate on the definition of the second layer. For instance, Costa and MacCrae (1992) produced 6 facets for each of the five factors (Costa and MacCrae 1992). Moreover, The Abridged Big Five Dimensional Circumplex (AB5C) defines each facet as a blend of two of the big five factor (Hofstee, De Raad, and Goldberg 1992).

The Big Five Aspect Scales (BFAS), considers all facets defined by(Costa and MacCrae 1992) and (Hofstee, De Raad, and Goldberg 1992) and defines 10 *aspects* as an intermediate layer between factors and facets. We draw from BFAS to map the Big Five to agent behaviors.

Agents with Personality

Our work focuses specifically on communicating personality via an agent's actions in a plan-based multi-agent environment, but numerous researchers have incorporated personality research into virtual agents.

Gebhard and Kipp (2006) define a layered model of affect with emotions as short-term, mood as mid-term, and personality as long-term affect, where personality determines the initial mood before computing mood changes (Gebhard and Kipp 2006).

André et al. (1999) propose an affective agent-user interface based on the Big Five, but rely on the same high-level behavior to convey personality traits. The authors also reduce the dimensions into extraversion and agreeableness to focus on social interactions and neuroticism to control the influence of character emotions (André et al. 1999).

Egges, Kshirsagar, and Magnenat-Thalmann (2003) uses the Big Five to simulate conversational virtual humans with personality, emotions, and mood (Egges, Kshirsagar, and Magnenat-Thalmann 2003). They assume the goals, standards, and attitudes of the agent are fully defined in the domain.

SPOT trains a neural network on a set of rules based on known behavioral predispositions of the Big Five and common sense about human behaviors (Poznanski and Thagard 2005). For instance, crying is the response of a highly neurotic personality in a stressful situation and a bad mood.

Our proposed personality model best compares to that of Bahamón and Young (2013). They define a mapping between personality traits and actions as a domain independent knowledge base. Actions are automatically matched with different personality traits using a set of predefined rules.

Our work is different from the research above in several ways. First, virtual human research often focuses on a single virtual agent, while ours is applicable to multi-character narratives. Second, the majority of work on virtual humans focus on fine-grained physiological manifestations, such as facial expressions, gestures, and dialog, whereas we focus on communicating personality via higher-granularity actions.

Many models only utilize personality in service of affect, particularly determining the initial mood of each agent (Mehrabian 1996). Those models which do consider the effect of personality on agent behavior often focus only on a specific subset of the five factors.

Finally, most automatic story generation systems that reason about personality map personality traits to specific actions or patterns; e.g. they assume specific actions have been labeled as aggressive or helpful. This makes their models more domain dependent and increases author burden. We seek to characterize actions automatically based on properties already available in narrative planners.

The Big Five

We draw primarily from the Big Five Aspect Scales (BFAS), which define 10 *aspects* (2 per factor) (DeYoung, Quilty, and Peterson 2007). Our goal is to oversimplify rather than overlook. In other words, although we simplify each factor, we consider all 10 facets defined by BFAS, and thus all five factors of the Big Five. This contrasts with other work (e.g. Bahamón and Young, 2017), which chooses to overlook rather than oversimplify.

Openness to Experience - Highly open individuals -with high *scores*- are abstract thinkers and motivated by intrinsic interest. Low scores are observed in conventional and conservative people. The first BFAS aspect of openness to experience is (confusingly) named openness, which focuses on imagination, creativity, and interest in art, music, and nature. The other aspect, aptly named intellect, is associated with ideas, e.g. enjoying philosophical discussions and solving complex problems.

Conscientiousness - High scores are observed in organized and persistent individuals who prioritize order and efficiency. They are motivated by being active and orderly, in contrast to low conscientious people, who are messy, lazy, and care-free. The two aspects of this factor are industriousness and orderliness. Industrious individuals are self-disciplined and diligent, while orderliness indicates attention to tidiness and dutifulness[1].

Extraversion - Extroverts are social, active, optimistic, and motivated by enthusiasm. Introverts are reserved, shy, and quiet. The aspects of extraversion are enthusiasm and assertiveness. Enthusiasm indicates a tendency to interact with others and make friends, whereas assertiveness reflects a strong personality, who takes charge and leads the way.

Agreeableness - Individuals with high scores are altruistic, cooperative, and empathetic. Low scores are observed in competitive, distrustful, and uncompromising people. Its aspects are compassion and politeness. Compassionate individuals show sympathy and tend to take interest in others' feelings[2], whereas, individuals lacking politeness are disrespectful and pursue their own goals at the expense of others.

