Shakespearean Spatial Rules

Christine Talbot and G. Michael Youngblood
Game Intelligence Group
University of North Carolina at Charlotte
{ctalbot1, youngbld}@uncc.edu

ABSTRACT

Many games and simulations utilize modularized low-level code to move characters about in an environment. This requires extensive technical skill to translate basic high-level actions, as well as extensive time to write code, which includes very detailed instructions on what and when actions will occur across all agents. Other options exist such as mocap files; however, most are not highly dynamic, concerned with spatial positioning, or require human intervention to solve the problem.

This paper presents an approach that utilizes play-scripts and natural language processing, along with some spatial reasoning rules to control characters in a virtual environment. Rules around grouping of characters, conversational space, theatre, and general behaviors are key in fully interpreting a play-script into movements on stage. These rules help us to achieve similar blocking for the Shakespearian play Hamlet, performed by virtual characters, as the director Sir Gielgud produced for his 1964 production of Hamlet.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Miscellaneous

General Terms

Algorithms, Measurement, Performance, Design, Experimentation, Human Factors, Languages, Theory

Keywords

Shakespeare; Hamlet; Spatial Reasoning; Natural Language Processing; Plays; BML Realizer; SmartBody; Speech Act Theory; Agent Reasoning; Agent Planning

1. INTRODUCTION

Current research focuses mainly on positioning characters side-by-side.There has been limited work on positioning virtual characters within a scene to support the current actions being performed. Most of this work is focused on nonverbal behaviors and interaction with humans to make the virtual characters seem more realistic.

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

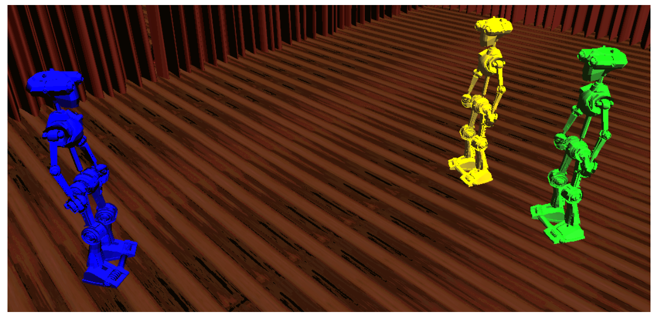

Figure 1: 3D Enactment of Hamlet in Unity Using the SmartBody BML Realizer

The movie industry has also utilized animated and virtual characters based on real actors' movements recorded via mocap files. This group comes closest to taking into consideration the implications of spatial reasoning for controlling virtual characters. Their methods of recording motions as they are being performed by actors provides intricate details for replaying the motions. However, it comes with several drawbacks, such as expensive tools, good actors, and creation of realistic environments to perform in. It is not very dynamic and every situation must be recorded for the exact situation being simulated.

The gaming industry relies on modularized low-level code to move characters about in an environment. This requires extensive technical skill to translate high-level actions, as well as extensive time to write all of that code. Most movement is hard-coded on what can be done and when it will occur.

A newer option includes a Functional Markup Language (FML) [33], Behaviour Markup Language (BML) [3], and BML Realizers [34] like SmartBody (Figure 1) [32]. These also require some lower-level coding, but begin to abstract and parameterize the motion of the characters. It creates more dynamic and repeatable motions for characters.

The problem is that this method still requires a gamewriter to write very specific and detailed steps. With BML, one must specify where the character looks, when they look there, how they move, when they move, and when they should pick-up/put-down objects. This can be very time-consuming, as we saw with our prior work [29] where it took over four hours to write BML for a ten minute scene. Not everyone is doing this by hand, but many of the people automating the BML commands are focusing on generating characters that stand side by side and talk to each other.

We are more interested in where they will be when they interact, as well as their spatial movement throughout a scene.

As humans, our approach for giving directions is much more vague than any of the previously mentioned approaches for describing motion. For instance, we do not typically give exact information on how far to go, do not specify common sense things like the road curves left, and do not remind folks to take the elevator to get to the third floor.

So the question is, how can we control our characters with the deep-level control that the mocap files give us while providing authors a more natural and high-level way of describing the actions the characters should take? We observed that play-scripts (used by almost all theatre and movie productions) provide this capability already [1].

In play-scripts, we give high-level directions to the actors for where to go, what to do, and what to say. Many people writing dialogues in interactive media are essentially playwrights. Therefore, writing play-scripts is already a part of their education. We do not have to invent a new way to write spatial directions.

We translated character motions from play-scripts by utilizing some basic part of speech tagging and named entity recognition natural language processing techniques.We then applied some spatial rules (the focus of this paper) to provide more realistic blocking and movement of the characters throughout the play. We compared a manually mapped baseline of Sir Gielgud's Hamlet on Broadway in 1964 [26] to both a strict natural language processing translation of motion [29] and our rule-based motion versions of this same act of Hamlet.

