

Using Influence and Persuasion to Shape Player Experiences

David L. Roberts*
School of Interactive Computing
Georgia Institute of Technology

Merrick L. Furst†
School of Computer Science
Georgia Institute of Technology

Charles L. Isbell‡
School of Interactive Computing
Georgia Institute of Technology

Brian Dorn§
School of Interactive Computing
Georgia Institute of Technology

Abstract

In this paper we present a pilot study on the use of *influence* to guide player behavior in an interactive storytelling environment. We discuss the need for authoring support in drama management systems and present *computational models of influence* as a method to meet that need. To investigate some of the benefits of using influence, we implement an authoring tool and web-based mixed-media choose-your-own-adventure-style storytelling system, using it to conduct a pilot study. The pilot study is based on hand-authored static content, but the results indicate promise for generalizing to dynamic content. Further, we provide an analysis of the data that indicates the potential effectiveness of *influence statements* on affecting player decisions and comment on players' perceptions of self-agency with or without those statements.

1 Introduction

Commercial computer games produced today can be the result of tens, if not hundreds, of person years of a coordinated effort among artists, programmers, designers, and authors. An important goal in developing new technologies for authoring such games (and other similar interactive virtual experiences) is to provide a paradigm that preserves expressive power for authors without increasing either their authorial effort or their need for advanced technical expertise.

One approach to realizing this goal is to implement an omniscient coordinator that tracks the player's experience and adapts the environment to bring about a targeted progression of events. Such a *drama or experience manager* is tasked with guiding the player in a dramatic experience prescribed by the author [Laurel 1986; Riedl et al. 2008].¹

In the past two decades a large number of approaches to drama management have been developed (see [Mateas 1999] for a survey of early work and [Roberts and Isbell 2008] for a survey of more recent work). Aside from the environment and story itself, three design problems must be solved to fully implement a drama management system:

1. The system must have a way of representing the state of the narrative and the author's goals for the experience. In addition,

*e-mail: robertsd@cc.gatech.edu

†e-mail: merrick@cc.gatech.edu

‡e-mail: isbell@cc.gatech.edu

§e-mail: dorn@cc.gatech.edu

tion, the system must have a way to reason about the player's behavior in the environment in order to select appropriate narrative goals.

2. The system must have actions that provide it a way to affect the environment. More importantly, the system must be able to reason about how the actions it takes will affect both the player's experience and achieving the author's goal.
3. The system must have a way to ensure consistency of the actions it takes given the current state of the environment.

To date, the bulk of work on drama management has been focused on the first problem. Various approaches have been designed and to varying degrees implemented and tested in simulation or with actual game environments. The approaches have used various representational schemes from bags of beats [Mateas and Stern 2005], to formal decision processes [Nelson et al. 2006], to dynamic decision networks [Mott and Lester 2006]. Similarly, the algorithmic reasoning solutions have run from AI planning [Young et al. 2004], to optimization [Weyhrauch 1997], to advanced statistical techniques [Bhat et al. 2007].

Despite the significant representational and computational power provided by those approaches (and others), the systems have relied heavily on the author to implement solutions to the second and third problems. In this paper, we will discuss our first steps toward automatically generating the solutions to those problems. Our theory is based on designing *computational models of influence* to allow the system to reason about how best to shape the player's experience and automatically create utterances that are both meaningful in the environment and will persuade the player to behave in a particular manner or make desirable decisions.

In the next section, we will provide some background on theories of influence by discussing six principles of influence from social psychology that will form the basis of the models we are developing. The majority of this paper will be devoted to presenting our storytelling architecture, authoring tool, and the design of our study. Beyond that, we will present results from our initial pilot study that lend support to some of the benefits that will result from using theories of influence in the automatic implementation of drama manager actions, focusing especially on one of the six principles of influence. We will conclude the paper with a discussion of future directions for further studies.

2 Influence Theory

Because interactive experiences are marked by a strong social context, we have to move beyond simple physical manipulation of an environment to fully engage in the management of these experiences. Doyle and Hayes-Roth have had success using a "sidekick" that directs the player's attention to different parts of the environment [Doyle and Hayes-Roth 1998]; however, their approach was not automated and relied on the relationship of the player's character to the sidekick, not on a theoretically-grounded approach to

¹For the sake of consistency with the majority of the literature, we will refer to such managers as drama managers (DMs).

manipulating the player’s attention. In contrast, we turn to the theories of *influence* and *persuasion* from social psychology and to the theories of *behavioral economics*. While there are a number of takes on those theories (a complete discussion of which is well beyond the scope of this paper), we have based the work described in this paper on the theories of Cialdini [1998] and Ariely [2008]. Our goals in developing computational models of influence are to: 1) benefit authors by providing tools designed to influence players to buy into the adoption of goals consistent with the author’s; 2) reduce the burden on authors by enabling them to specify goals abstractly, relying on the principles of influence to bridge the gap to a concrete implementation in the virtual environment; and 3) accomplish (1) and (2) without the player perceiving any decrease in (and preferably an increase in) self-agency. In this paper, we will present the results of a pilot study that verify the efficacy of this approach for addressing the first goal. Our work on the second and third goals is ongoing and will be discussed in the conclusion of this paper.