Neuroticism - Neurotic individuals are motivated by their desire to decrease stress and anxiety. High neuroticism (also referred to low emotional stability) indicates strong emotional reactions to external stimuli and sensitivity to threat and punishment. Low scores represent less observable reactions and mood changes. BFAS distinguishes between two aspects of neuroticism, withdrawal and volatility. Withdrawal refers to the inward expression of negative emotions like depression and anxiety. In contrast, volatility is the outward expressions of negative emotions in forms of anger and panic[3].

A Personality Model Inspired by The Big Five

We extend Shirvani, Farrell, and Ware's (2018) BDI-based narrative planner to include a personality model inspired by the Big Five but using only features already measurable in that planner.

[1]While individuals with high intellect are generally better at making plans, highly conscientious ones are best at implementing those plans.

[2]While enthusiasm rewards social affiliation, compassion reflects affiliation driven by empathy.

[3]Withdrawl can be compared to the behavioral inhibition system, and volatility to the fight or flight instinct (DeYoung, Quilty, and Peterson 2007).

The Story Domain

We use the following domain in our evaluation and throughout this paper to provide examples.

Tom is very sick with the flu. The nearby town has an alchemist, who has a recipe for healing potions. The alchemist makes a bottle from a batch of medicinal herbs and sells them for a silver coin, but she is out of healing potions and medicinal herbs! Tom has a silver coin, and there is also a single batch of herbs in the cottage. The alchemist is sound asleep and the town guard patrols the streets.

Potential Narrative Planning Features

In order to minimize authorial burden, our proposed personality model must be independent of the story domain provided by the author. For instance, we cannot represent that conscientious and open characters are concerned with conventionality, since we would need the author to label actions with their level of conventionality. We have intentionally chosen to limit ourselves to structures already provided by narrative planners. Here we briefly describe the features of narrative planning that have the potential to be good markers for different personality traits.

Consenting Characters: An action is defined by a set of preconditions, effects, and consenting characters (Riedl and Young 2010). Consenting characters are the ones responsible for executing the action, e.g. both Tom and the alchemist are consenting characters for buying the medicine.

An action may include characters without their consent, e.g. Tom poisoning the alchemist's drink. If a character is not consenting to an action, but included in its effects, it almost always means that something is being done to that character against their will, since the action does not benefit them in pursuit of their goals[4].

Agent Plans: An action is *explained for an agent* if it is part of a sequence of actions which the agent believes can be executed to achieve its goal. A sequence of actions is *explained* (in general) when all actions are explained for all of their consenting characters (see Shirvani, Farrell, and Ware 2018 for full details). A narrative planner can reason about a sequence of explained actions even when they are not part of the actual story that gets told. These sequences define possible worlds that would have made sense if they had happened (even if they didn't).

For the sake of this paper, we will define an *agent plan* as any sequence of explained actions that the agent believes will achieve one of its goals. Our model of personality is a way of choosing between agent plans to find the one which best fits the agent's personality.

Anticipation: An agent does not have to be a consenting character for every action in their plan (Shirvani, Ware, and Farrell 2017). For example, because Tom believes the alchemist wants to make money, he expects that if he brings the herbs to her she will brew a potion and sell it to him.

[4]We constrain this to actions that include those characters in their effects, since we distinguish those actions from ones where the aurhor includes other characters to establish preconditions, e.g. if Tom commits a crime and the alchemist is there to observe it, Tom is now a criminal.

Changing Minds: An agent can switch between plans, representing the agent chaning their mind. Recall that Tom can give the herbs to the alchemist, have her brew the potion, and then buy it. Tom can also break into the alchemist's house, steal the recipie for the medicine, and brew the potion himself. Both are valid agent plans for Tom, because they are sequences of explained actions that will achieve his goal. Therefore, this is also a valid agent plan: Tom breaks into the house, then leaves the house, knocks on the door to wake up the alchemist, then Tom gives the alchemist the herbs, has her brew the potion, and Tom buys it. In this agent plan, Tom changes his mind about how to achieve his goal, but it is still composed only of explained actions that can eventually achieve his goal.

Conflict: A conflict exists for an agent's plan when they can imagine a way that it might go wrong. Specifically, an agent plan has a conflict for every sequence of explained actions the agent can anticipate and which would render the plan impossible to execute. This model is inspired by Ware and Young's (2014) plan-based model of narrative conflict.