These techniques can be applied more broadly since they only rely on the components that are inherent to play-scripts, movie scripts, and television scripts. The only scene-specific setups are ones based on identifying the characters and starting positions of key props within the scene–all of which are part of the manual setups of any scene for any play.

Previously, we found that the natural language processing techniques were able to capture most of the character position changes from the play-script, however did not do too well with the gazing directions [29]. We expect that after adding in the rules engine, we will find an improvement in gazing directions with respect to the baseline from the Broadway production's video. The character traces should also show slight improvement in character positioning due to the conversational, grouping, and theatre rules that are applied.

2. BACKGROUND

William Shakespeare has written at least three of the top ten most-produced plays in North America, despite the fact that most lists explicitly exclude Shakespeare's plays from their top ten lists as it would be unfair [12]. Shakespearean plays are also notorious for not containing many director annotations to assist with enacting his plays. This has lead to the many different interpretations of Shakespeare's plays over the last 400 years [21].

We focused on one particular famous production of one of Shakespeare's plays from 1964. Sir John Gielgud directed Hamlet on Broadway with Richard Burton playing Hamlet. This production ran for 138 performances, setting the record as the longest-running Hamlet ever to play New York [26]. It was filmed during three successive stage performances in June/July 1964 by Electronovision, Inc. [5]. In addition,

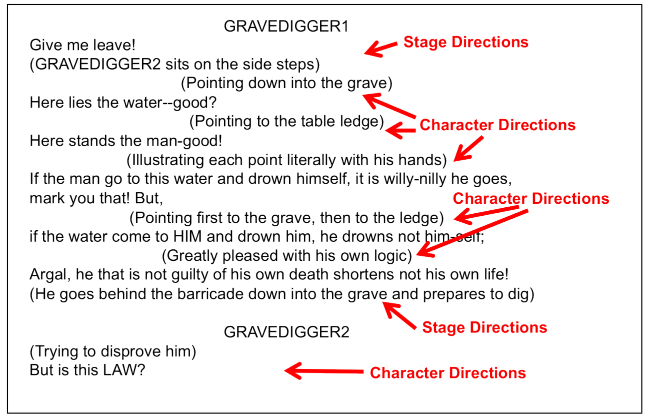

Figure 2: Play-Script Excerpt from Sir Gielgud's Hamlet on Broadway 1964

Richard Sterne (another actor in this particular production) published a book with very detailed director's annotations and notes for the entire play [26]. This additional level of detail provided us a detailed baseline to compare our work.

Play-scripts, such as the one used from Sir Gielgud's Hamlet, provide most of the direction and motivations to the actors regarding the director's intended interpretation of the play. The key component of interest in play-scripts stems from the annotations regarding what the actor should be doing and how. These annotations can be broken up into scene, stage, and character directions as seen in Figure 2. Scene directions are found at the beginning of the scene and setup the stage and characters for the scene. Stage directions describe what needs to occur on-stage during a scene, such as entrances, exits, major movements of characters, and so forth. Finally, character directions provide details to the motivation of a line and how it should be performed. [1]

Play-scripts provide a natural way of directing actors and characters, including any relevant spatial directions. Therefore, these two components (Sir Gielgud's Hamlet and play-scripts) are key to our approach in this work as they provide both a baseline and a user-friendly structure to communicating spatial movements in a scripted environment.

3. RELATED WORK

The focus of much research has involved virtual characters; however, very little of this work has investigated spatial movement of those characters. Therefore, we focus this section on the spatial reasoning in the natural language processing, cognitive psychology, and robotics research areas.

In the natural language processing community, many researchers are working towards better understanding of the written and spoken word. There is quite a bit of work in niche areas for natural language understanding, such as a focus on spatial language expressions. These examine different prepositions, which indicate the temporal, spatial, or logical relationship of objects to the rest of the sentence (e.g., *in, on, near, between*). For instance, Regier built a system that assigns labels such as "through" or "not through" to movies showing a figure moving relative to a ground object for learning how we qualify the particular term "through" [10]. Kelleher and Costello [13] and Regier and Carlson [23] built learned models for the meanings of static spatial prepositions such as "in front of" and "above" while Tellex focused on "across" [30].

Some groups are pursuing the complexities of spatial cognition within language on object representations and geometry, as well as the number and structure of the objects utilizing the prepositions that situate them in space [16]. Kelleher also proposed a framework for understanding prepositions primarily around the closeness of objects and the visual representation of those objects [13]. His research explores how humans describe where objects are within space, which is key in extracting spatial information from natural language. This information has been used by other methods, such as WordsEye, which takes natural language to draw a scene utilizing the spatial locations described in text [7].