Using computational models of influence, authors can create situations that will influence a player to adopt a long term goal rather than simply accomplish a short term task. Take, for example, the task of trying to convince a player to mow her lawn. In the past, this type of goal would be achieved through the physical manipulation of the environment using a type of “lock-and-key” approach—progress in the experience is halted (the lock) until the player performs the desired task (the key). In this example, a lawnmower might be placed in view of the player and (metaphorical) walls would be constructed to prevent the player from performing other (meaningful) tasks. If the player does not mow her lawn, eventually she will realize there isn’t much else to do. Once the task is completed, the walls are taken down and she can proceed.

Now, suppose it is desirable for the player to mow her lawn on a regular basis, not just once. One option is to repeatedly use the lock and key approach. While that may be successful, it is somewhat limiting and may become repetitive or mundane for the player. By carefully constructing a social interaction in the environment based on a model of influence, the drama manager can persuade the player to adopt the goal of mowing her lawn regularly. In doing so, the need to repeatedly manipulate the environment is eliminated but the player continues to comply with the author’s goal.

It should be noted that the traditional approach of physical manipulation of the environment can place players in situations that may change their mental or emotional state. It is relatively straightforward to arrange for the transfer of knowledge from a non-player character (NPC) to the player; however, simply imparting knowledge to the player is not sufficient to increase the likelihood that she will choose to adopt a specific narrative goal. By contrast, using influence theory, the drama manager will be able to decide how to change the player’s mental or emotional state without using detailed pre-authored content.

The concepts discussed in the following two subsections are excerpted and adapted from Cialdini’s book *Influence: the Psychology of Persuasion* [1998] and are discussed by Roberts *et al.* as well [2008].

2.1 Click Whirr

All species—including humans—have certain built-in mechanical responses to specific stimuli. In animals, these responses take on many forms. For example, a certain large species of fish maintain a symbiotic relationship with another smaller fish known as a Bluestreak Cleaner Wrasse. The Wrasse eats parasites and dead tissue from the underside of the larger fish. The Wrasse will perform a dance in front of the larger fish which will activate its mechanical response and cause it to become perfectly still and wait to be

cleaned. The Wrasse will then approach and clean it, obtaining an easy meal while providing a service to the larger fish.

These responses have been called *click*, *whirr* responses to represent the mechanical click of a recorded tape loading and the whirring of it as it is played. In animals, it is believed that these click whirr responses are instinctual and are free from social context. On the other hand, in humans it is believed that these responses are developed from psychological principles or social stereotypes that we learn over time. In fact, these learned responses in humans are thought to be coping mechanisms. We use them to reduce our cognitive burden when dealing with the ever-increasing complexity of stimuli we are faced with on a daily basis.

In order to use these principles effectively for interactive experiences, we need only to hit upon the trigger features that cause humans to play their recorded tapes. For example, a third species of fish, the Saber-Toothed Blenny, has learned to take advantage of the symbiotic relationship exhibited by the other two species simply by performing the dance to induce passiveness in the larger fish. When the larger fish enters its catatonic state, the Blenny will swim up and take a bite from the larger fish to obtain a free meal and swim away before it can be attacked. The amazing thing about using the principles of influence to the DM’s advantage is that to do so requires minimal effort. As a result, a player willingly complying with the DM’s (and therefore author’s) wishes will tend to see their actions as a result of either their own choices or of natural forces rather than the influence of an exploiter [Cialdini 1998].

2.2 Tools of Influence

In the long run, we will focus on six principles of influence. These principles have been identified by years of research in the field of social psychology and behavioral economics and are frequently employed as sales tactics by savvy marketers:

- **Reciprocation:** give and take; when someone does something for us we feel obligated to return in kind.
- **Consistency:** we have a near-obsessive desire to appear consistent with what we have already done or said.
- **Social Proof:** we look to others, especially those similar to us, to determine the appropriate action to take.
- **Liking:** the more we like someone, the more willing we are to acquiesce to her requests.
- **Authority:** we have a deep sense of duty to authority.
- **Scarcity:** something that, on its own merits, holds little appeal to us will become decidedly more enticing if it will soon become unavailable to us.