Mapping the Big Five into Narrative Planning

Given multiple plans that could achieve their goal, an agent should choose the one which best fits its personality. Table 1 presents narrative planning features, the personality factor they represent, and the BFAS aspects they are associated with. We calculate these 12 features for each plan an agent considers. Then, using these 12 features, we calculate a single utility value that expresses how stongly the agent should prefer that plan.

We must note that features marked as (R) are negatively correlated with their corresponding aspect. For instance, feature 3 in Table 1 is "# of actions in a plan (R)." This means highly conscientious agents try to *minimize* the number of actions in their plans, but low conscientious agents *maximize* and choose the longest plan. On the contrary, feature 5 is positively correlated to conscientiousness, so highly conscientious agents maximize the number of actions with themselves as consenting characters (and vice versa).

Openness to Experience We focus on creative thinking to model openness. We define a creative plan to be the one which contains the most original action. Of all the actions that appear in all the plans an agent could use to achieve their goal, the agent considers the plan containing the action which appears least frequently. Specifically, the value of feature 1 is calculated as follows:

$$1 - \min_{i=1...m} \sum_{j=1}^{n} \frac{Occurences(a_i, p_j)}{Length(p_i)}$$

Where, p_j represents a plan in the set of n plans $\{p_1...p_n\}$ the agent is considering and a_i is an action in the plan with m actions. $Occurences(a_i, p_j)$ is the number of times the action type appears in p_j and $Length(p_j)$ is the number of actions in p_j.

Next, individuals with high intellect have the capacity to form better plans and let their hypotheses die in their stead

Table 1: Selected planning features for the big five factors

Facet	Feature
Openness	1.The minimum action likelihood in a plan (R)
Intellect	2.Probability of success of a plan
Industriousness and orderliness	3.# of actions in a plan (R)
	4.# of times the agent changes their mind (R)
	5.# of actions with self as the consenting character
Enthusiasm	6.# of actions including others with their consent
Assertiveness	7.# of actions including others without their consent
Compassion	8.# of actions including others with their consent
	9.# of goals achieved for other characters
Politeness	10.# of actions including others without their consent (R)
	11.# of conflicts created for other characters (R)
Withdrawal and Volatility	12.# of times the agent changes their mind

(Popper and others 1948). Therefore, they try to find a possible world for each character in which their plan causes a conflict with their own.

The probability of success (feature 2) of a plan is defined as follows.

$$1 - \sum_{i=1}^{n}\sum_{j=1}^{m} \frac{Conflict_{a_j}(c_i)}{n \cdot m}$$

Where, c_i is a character in the set of n characters and a_j is an action in the set of m actions. Accordingly, $Conflict_a(c)$ is 1 if c is an NPC, whose plan conflicts with a and otherwise 0.

The lowest action probability in our story domain is buying the recipe from the alchemist. Therefore, a highly open character does so and makes the medicine himself, which is highly unlikely to fail as well. However, with low openness, Tom tries to steal the recipe from her and gets arrested[5].

Conscientiousness Markers of industriousness include persistence and efficiency. Therefore, highly conscientious agents tend to prefer plans that are shorter (feature 3) and once they have chosen a plan, they stick to it and rarely change their mind (feature 4). More specifically, a highly conscientious agent tries to minimize the number of times they change their mind and "finish what they started" (DeYoung, Quilty, and Peterson 2007).

Regarding orderliness, we focus on the need for order and perfection, to make "everything just right" and to have "every detail taken care of" (DeYoung, Quilty, and Peterson 2007). In addition to those markers, the negatively correlated markers of industriousness, being lazy and postponing decisions, all reflect relying on self rather than others (feature 5). Accordingly, highly conscientious agents tend to maximize the number of actions with themselves as one of the consenting characters, whereas low conscientious agents wait around for other characters when they can expect them to take helpful actions.

The shortest plan to achieve the goal with maximum self-reliance is breaking into the shop, stealing the recipe, and making the medicine. On the contrary, if Tom tells the alchemist about the herbs, the plan will mostly rely on the alchemist to go, collect the herbs, come back, and make the medicine. What makes the plan even longer is to poison her drink and steal the herbs and the medicine when she comes back, unknowingly drinks the poison, and faints.

Extraversion Extroverts find social interactions rewarding and thus tend to include as many other characters in their plans as possible.

Consequently, we represent enthusiasm (feature 6) with the number of actions including other consenting character and assertiveness (feature 7) as the number of actions with other characters in their effects without their consent.