From the perspective of cognitive psychology of language, Coventry describes spatial language and how humans describe different situations using prepositions, such as a pear being in a bowl or not. He elaborates with many different prepositions such as *in*, *on*, *near*, *far*, *at*, and *between* [6]. However, these prepositions are very dependent on the frame of reference used for the spatial description. Describing spatial locations using an intrinsic, absolute, or relative frame of reference can dramatically change the interpretation of the same sentence [17]. Stating "a ball is in front of the chair" can mean different things depending on which way the object is facing, where the observer is, or what global spatial reference that is being used—all with respect to which reference the person describing the spatial relationship is using.

Once we are able to determine the frame of reference being used for the spatial descriptions, we can utilize methods of mapping objects based on cardinal directions as described in Frank's work [8]. Other methods include the use of spatial templates to identify acceptable locations with respect to a given object for a particular preposition [19], and vector sum models [23] to formalize spatial relationships.

Roboticists have pursued an understanding of spatial language primarily to understand verbal instructions. This can be seen in many works, such as Brooks's thesis where he attempted to train a robot to be an actor using verbal directions. The robot could not speak, but shrugged if he did not understand the directions [4]. This is a different approach to teaching a character to enact a scene of a play; however, Brooks' approach required a more detailed and lower-level of communication to his robot than is typically found in a play-script. David Lu and Bill Smart's work with robots in theatre has focused around mimicking actor's movements with robots to help incorporate social interactions into robots without explicitly programming them [20]. They used actors to perform specific scenarios and replicated them on robots, making their movements more believable. These were generalized to similar situations and to robots that could not physically replicate the original motions.

Other research utilizing virtual agents focuses primarily on the conversational and nonverbal domains, such as Thespian [24], Virtual Storyteller [31], and Stability and Support Operations (SASO) [15]. The emphasis appears to be more on the speech and emotional interaction with humans or other characters. However, with the growing focus on realistic virtual environments, the spatial domain is becoming a more critical component in creating that realism.

Markup languages, such as Behavior Markup Language (BML) [22], are making it possible to abstract the control of virtual characters. BML abstracts the physical realization of behaviors and movements, along with their constraints. It is not concerned with the intent behind the movements [22] .BML is structured like a typical XML message. One can control what is done, when it is done, and what runs concurrently with other commands. However, it is often at such a low-level that this can be extremely time-consuming to build, especially for things like non-verbal behaviors (eye saccade, gesturing while speaking, head nods, and so forth). Plus, writers must be fluent in these technical languages, plan out specific points and marks within the environment, and convert the more fluid, natural descriptions into more concrete commands with fewer human assumptions.

4. METHODOLOGY

4.1 Baseline

In our previous work [29], we utilized the Electrovision video [5] and annotated play-script [26] to hand-map the movements and positions of the characters in the Graveyard scene on stage (Hamlet ACT V, SCENE 1). We used this mapping as our ground-truth to compare a basic natural language translation of the same annotated scene. The natural language module was based on a simple part of speech tagging and named entity recognition process that focused primarily on the scene and stage directions within the play. It takes a command, such as:

"GRAVEDIGGER1: (Pointing down into the grave)"

and translates it into

actor=GRAVEDIGGER1
action=POINT
target=GRAVE

This information was translated directly into a BML command for GRAVEDIGGER1, such as:

<gesture lexeme="POINT" target="GRAVE" />

This produced a reasonable translation of the movements of the characters on the stage, however did not do well with gaze and facing directions [29].

4.2 Rules

With this work, we look to expand upon the natural language processing to incorporate rules to better our translation of motion. We have pulled from many different areas to encompass the types of rules that are typically utilized when performing plays. We have categorized these rules into four basic areas:

1. Grouping Spatial Rules
2. Conversational Spatial Rules
3. Theatre Rules
4. General Rules

In the next few sections we discuss what is involved in each of these rule groups to provide a background for our work.

4.2.1 *Grouping Spatial Rules*

Jan describes six different forces that affect when/why a person may shift position when in a group of people; however, the main reason that could affect the positioning of characters in a play is that one person is too close to others to be comfortable, or proxemics [11]. Hall describes four different zones that personal space is divided into: intimate, personal, social, and public zones [9]. The actual distances involved in each zone differs for each culture and its interpretation may vary based on an individual's personality. If the speaker is outside the participant's comfort area, the participant will move toward the speaker. Similarly, if someone invades the personal space of a participant, the participant will move away [11]. Also, when there are several people in a conversation, they will tend to form a circular formation. This provides a sense of inclusion for all participants and provides a better view of all members while conversing [14].

4.2.2 Conversational Spatial Rules

Older research from psychology shows that people prefer to be across from one another than side-by-side in most situations, but there is importance to the surrounding area for determining the distance that is comfortable [25]. Also, friendship and attraction can affect the spatial distances between people by decreasing them, while negative attitudes may not have much affect on the spatial distances [27].

According to studies reviewed by Sundstrom, comfortable face-to-face distance for speaking while sitting is approximately five feet and comfortable face-to-face conversation standing is approximately three feet [27]. He also discusses the effects of spatial invasion for character behaviors and movements and provides a nice overview of multiple research efforts looking at conversational space for both sitting and standing positions [27].