These principles provide a foundation for understanding how to create powerful tools to effect behavioral change in players. The models we are implementing are computational realizations of those tools. Used properly, each of these principles by themselves, or in combination with another, can greatly increase the likelihood of someone complying with a request. It is worth noting that these principles can never guarantee compliance. We believe this is an important feature of our approach: while the careful application of these principles of influence can greatly increase the chances of a player choosing to act in a manner the author prescribes, she always has a choice. Thus, the affordance for self-agency is strictly preserved.

Note that each principle can potentially be used in more than one way. For example, scarcity can be used to entice a player to obtain a particular object or to convince her that certain information is




Figure 1: A schematic of the storytelling system. The player begins with the consent information and a brief demographic survey, then iterates through story events consisting of text-video-text-question tuples, and finishes with a brief exit survey.

more important. Liking can actually be employed as a function of friendship, reputation, or physical attractiveness. Authority can be asserted merely with a title like Dr. or can be a function of celebrity.

In addition to social psychology, many of these principles of influence are grounded in the theories of behavioral economics [Ariely 2008]. For example, the principle of social proof is often used in determining the appropriate price for a good. Similarly, consistency is referred to as *setting an anchor* and is discussed at length in the behavioral economics literature. Setting an anchor is something that humans do subconsciously when initially determining the value of something for which they are previously uncommitted. This could be an object like a TV or a time commitment like volunteering at an animal shelter. In a social context, setting an anchor refers to making a determination as to what some agreement is worth. For example, suppose we ask a player, “Is helping animals in need worth two hours per week of your time?” By answering yes to that question, the player has set an anchor. Now, suppose the player has said yes to the first question and we now ask her, “Are you willing to spend an hour per week volunteering at an animal shelter?” Because she has already anchored to two hours per week (even if only in theory), she is far more likely to agree to spend the hour per week at the shelter. The notion of setting an anchor in behavioral economics is a specialization of the commitment principle in social psychology.

3 System Architecture

Our framework for evaluating our approach is a web-based mixed-media choose-your-own-adventure-style interactive storytelling system. Our system displays a sequence of authored text and videos that comprise narrative units, or events, that are linked together by explicit decision points for the player. The videos have been obtained from YouTube², a free online repository for streaming video. A player’s experience will advance according to the procedure depicted in Figure 1 (more details on certain aspects of the system are provided in the next section on study design):

- 1) **Consent information:** Upon arriving at the landing page, the player is presented with the study’s consent information and asked to acknowledge it.
- 2) **Demographics:** The player is asked to answer a set of basic demographic questions. Upon completion, the player is presented with brief instructions for continuing.

²<http://www.youtube.com/>

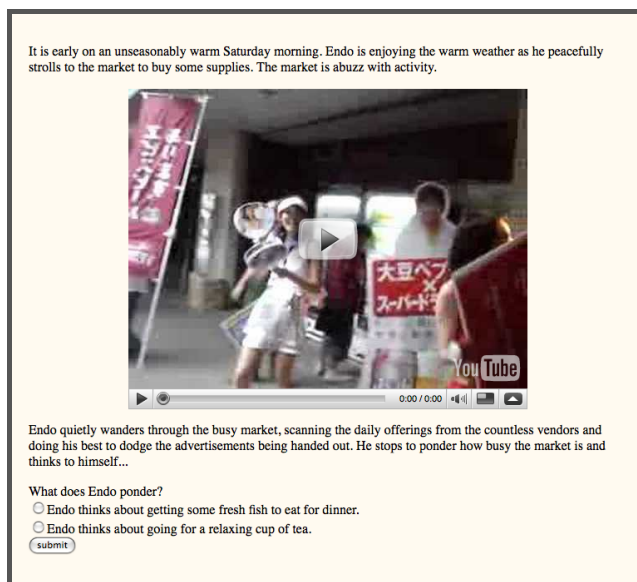


Figure 2: A screenshot of the storytelling system during a story event. The player sees text containing story information, a video, more text with story information, and a question-response set that allows the player to make choices for the main character.

- 3a) **Story:** The player sees author-provided text that describes the beginning of a story event (referred to here as “pre-text”).
- 3b) **Video:** The player is presented with a video from YouTube presenting information supplemental to the story. Sightings, sounds, and action can be conveyed much more efficiently. The video does not start playing automatically; it is completely optional for the player to view it.
- 3c) **Story and selection:** After the video, the player is presented with a short bit of text (referred to here as “post-text”) and a multiple choice question. The question solicits a decision from the player that will drive the narrative. For the purposes of this study, there are always two answers to the questions, but there is nothing in the architecture design that prevents more options. At this point, the system cycles back to Step (3a) and displays another text-video-text set or, once the story has reached a conclusion, moves on to the final step.
- 4) **Exit survey:** The player is presented with a set of five statements and is asked to indicate their level of agreement.