Once again the plan that includes other characters the most is for Tom to tell the alchemist about the herbs and then proceed to poison her[6]. In contrast, an introvert would break into the house, steal the recipe, and make the medicine, which results in evading all interactions with other characters.

Agreeableness Compassionate agents have the tendency to help others achieve their goals. Therefore, a high agreeable agent prefers actions that either help others along the way -explained by their consent to the action (feature 8) or directly satisfy their goals as a side effect (feature 9).

On the other hand, with low politeness, agents take advantage of others without their consent, thus, satisfying their own goal but not others (feature 10). Those agents are also very competitive and seek conflict, which means that low agreeable agents will even go out of their way to thwart the plans of other agents (feature 11). This feature is calculated as follows.

$$\sum_{i=1}^{n}\sum_{j=1}^{m} \frac{Conflict_{c_i}(a_j)}{n \cdot m}$$

Where, c_i is a character is set of n characters and a_j is an action in set of m actions. Therefore, $Conflict_c(a)$ is one if a is an action that causes a conflict with c's plan and otherwise zero.

As expected, a highly agreeable character would simply collect, give the herbs to the alchemist, and buys the medicine from her. While, a disagreeable Tom attempts to poison her.

[5] the most probable action with the lowest probability of success

[6] It's worth mentioning that the act of poisoning the alchemist is meant to demonstrate assertiveness, taking charge and taking control. However, a high score in agreeableness would make Tom avoid this plan.

Neuroticism Although neuroticism is one of the factors that mainly represents affective responses, we target its side effects on planning and thought processes. A neurotic agent tends to be indecisive and often change their mind.

With a high neuroticism score, Tom frequently changes his mind after taking certain actions. For instance, he would break into the shop to steal the recipe but change his mind, leave the shop, and proceed to knock on the door[7].

Calculating the Utility Value

Our model begins by finding a set of plans that an agent could use to achieve its goals. This is done via aforementioned planners. We then calculate the utility value of each plan as follows:

1. Compute the value of every feature presented in Table 1.
2. Compute the total feature value of each factor as the weighted sum of its features (the sum of the weights equal to 1). Here, we consider each feature to have the same weight.
3. Multiply each total feature value by its factor score in [-1, 1].[8].
4. Normalize values for all features to [0, 1] to form a 5 by 1 vector.
5. Utility of a plan is the Euclidean norm of that vector.
6. Finally, plans are sorted based on their utility value, which expresses the character's preference for each plan.

Evaluation

In this paper, we investigate three hypotheses to evaluate whether people reading stories generated with our model can percieve the agent personalities we intend. For each of the five factors:

1. When presented a short story where the protagonist displays that factor, subjects answer questions about the protgaonist's personality that demonstrate they percieved the factor as intended.
2. When presented the story above, subjects can recognize other stories where the protagonist demonstrates that same factor.
3. Those subjects who perceive the factor as intended (from Hypothesis 1) recognize other stories that demonstrate that same factor (as in Hypothesis 2).

[7]Of course, a high conscientiousness score decreases the number of times Tom changes his mind. However, as we will describe in the next section, neuroticism has a higher impact on this feature than conscientiousness.

[8]If the score of a factor is 0, the value of its features will be ignored. However, instead of ignoring those features, the agent must exhibit the behaviour of an average character. For instance, an agent with conscientiousness score of 0 should neither prefer short or long plans, but rather a plan with an average length. Therefore, we calculate the average value of each feature over the set of the considered plans. Subsequently, the corresponding value of feature f is computed as $-0.1 \times |Value_f - Average_f|$. The sign of the multiplier ensures prioritizing the closest values to the average and its magnitude gives more weight to non-zero factors.

We want to demonstrate that subjects can percieve both the high and low scores of each factors, so we tested 10 conditions in total: one where Tom has high openness, one where he has low openness, one where he is highly conscientious, etc. In each condition, Tom has either a high value (1) or low value (-1) for one factor and average values (0) for the other four. While we would eventually like to test all 243 combinations of high, low, and average features, this initial work is limited to the 10 described above.

Experimental Design

Subjects first read a brief description of the domain, introducing the people, places, things, and actions. Then they were shown four possible stories[9] that might happen, depending on how Tom decided to act.

The set of four agent plans that Tom chose between were based on his personality (which was different depending on the condition). We included a plan with high utility (good fit for Tom), the plan with the lowest utility (bad fit for Tom), the plan for an average agent (i.e. best fit for an agent with a value of 0 for all factors), and a randomly chosen plan that was not a duplicate of the other three.