4.2.3 Theatre Rules

In the theatre, there are special rules and conventions when staging a play. Many of these guidelines revolve around engaging the audience and visibility onstage. To help with this, the stage is often split into nine areas upon which basic theatre rules are based. They consist of upstage, stage right, stage left, downstage, and combinations of each as shown on the bottom right of Figure 6.

Being downstage (near the audience) is a stronger position than being upstage and should be held by the most important characters in the scene. Also, because we tend to read left to right, downstage right is the most powerful position onstage as audiences tend to look left first, then scan right when watching a play. The more important a line is, the more likely an actor is to fully face the audience, although the most common position is a one-quarter (or 45° angle from the audience) body position as it ensures the audience can see all the characters on the stage properly. Actors should never turn their back to the audience. [2]

Moving onstage can cause many issues including upstaging and covering. Both of these issues should be avoided, which in turn provides additional rules to characters on the stage. Upstaging is where one actor takes a position further upstage, or above a second actor, which causes the second actor to face upstage/away from the audience. Therefore this must be avoided to ensure actors do not present their backs to the audience, especially if both characters are just as important to the scene [18].

Covering occurs when one actor blocks the audience's view to a second character onstage. If this does happen, the covered actor should adjust to provide visibility of him/herself to the audience by counter-crossing (performing a movement

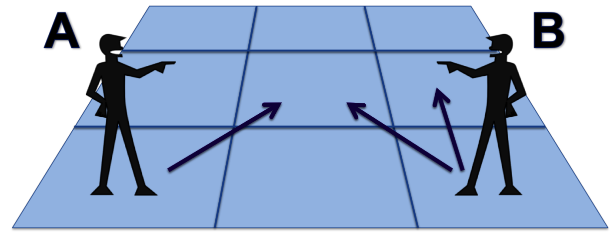

Figure 3: Example of a Counter-Cross: Actor B could move to Center Stage or to Right Stage to Counter Being Upstaged by Actor A

in the opposite direction of the other actor—see Figure 3). When making these changes, actors should cross downstage from other actors unless their movement should not be noticed by the audience. Finally, when crossing the stage, it will take two separate crosses (movement from one area of the stage to another) to cross upstage–one to the left or right, turn in, then the second to cross upstage [2].

4.2.4 General Rules

The last group of rules encompasses all those things that we often think of as common sense. For instance, when we are walking we are usually looking at where we are headed. Similarly, when we pick up or point to an object, we tend to look at it; and when we are listening to someone, we look at the speaker. When someone points to something or something/someone moves, we are usually drawn towards looking at that person or object. If someone wants to pick up an object, they need to be close to it. Finally, characters should always perform natural movements and not have their gaze or orientation jump from one position to another.

4.3 Architecture

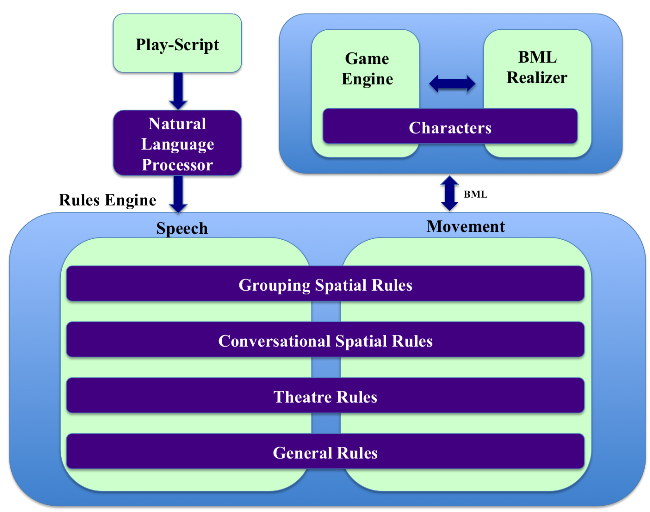

Figure 4: Rules Engine Architecture

When we put all these rules together, we are able to formulate an intricate engine to control the movements of the characters to present a realistic interpretation of the play, similar to an actor. We combined the use of a standard annotated play-script with a natural language processor, which