All of the text, question-response sets, and videos are authored ahead of time. We have opted to obtain the videos for our system from YouTube to ease the authoring process and to enable authoring studies in the future. Each video is given a unique identifier that is stored in a mysql database along with the html code to embed the video, any meta-data we associate with it, and information on how it is used in the narrative. Additionally, the database is populated with the question-response groups and the story text we have authored. Figure 2 is a screenshot of our production system in Step (3) during the first story event.

Because directly updating database tables to create story events and transitions based on answers to questions is burdensome, we have also implemented an authoring tool. The tool provides a graphical representation of story events, transitions between them, and a way to update the associated text, video, question, and answer data. The story is depicted graphically as a directed acyclic graph where vertices represent story events and have associated pre-text, video,

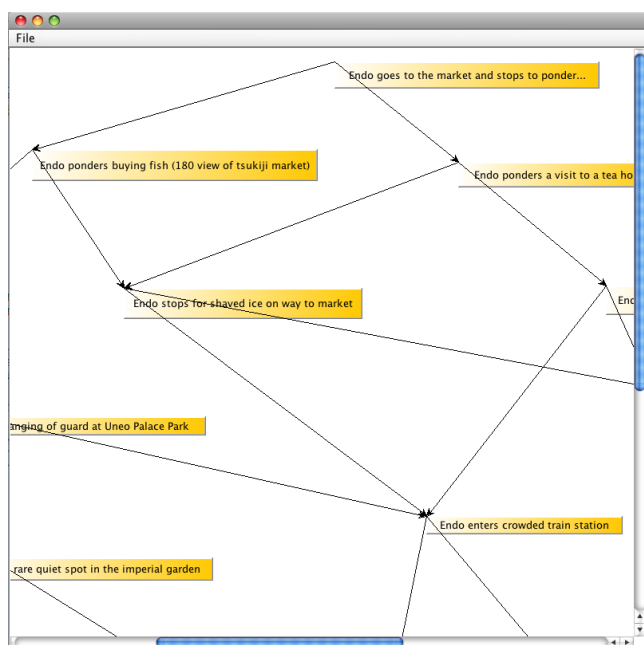


Figure 3: The graphical representation of the story structure in the authoring tool. The vertices represent story events and the edges represent the possible transitions between the story events.

post-text, and question entries. Edges represent transitions and map a story event and a question answer to the next story event.³ The graph representation of part of the story used in our study is presented in Figure 3.

To edit the story information associated with an event or a transition, there is a table of key-value pairs that appears alongside the graph. By changing a value in the table, all of the appropriate database updates are made to the relevant tables. Also, this approach allows for recycling text or videos across multiple events. If the author wishes to use a video entered into the database for an earlier event in the event they are currently editing, they can simply enter the unique video identifier assigned by the database and all the relevant fields for the event will be updated. The same holds true for pre-text, post-text, and questions. This design of the authoring tool was sufficient for the purposes of creating a story for our study, but we have updates planned prior to conducting the authoring studies described in Section 6.

4 Study Design

The story presented in our study takes place in Japan. The main character, a young man named Endo, is going about his business on a lazy Saturday. What he experiences during the day is, in part, controlled by the decisions the players make for him. The story consists of 24 events (not all of which will occur in any given episode) that are based on 16 YouTube videos, 24 authored bits of pre-text, 18 authored bits of post-text, 14 questions, and 28 answers. There are 44 transitions that link the events in various orders. The story always begins with Endo in a market and the subsequent events take place in various other parts of the city. Some of the events include going for a cup of tea, shopping for a knife, shopping for a camera, walking in the imperial palace garden, and buying fresh fish at the fish market. The first 3-4 events of the story are all very diverse.

³The concept of a GUI story editor is not new (*c.f.* [Medler and Magerko 2006]).

The conclusion of the story is always the same regardless of the choices made near the beginning.

4.1 Recruiting and Data Selection

We recruited study participants from five different locations including a special interest internet forum and mailing lists at our institution as well as two of our peer institutions. Recruiting messages asked people to help out with a “mixed-media interactive storytelling system” and the focus of the study was intentionally left vague. Further, the consent information presented to the recruits was intentionally left vague as well. It described the study as focusing on the emotional response of players to videos, even though that was not exactly our goal. We chose to do this because we did not want to bias our results by prepping participants to be on the lookout for influence messages.

We capped the number of participants at 75. That is, of all the people that clicked on the link in our advertisement, once 75 of them had filled out (or declined to answer) the demographic survey we no longer accepted data for analysis. Of those 75 participants, 71 continued to begin the actual story, and 59 completed an entire iteration of the story and answered (or declined to answer) the exit survey. Additionally, although we will not report on the data because it is beyond the scope of this study and paper, 12 of the participants continued on to begin the story at least once more.