After reading these four possible stories, we narrated which one actually happened, which demonstrates Tom's personality through his choice.

Then subjects responded to statements about Tom's personality using a 5 point Likert scale. These statements are adapted from various Big Five personality inventories. Since we only use existing planning features to simulate a simplified version of the Big Five, we selected the markers that best captured our simulated traits. Table 2 presents the questions asked based on the shown personality factor. The table also shows whether we directly use an item from previous research, or change or reverse it to better represent our features. Some items are also changed to include an adequate number of reverse markers for each factor.

Finally, subjects were shown four new stories and asked which one they thought Tom would choose. These four stories were all different from the previous four. They included one that reflected a high utility plan for Tom, one of low utility, the first story generated by the Glaive narrative planner (which makes no attempt to account for personality), and a randomly chosen story that was not a duplicate.

Results

We generated 26 stories and collected results for 228 subjects using Amazon's Mechanical Turk. Each participant viewed one of the 10 conditions. At least 40 subjects evaluated each factor (at least 20 for the high openness, at at least

[9]We showed 4 actual stories, not 4 agent plans. For example, if Tom planned to steal the potion from the alchemist, it is possible that the actual story told invovled Tom getting arrested by the guard and failing to obtain the potion. We did this to ensure readers percieved possible conflicts.

[10]DeYoung, Quilty, and Peterson (2007)

[11]Hofstee, De Raad, and Goldberg (1992)

[12]Goldberg (1992)

[13]Costa and MacCrae (1992)

Table 2: 5-point Likert scale questions

	Question
O	"Finds creative solutions to problems." from "Needs a creative outlet."[10]
	"Tends to analyze possible outcomes of his plans." from "Tends to analyze things."[11]
	"Has difficulty coming up with excellent plans (that rarely fail).(R)" reverse of "Has excellent ideas."[12]
	"His ideas are ordinary and hardly unique.(R)" from "Considers himself an average person."[11]
C	"Gets things done quickly."[10]
	"Wastes his time.(R)"[10]
	"Makes plans and sticks to them."[10]
	"Does just enough work to get by and rather relies on others.(R)" from "Does just enough work to get by."[13]
	"He wants everything to be just right so he prefers to do things himself." from "Wants everything to be just right."[10]
E	"Keeps others at a distance.(R)"[10]
	"Finds it difficult to approach others.(R)"[10]
	"Feels comfortable around people.[13]"
	"Takes charge."[10]
	"Has an assertive personality."[10]
A	"Takes advantage of others.(R)"[10]
	"Avoids conflict." reverse of "Seeks conflict."[10]
	"Is out for his own personal gain.(R)"[10]
	"Likes to do things for others."[10]
	"Can't be bothered with other's needs.(R)"[10]
N	"Is filled with doubts about things."[10]
	"Is NOT easily discouraged.(R)" reverse of "Is easily discouraged."[10]
	"Rarely changes his mood.(R)" reverse of "Changes his mood a lot."[10]
	"Does not know why he does some of the things he does."[11]
	"Does things that he later regrets."[11]

20 for low openness, etc.). All stories were the same for the participants viewing the same condition except for the random stories.

Hypothesis 1 is that subjects answer personality questions about Tom consistent with the personality our model expressed. For each question answered by each subject, we define success as agree or strongly agree if the marker is positively correlated (or disagree or strongly disagree if it was negatively correlated) with that factor when they saw Tom show a high value for that factor. When they saw Tom show a low value of that factor, success was the opposite. We use a binomial exact test to detect if we observed more successes than we should expect to see by chance. The *p*-value and effect size (expressed as relative risk) for each factor are given in Table 3. We rejected the null hypothesis at the $p < 0.05$ level for 3 factors, and at the $p < 0.1$ level for the other two.

Table 3: Experiment Results

	Hypothesis 1		Hypothesis 2	
	p-value	Effect Size	p-value	Effect Size
O	0.072	1.160	**0.026**	1.60
C	**0.016**	1.160	**0.001**	1.73
E	**0.024**	1.167	**0.014**	1.61
A	**0.048**	1.167	$\prec$**0.001**	2.80
N	0.063	1.128	**0.002**	2.04

Hypothesis 2 is that when subjects choose a story for Tom they will choose the one that best expresses Tom's personality according to our model. We define success as choosing the best story out of the four presented. The *p*-value and effect size for a binomial exact test for each factor are given in Table 3. We rejected the null hypothesis at the $p < 0.05$ level for all factors.