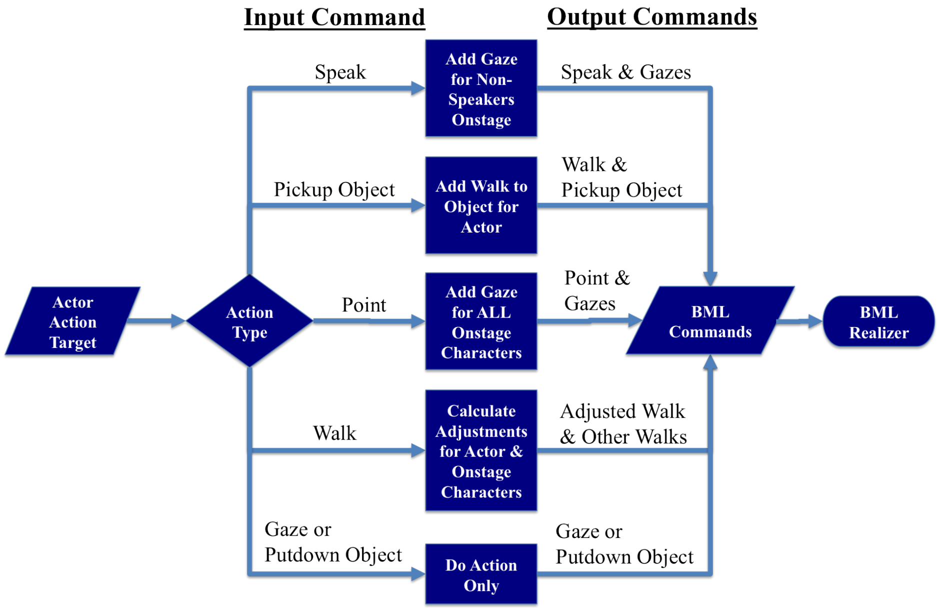

Figure 5: Logic in the Rules Engine

utilizes a part of speech tagging and named entity recognition module to extract the high-level movements of the characters.

These movements were fed into our rules engine (as seen in Figure 4) to adjust the motion based on these rules:

r_1: Characters should face the audience as much as possible, and avoid turning their back to the audience

r_2: Characters should face the person speaking

r_3: Characters with higher importance or larger roles should be placed slightly closer to the audience relative to lesser role characters

r_4: Characters should try to stay closer to center line as much as possible to improve visibility for the maximum portion of the audience

r_5: Characters should avoid unnatural movements by adhering to basic frame coherence rules, such as not having their gaze or orientation jump from left to right immediately

r_6: Characters should maintain appropriate personal space based on inter-character relationships within the play

r_7: Characters should be next to an item they wish to pick up

As the natural language processor identifies the action that needs to be performed, it sends it into our rules engine as an actor-action-target command. From there, our rules engine applies these seven rules to the action, translating it to one or more BML commands that are sent to the BML Realizer and Game Engine. A high-level overview of the process flow can be seen in Figure 5.

For speech commands, the rules engine adds additional commands for each onstage character to look at the speaker. This angle is adjusted based on the current position of the characters to ensure no one is looking more towards backstage than the audience. The speaker's gaze is also adjusted to look at the last speaker, assuming that character is still onstage.

With walk or locomotion commands, the rules engine takes into consideration the position of all the characters onstage to determine the best destination with respect to the requested target. Each character's overall importance to the scene was prioritized such that every character's importance relevant to every other character was clear, such as below:

$$Hamlet > Gravedigger1 > Gravedigger2 > Horatio$$

As can be seen above, Hamlet was the most important character in the scene, followed by Gravedigger1. This prioritization was used to determine who should be closer to the audience at any point of time. If the action's actor defined by the natural language processor (actor character) had a higher priority than one or more characters onstage, then the lower priority character(s) were moved to adjust for the relocation of the actor character, ensuring the distance to the audience was shorter for the higher priority character(s).

Also, when characters were directed to approach another character, the target locations were adjusted to accommodate any grouping or conversational space. If they were approaching a single character, they were directed to stop at approximately three feet from the other character. If they were approaching two or more characters, they were instructed to maintain an arc-like configuration facing the audience and maintain three feet from the closest character.

These character spacing adjustments were performed only once per annotation which incurred a walk command. This prevented characters from constantly adjusting and creating unnatural movements onstage, as well as aligned the timings of the movements with the intended actions within the play. When a command is sent for a character to pickup an object, the rules engine will check to see where the character is on stage with respect to the target object. If they are not