Of the 71 participants who began the story, we threw out the data from 11 trials where the participants did not complete the story. Because the major motivation behind this paper was to determine the rate at which players entered a particular story state, including partial stories did not make sense. We found that the difference in the rate of study participants dropping out of the study was not statistically significant between the groups we will define as control ($\frac{5}{38}$) and treatment ($\frac{6}{33}$) groups ($p = 0.5298$, see “excluded” in Table 1 in Section 5).

4.2 Treatments

As mentioned in Section 2, the goal of this study is to provide data for assessing the claim that there is potential for influence statements to bring about desired behavior in players. In addition, we have data, which will be discussed in Section 5, that indicates the effect of using influence models on a player’s sense of self-agency.

Our design includes two groups that players are randomly assigned to at the beginning of each episode: a control group and an influence treatment. To limit the number of variables changed in the treatment group, before running the study we selected one of the 24 story events as our “goal”. In this case, the goal was that the player chooses for Endo to buy fish at the fish market. There are five story events that induce three possible paths the player can take that will result in Endo buying fish and four possible paths that do not result in Endo buying fish. With the exception of the first story event, the remaining four events were associated with a hand-authored influence statement of no more than 3 sentences in one of four places: before the pre-text; after the pre-text and before the video; after the video and before the post-text; or after the post-text. Figure 4 is a screenshot of the system in the influence treatment where an influence statement has been inserted after the post-text.

To further limit the number of variables changed in the treatment group, we chose to base the statements on only one of the six principles of influence presented in Section 2: **scarcity**. To effectively use scarcity as a tool of influence, the player must perceive some value in the character performing a particular behavior, and then perceive that access to that behavior will no longer be available at some point in the future. The statements were designed, by us, to

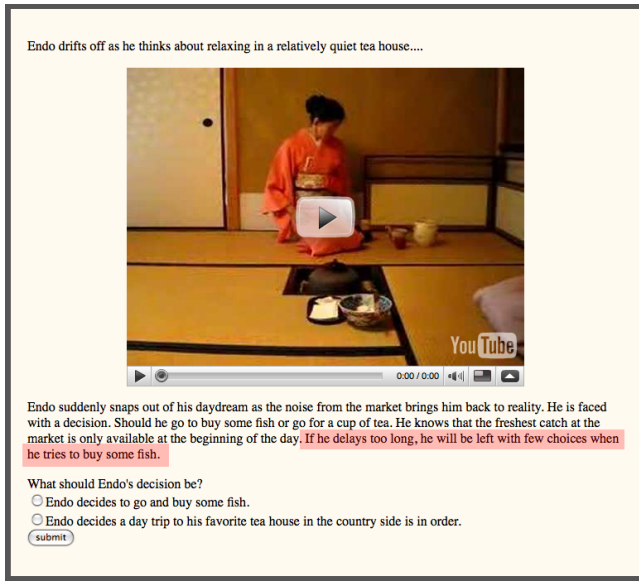


Figure 4: A screenshot of the storytelling system during a story event in which an influence statement (highlighted here) has been included. The presentation of the text, video, and question-answer sets remains the same as for the control group; however, an influence statement in the form of at most 3 sentences is added to try to shape the player’s decision making.

convey both of those things. Here are two examples: 1) “[Endo] knows that the freshest catch at the market is only available at the beginning of the day. If he delays too long, he will be left with few choices when he tries to buy some fish.” and 2) “While he waits, Endo checks his watch. With each passing minute, the chances of Endo finding a good fish to buy at the market are decreasing rapidly.”

4.3 Surveys

To analyse the similarity of the people in the control and treatment groups, we had all of the participants take part in a brief demographic survey before beginning the story and an exit survey after they completed it. They were allowed to skip the surveys or omit answers to any of the questions. We used the demographic survey to compare the populations on completely independent variables. The exit survey was used to compare the populations’ reactions to the story and their role in the plot progression.

The demographic survey consisted of four questions about education level, age, the amount of time spent playing computer games per week, and the amount of time spent on the internet per week. The subjects were asked to fill out a web form containing multiple choice questions for each of the four topics.

The exit survey consisted of five questions about the subject’s impressions of the system, their ability to control the plot progression, and their connection with the character. In each of the five questions, the participants were asked to indicate their level of agreement on a five point Likert scale [Likert 1932]. The prompts were designed to try to quantify the effect, if any, that the use of influence statements has on the player’s experience.