Hypothesis 3 is a combination of the other two. We define success as a subject answering questions about Tom as expected *and* choosing the story which best expresses Tom's personality. These results were only significant ($p < 0.05$) for agreeableness (effect size 2.78) and marginally significant ($p < 0.1$) for conscientiousness (effect size 1.64) and extraversion (effect size 1.60). We interpret these results as follows. Although readers can capture a mental picture of a character's personality similar to our model, they cannot perfectly match that picture to both the presented markers (Table 2) and other stories. This indicates that our set of markers can be improved to better represent each factor.

Though many tests were significant, effect sizes were relatively low. We attribute some of this to the high noise collected from Mechanical Turk data.

Conclusions and Future Works

In this paper, we proposed a personality model based on the Big Five to extend BDI-based narrative planners. Our goal was to enhance the immersion of interactive virtual environments by simulating more unique characters.

In order to make our personality model fully domain independent, we limited ourselves only to 12 existing narrative planning features that can represent the 10 aspects of the Big Five. The experiment results of our study indicated that a human audience can identify specific behavioral patterns and also recognize stories with similar traits.

In the future, we plan to use our personality model to evaluate stories that combine different pairs of the Big Five and construct a player model for experience management.

Finally, since most Big Five factors have a strong impact on character emotions, we aim to further build upon our proposed model to include affect and coping mechanisms.

Acknowledgments

This research was supported by NSF awards IIS-1647427.

References

André, E.; Klesen, M.; Gebhard, P.; Allen, S.; and Rist, T. 1999. Integrating models of personality and emotions into lifelike characters. In *International Workshop on Affective Interactions*, 150–165. Springer.

Bahamón, J. C., and Young, R. M. 2013. CB-POCL: a choice-based algorithm for character personality in planning-based narrative generation. In *2013 Workshop on Computational Models of Narrative*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Bahamón, J. C., and Young, R. M. 2017. An empirical evaluation of a generative method for the expression of personality traits through action choice. In *Thirteenth AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment*, 144–150.

Bahamón, J. C.; Barot, C.; and Young, R. M. 2015. A goal-based model of personality for planning-based narrative generation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence, Student Abstract Track*.

Costa, P. T., and MacCrae, R. R. 1992. *Revised NEO personality inventory (NEO PI-R) and NEO five-factor inventory (NEO-FFI): Professional manual*. Psychological Assessment Resources, Incorporated.

DeYoung, C. G.; Quilty, L. C.; and Peterson, J. B. 2007. Between facets and domains: 10 aspects of the big five. *Journal of personality and social psychology* 93(5):880.

Egges, A.; Kshirsagar, S.; and Magnenat-Thalmann, N. 2003. A model for personality and emotion simulation. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, 453–461. Springer.

Gebhard, P., and Kipp, K. H. 2006. Are computer-generated emotions and moods plausible to humans? In *Sixth International Conference on Intelligent Virtual Agents*, 343–356. Springer.

Goldberg, L. R. 1992. The development of markers for the big-five factor structure. *Psychological assessment* 4(1):26.

Hofstee, W. K.; De Raad, B.; and Goldberg, L. R. 1992. Integration of the big five and circumplex approaches to trait structure. *Journal of Personality and Social Psychology* 63(1):146.

Mehrabian, A. 1996. Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology* 14(4):261–292.

Popper, K. R., et al. 1948. Objective knowledge: An evolutionary approach. *Ox-43*.

Poznanski, M., and Thagard, P. 2005. Changing personalities: towards realistic virtual characters. *Journal of Experimental & Theoretical Artificial Intelligence* 17(3):221–241.

Riedl, M. O., and Young, R. M. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research* 39:217–268.

Samuel, B.; Reed, A.; Short, E.; Heck, S.; Robison, B.; Wright, L.; Soule, T.; Treanor, M.; McCoy, J.; Sullivan, A.; et al. 2018. Playable experiences at AIIDE 2018. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Shirvani, A.; Farrell, R.; and Ware, S. G. 2018. Combining intentionality and belief: Revisiting believable character plans. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

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

Teutenberg, J., and Porteous, J. 2015. Incorporating global and local knowledge in intentional narrative planning. In *International Conference on Autonomous Agents and Multiagent Systems*, 1539–1546.

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

Ware, S. G.; Garcia, E.; Shirvani, A.; and Farrell, R. 2019. Multi-agent narrative experience management as story graph pruning. In *Fifteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.