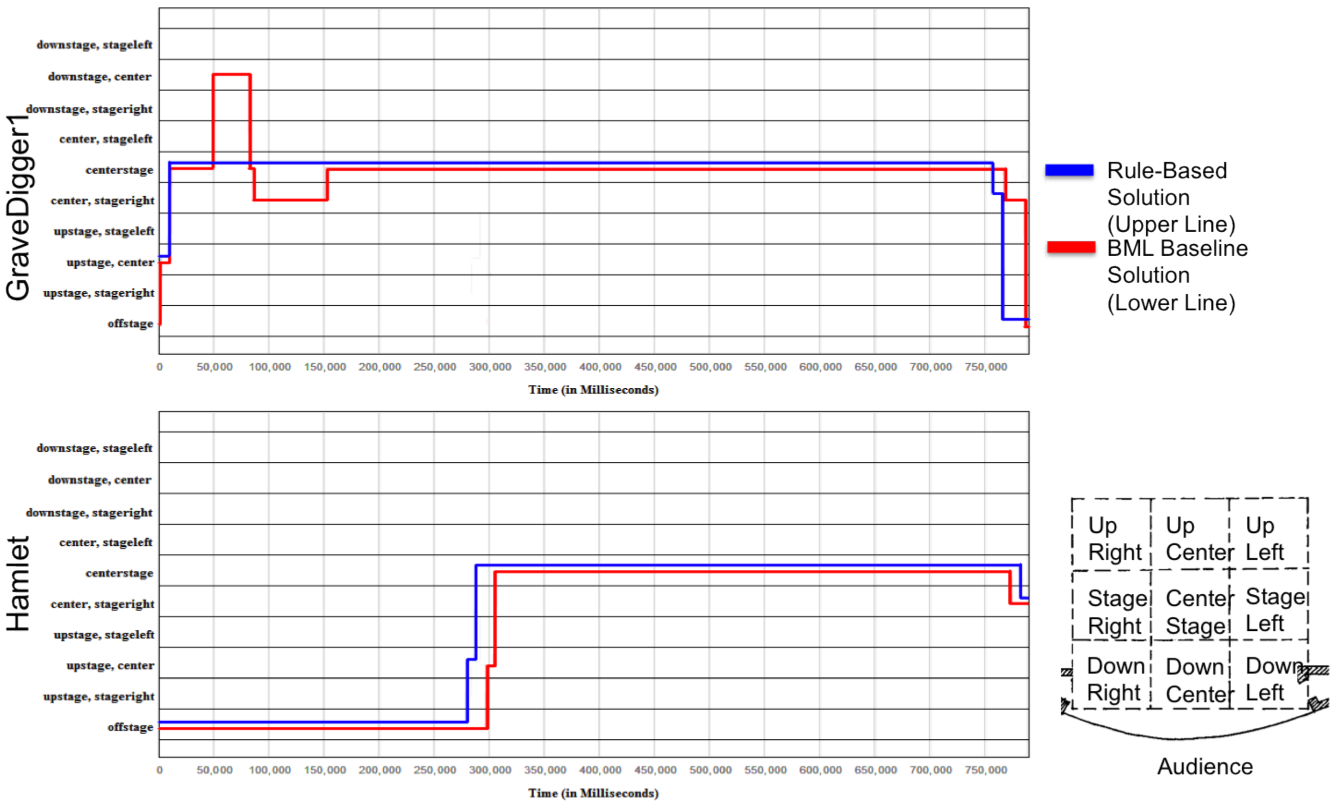

Figure 6: Comparison of Character Traces for Position Over Time (ms) and Stage Grid Diagram

near the object, they will walk to the object before trying to pick it up. If this movement conflicts with any of the aforementioned stage locations based on character importance, the other character(s) will receive a walk command to move them to an appropriate location.

Finally, as a character pointed to a target, the characters that are onstage are directed to look at what the character is pointing to. With gazing and releasing objects, the BML Realizer handled ensuring appropriate frame coherence for the characters and did not require any additional logic before performing the action(s). Therefore, these commands were submitted directly to the BML Realizer and Game Engine for controlling the characters.

5. EXPERIMENTATION

We took the character traces from both our ground truth (hand-coded BML based on the Electronovision video [5]) and our natural language processor with a rules engine and compared them. We wanted our new method to result in character positioning as close to our baseline as possible; however, we did not want to penalize for being "close enough."

As can be seen in Figure 6, overall we were able to position characters on the stage well, despite the natural language processing issues that come with any machine translation. During analysis, we split the stage into the nine squares to represent the nine general locations on the stage– combinations of: upstage, downstage, center-stage, stage-right, and stage-left (as seen in the lower-right of Figure 6).

We found that our method was able to position the characters within 0.12 squares (Euclidean distance) of our baseline BML method and placed them correctly 88.9% of the time on the stage. The other 11.1% of the time, the characters in the video added their own unannotated movements to what was directed by the director. For instance, near the beginning of the scene, Gravedigger1 walks towards the audience, then turns around and heads back towards the grave. This movement was not annotated in the play-script and therefore was not performed by our rules-based characters. This highlights one aspect of the actor's initiative to improvise despite the directions provided by the script.

For gaze, we divided the directions into the four basic gaze directions: towards the audience, stage-right, stage-left, and upstage/backstage as can be seen in the lower-right of Figure 7. Here we found our results did not match as well (as seen in Figure 7), with the gaze being correct only 52.7% of the time and, on average, within 0.53 quadrants of our baseline gaze direction.