Variable	χ^2	d.f.	p	Fisher p
Dropout Rate:				
excluded	0.7133	1	0.3980	0.5298
Demographics:				
education level	2.8791	7	0.8960	0.9106
age	4.8730	7	0.6755	0.7521
game hours	4.1784	5	0.5240	0.6484
internet hours	4.3739	6	0.6262	0.6791
Exit Survey:				
sense of control	9.5556	5	0.0889	0.0464
adapt	2.7028	5	0.7457	0.7798
manipulated	3.8408	5	0.5726	0.6436
connection	3.3411	5	0.6476	0.6888
engagement	3.6109	5	0.6067	0.6624
Goal Realization:				
achieve goal	5.4289	1	0.0198	0.0286

Table 1: The results of a χ^2 analysis and a Fisher’s Exact Test on participant exclusions, demographic survey data, exit survey data, and goal realization.

5 Results

In short, our results are rather encouraging, although not definitive. As our ultimate goal is to use models of influence to automatically generate statements a drama manager can use to shape a player’s experience in a game environment, we opted to conduct our study in a pseudo-game environment where we were able to control for many of the variables present in a full game environment. Despite not controlling for everything, we can still claim with confidence that using statements based on the principle of scarcity in an interactive storytelling environment can lead to a statistically significant increase in the frequency with which an author’s goal is realized compared to a control group. There are various alternative explanations. The increase could be attributable to having an extra sentence, more frequent use of the word fish, etc. To fully eliminate these possibilities is the subject of a future study in a more controlled group.

Table 1 summarizes statistics from χ^2 and Fisher’s Exact Test analyses of participant responses. Due to our small sample size and response distributions, there were a number of cases where the expected values in the χ^2 computation were less than five. In such cases, Fisher’s Exact Test serves as a more appropriate statistical test [Fleiss et al. 2003] and the results can be interpreted in the same way as those of χ^2 . Accordingly, we rely largely on the p values indicated by Fisher’s Exact Test.

The analysis presented here was conducted on the data we obtained including a category for “no answer.” To be sure, we re-ran the analysis on the same data set where we excluded missing responses. We found that while the p values may have changed, the statistical significance of the results at $\alpha = 0.05$ were unchanged with one exception. As will be discussed below, the significance of the “sense of control” Likert scale responses (see “control” in Table 1) were marginally significant ($p = 0.0889$) according to the χ^2 analysis but significant at $\alpha = 0.05$ according to Fisher’s Exact Test ($p = 0.0464$) when “no answer” was included in the analysis. When “no answer” was excluded from the analysis, both the χ^2 ($p = 0.0490$) and Fisher’s Exact Test ($p = 0.0295$) indicated significance at the $\alpha = 0.05$ level.

5.1 Goal Realization

The major finding of this study is that inserting influence statements based on the principle of scarcity can have a statistically significant effect on the frequency that goals are realized. Of the 59 participants who completed the story and took part in the exit survey, 33 were randomly assigned to the control group and 27 were randomly assigned to the treatment group. In the control group, 17 out of the 33 (or 51.5%) participants bought fish during the story whereas in the treatment group, 21 out of the 26 (or 80.8%) of the participants bought fish. This difference is significant with $p = 0.0286$ (see “goal” in Table 1).

5.2 Similarity

To ensure that the significance of the goal realization result was not biased by an unbalanced population between the two groups, we compared the distribution of answers to each of the four demographic survey questions. We found the distributions did not differ significantly for any of the four questions. In addition, to justify the exclusion of incomplete data, we compared the rate of participant drop out between the two groups. We found that five out of the 38 (13.2%) participants in the control group did not complete the story whereas seven out of the 34 (20.6%) participants in the treatment group did not complete story. This difference was not significant ($p = 0.5298$).

5.3 Exit Survey

We found that participants rated their experience similarly in four out of the five categories in the exit survey. While not statistically significant ($p = 0.7798$), it seems participants in the control group indicated a slightly higher agreement that the story was adapted specifically to them (see Figure 5). Again, while not significant ($p = 0.6436$), participants in the treatment group indicated feeling slightly less manipulated by the system (see Figure 6). Insofar as this finding is significant, it might be attributable to the nature of how influence works: an exploiter (the author) can put an exploited person (the player) in a situation where they feel they have made decisions for their own reasons, free from the persuasion of an outside source, even if that is not at all the case.

The last two questions of the exit survey pertained to how the player connected with the character and engaged the system. There was no significance to the minor variations observed in the response distributions for those two questions (connection $p = 0.6888$ and engagement $p = 0.6224$).

5.4 Sense of Control

We did observe a significant difference in the responses of players in the control and treatment group ($p = 0.0464$) when asked about their feeling of control over the story progression. The full distribution of answers by group is presented in Figure 7. There are two points worth mentioning about the distributions. First, participants in the control group did not indicate strong feelings in either direction, tending to either agree or disagree (with only one participant selecting “neither agree or disagree”). On the other hand, some of the participants in the treatment group indicated either a stronger level of agreement with four responses in the strongly agree category and a stronger level of disagreement with three responses in the strongly disagree category. The cause for this greater variability of responses in the treatment group is unclear.

Second, note in Figure 7 the discrepancy between participants in the control and treatment groups who indicated overall agreement (either agree or strongly agree). In the treatment group, 13 out of 24




Figure 5: The distribution of responses in the exit survey pertaining to the player’s perception of the story being adapted to them. (“na” indicates no answer given, “sd” indicates strongly disagree, “d” indicates disagree, “no” indicates neither agree or disagree, “a” indicates agree, and “sa” indicates strongly agree)




Figure 6: The distribution of responses in the exit survey pertaining to the player’s perception of being manipulated. (“na” indicates no answer given, “sd” indicates strongly disagree, “d” indicates disagree, “no” indicates neither agree or disagree, “a” indicates agree, and “sa” indicates strongly agree)

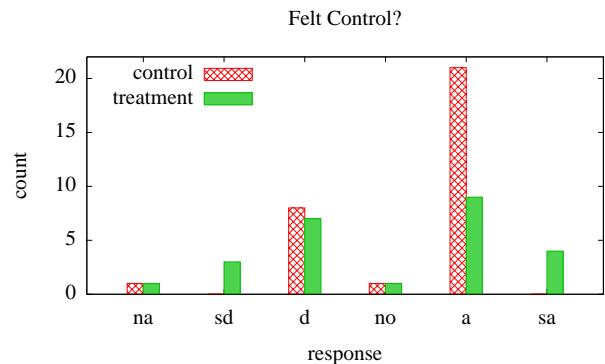


Figure 7: The distribution of responses in the exit survey pertaining to the player’s perception of control over the story progression. (“na” indicates no answer given, “sd” indicates strongly disagree, “d” indicates disagree, “no” indicates neither agree or disagree, “a” indicates agree, and “sa” indicates strongly agree)

(54.1%) indicated overall agreement whereas 21 out of 30 (70.0%) in the control group indicated overall agreement. This difference is mirrored for disagree as well. We believe we have an explanation: the control group in this study is equivalent to no drama manager; therefore, it is unsurprising that participants perceived a change in their degree of control in the treatment group.

An alternative explanation revolves around story content. Players who did not purchase fish entered a portion of the story where there were more options immediately available to them. While the branching factor of the story graph is two for all paths, there is less coherence among events in the paths that do not pass through the “purchase fish” event; that is, if a player makes a decision to cause a particular story event, it happens immediately and they are presented with another set of options. The path that leads to and from purchasing fish required the player to re-commit to a decision they had made earlier. For example, the first time the player is given the choice to purchase fish, the next story event puts them in a situation where they pass the train station and must decide whether or not to continue with their original plan to buy fish or to enter the train station and head out of the city for the day. Anecdotally, a number of participants described this as “frustrating.” Because we did not include an interview or comment box in our survey, we cannot report any qualitative results to back up this explanation. Further, this characteristic of the story only became apparent to us in a *post hoc* analysis after speaking informally with a handful of study participants.

Given this anecdotal evidence and alternative explanation, an additional point of interest is the apparent lack of significant difference in a player’s feelings of being manipulated. The distributions of answers by group for feeling manipulated are presented in Figure 6. This data is very encouraging when taken in concert with the alternative explanation. If it is in fact the case that participants did not feel any more manipulated when influence statements were used and the change in feelings of control are attributable to story structure rather than influence, then we might have an indication that the player’s *sense* of self-agency is preserved when models of influence are used compared to using no drama manager (we already know the system’s affordance for player self-agency is preserved). This claim remains speculative at this point, but we plan to conduct additional experiments to try to verify that this claim can be made, perhaps by changing the influence goal, conducting a post-play interview with participants, or performing a “think-aloud” protocol with a small sample of participants to discuss their interpretation of the exit survey prompts.

6 Conclusion and Future Work

In this paper, we have presented an argument for using models of influence as tools to shape players’ experiences in interactive settings. We have presented a choose-your-own-adventure-style architecture that we have implemented and used to study and quantify the effect of influence statements based on the principle of scarcity. We have found that there is a measurable difference in players’ behavior when influence statements are used and that the difference is statistically significant. We have also presented data that describes the effect of the influence statements on the players’ perceptions of their experience. We have shown that in four out of the five categories surveyed, there was no statistical difference between the control group and the treatment group. In the one category that does show a statistical difference, we have presented two plausible explanations for the cause of that difference, and described future experiments we plan to run that will help to identify the proper explanation rigorously.