One key reason for some of the discrepancies in the character traces is due to the input utilized for the ground-truth vs our method. The ground-truth BML was written to include movements and motion that were not included in the play-script that our method utilized, but the actors performed. It included some movements based on what was seen in the video, but may not have fully encompassed all the gazes that occurred within the play due to user-translation error. Also, our rules were based on always performing adjustments with every command that was brought into the rules engine, whereas a real actor may not follow these rules 100% of the time. However, our rules did better than our prior version which just utilized a natural language processor by approx-

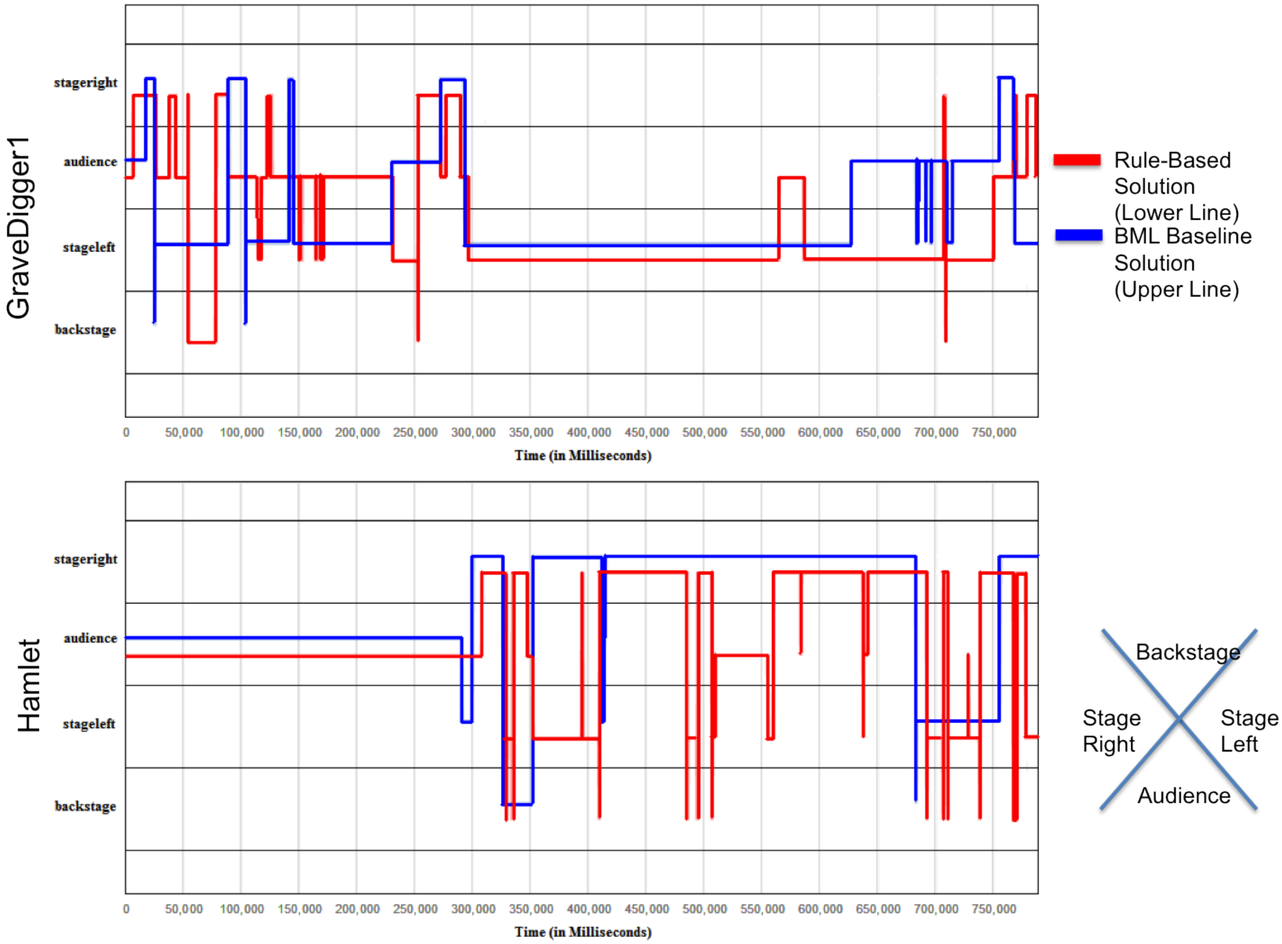

Figure 7: Comparison of Character Traces for Gaze Direction Over Time (ms) and View Angle Diagram

imately 15% for position and approximately 30% for gaze, even though it still incurred similar issues around duality of word meanings and pronouns found in our first experiment.

6. CONCLUSIONS

The results of this experiment show that adding rules helps with a better blocking of the play from a spatial perspective. It confirms our hypothesis that adding rules to control the spatial movements of our characters can more fully encapsulate the decisions actors and directors make when performing a play. We reduced our authoring time by four hours from our prior work by utilizing the natural language processing of the annotated play-script, and improved our blocking accuracy for both position and gaze with our rules-based approach.

This work has focused on the theatre, however many of the rules are also applicable to other applications of spatial positioning, such as games and virtual worlds. It does not apply the optimizations of theatre seating visibility (similar to multiple camera angles in television and movies) at this time, but this will be pursued in future work. Also, more robust natural language processing and more stringent rules will be required as we expand to additional Shakespeare plays due to the need to capture more complex spatial actions.