In addition to those experiments, we also plan to conduct authoring

studies. A major motivation for the development of our models of influence is to reduce the burden on the author. To study the effect the models may have, we plan to use the authoring tool described in Section 3 (in a modified form) as a platform for comparing the authorial effort in creating a story when the author can specify goals in a language we provide and allow the system to generate the drama manager action automatically rather than using hand-authored drama manager actions.

The output of that study will be a set of stories that have been authored both with a hand-created set of drama manager actions and with a set of system generated drama manager actions. Those stories will form the basis for comparing how players’ perceptions of their experience change between author-created actions and actions created using the influence model. These studies will enable us to re-verify the results obtained in this paper regarding the significant increase in goal realization in the treatment case.

Lastly, we are interested in performing a deeper analysis of the data, especially in situations where participants played more than once. In this study, we never intended to include data from more than one episode per participant, but we did allow participants to play as many times as they wished. There was a relatively high dropout rate in episodes subsequent to the first; nonetheless, six participants completed two trials and one completed four. We would like to use this data to examine the lasting effects of influence. For example, one of the participants was randomly assigned to the control group during their first episode and did not buy fish. In the second episode, the participant was assigned to the treatment group and did buy the fish. What if this were the other way around? What if they were in the treatment group during their first episode; would they still have a tendency to buy fish during their second episode even if they were no longer in the treatment group?

In conclusion, we are highly encouraged by the findings presented in this paper. We have presented a strong argument for further study into the efficacy of using influence tools for automatically implementing drama manager actions. The results presented in this paper indicate that this approach has the potential to be successful with further development. We have identified and presented a few techniques that will guide the development of future studies on this topic to enable us to more definitively verify our three main goals of providing tools that enable authors to effectively shape player experiences, reducing the author’s burden, and doing so without the player perceiving a decrease (and preferably perceiving an increase) in self-agency.

References

- ARIELY, D. 2008. *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. Harper Collins.
- BHAT, S., ROBERTS, D. L., NELSON, M. J., ISBELL, C. L., AND MATEAS, M. 2007. A globally optimal algorithm for TTD-MDPs. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS07)*.
- CIALDINI, R. B. 1998. *Influence: The Psychology of Persuasion*. Collins.
- DOYLE, P., AND HAYES-ROTH, B. 1998. Agents in annotated worlds. In *AGENTS '98: Proceedings of the Second International Conference on Autonomous Agents*, 173–180.
- FLEISS, J. L., LEVIN, B., AND PAIK, M. C. 2003. *Statistical Methods for Rates and Proportions*, 3 ed. John Wiley & Sons.

- LAUREL, B. 1986. *Toward the Design of a Computer-Based Interactive Fantasy System*. PhD thesis, Drama Department, Ohio State University.
- LIKERT, R. 1932. A technique for the measurement of attitudes. *Archives of Psychology* 140, 1–55.
- MATEAS, M., AND STERN, A. 2005. Structuring content in the Façade interactive drama architecture. In *Proceedings of Artificial Intelligence and Interactive Digital Entertainment (AI-IDE05)*.
- MATEAS, M. 1999. An Oz-centric review of interactive drama and believable agents. In *AI Today: Recent Trends and Developments. Lecture Notes in AI 1600*, M. Woodridge and M. Veloso, Eds. Springer, Berlin, NY. First appeared in 1997 as Technical Report CMU-CS-97-156, Computer Science Department, Carnegie Mellon University.
- MEDLER, B., AND MAGERKO, B. 2006. Scribe: A general tool for authoring interactive drama. In *Proceedings of the Third International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, 139–150.
- MOTT, B. W., AND LESTER, J. C. 2006. U-director: a decision-theoretic narrative planning architecture for storytelling environments. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS06)*, ACM Press, New York, NY, USA, 977–984.
- NELSON, M. J., ROBERTS, D. L., ISBELL, C. L., AND MATEAS, M. 2006. Reinforcement learning for declarative optimization-based drama management. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS06)*.
- RIEDL, M. O., STERN, A., DINI, D., AND ALDERMAN, J. 2008. Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4, 2.
- ROBERTS, D. L., AND ISBELL, C. L. 2008. A survey and qualitative analysis of recent advances in drama management. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4, 2.
- ROBERTS, D. L., ISBELL, C. L., RIEDL, M. O., BOGOST, I., AND FURST, M. L. 2008. On the use of computational models of influence for interactive virtual experience management. In *Proceedings of the First International Conference on Interactive Digital Storytelling (ICIDS08)*.
- WEYHRAUCH, P. 1997. *Guiding Interactive Drama*. PhD thesis, School of CS, CMU, Pittsburgh, PA. Tech Rept CMU-CS-97-109.
- YOUNG, R. M., RIEDL, M. O., BRANLY, M., AND JHALA, A. 2004. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development* 1, 1, 54–70.