7. REFERENCES

[1] Abigail-Nicole. How to Format a Stage Play–Script Frenzy. http://www.scriptfrenzy.org/howtoformatastageplay, July 2012.

[2] K. Baker. Stage Movement and Acting Rules. http://campuses.fortbendisd.com/campuses/documents/teacher/2010/teacher_20100525_1349.pdf, Jan. 2002.

[3] S. Bml. Wiki - BML 1.0 Standard - Mindmakers. http://www.mindmakers.org/projects/bml-1-0/wiki/Wiki, July 2012.

[4] A. G. Brooks. *Coordinating Human-Robot Communication*. PhD thesis, MIT, Jan. 2006.

[5] B. Colleran, J. Gielgud, W. Shakespeare, R. Burton, H. Cronyn, A. Drake, and E. Herlie. *Hamlet*, Electronovision, Inc., 1964.

[6] K. R. Coventry and S. C. Garrod. *Saying, Seeing and Acting: The Psychological Semantics of Spatial Prepositions (Essays in Cognitive Psychology)*. The Psychological Semantics of Spatial Prepositions. Psychology Press, 1 edition, Mar. 2004.

[7] B. Coyne and R. Sproat. Wordseye: An automatic text-to-scene conversion system. In *Proceedings of the*

28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01, pages 487–496, New York, NY, USA, 2001. ACM.

[8] A. U. Frank. Qualitative Spatial Reasoning: Cardinal Directions as an Example. *International Journal of Geographical Information*, 10(February 2012):269–290, 1996.

[9] E. T. Hall, R. L. Birdwhistell, B. Bock, P. Bohannan, J. Diebold, A. Richard, M. Durbin, M. S. Edmonson, J. L. Fischer, D. Hymes, S. T. Kimball, W. L. Barre, S. J. Lynch, Frank, J. E. McClellan, D. S. Marshall, G. B. Milner, H. B. Sarles, G. L. Trager, and A. P. Vayda. Proxemics [and Comments and Replies]. *Current Anthropology*, 9(2/3):pp. 83–108, 1968.

[10] International Computer Science Institute and T. Regier. The Acquisition of Lexical Semantics for Spatial Terms: A Connectionist Model of Perceptual Categorization, 1992.

[11] D. Jan and D. R. Traum. Dynamic Movement and Positioning of Embodied Agents in Multiparty Conversations. Embodied NLP '07, pages 59–66, Stroudsburg, PA, USA, 2007. Association for Computational Linguistics.

[12] B. D. Kabialis. When Shakespeare Dependence Hits Point of Diminishing Returns. http://www.berkeleybeacon.com/arts-and-entertainment/2012/2/2/when-shakespeare-dependence-hits-point-of-diminishing-returns, 2012.

[13] J. D. Kelleher and F. J. Costello. Applying Computational Models of Spatial Prepositions to Visually Situated Dialog. *Multiple values selected*, 35(2):271–306, Jan. 2012.

[14] A. Kendon. *Conducting Interaction: Patterns of Behavior in Focused Encounters*. Studies in Interactional Sociolinguistis. Cambridge University Press, Nov. 1990.

[15] P. Kenny, A. Hartholt, J. Gratch, W. Swartout, D. Traum, S. Marsella, and D. Piepol. Building Interactive Virtual Humans for Training Environments. *The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC)*, 2007(-1), 2007.

[16] B. Landau and R. Jackendoff. "What" and "Where" in Spatial Language and Spatial Cognition. *Behavioral and Brain Sciences*, 16(02):217–238, 1993.

[17] S. C. Levinson. Frames of Reference and Molyneux's Question: Crosslinguistic Evidence. *Language and Space*, pages pp. 109–169, 1996.

[18] P. B. Lile. Drama Stages and Stage Movements. http://www.graves.k12.ky.us/schools/gcms/plile/Drama%20Stages%20and%20Stage%20Movements.ppt.

[19] G. D. Logan and D. D. Sadler. A Computational Analysis of the Apprehension of Spatial Relations. *Language and Space*, pages pp. 493–529, 1996.

[20] D. V. Lu and W. D. Smart. Human-robot interactions as theatre. In *RO-MAN 2011*, pages 473–478. IEEE, 2011.

[21] A. Mabillard. The Chronology of Shakespeare's Plays. http://www.shakespeare-online.com/keydates/playchron.html, 2000.

[22] A. Marshall, H. Vilhjalmsson, S. Kopp, M. Kipp, M. Krieger, M. Wissner, P. Tepper, J. Homer, H. V. Welbergen, A. Hill, T. Bickmore, and J. Gruber. Behavior Markup Language (BML) Version 1.0 (Proposal). http://www.mindmakers.org/projects/saiba/wiki/Wiki/, 2011.

[23] T. Regier and L. A. Carlson. Grounding Spatial Language in Perception: An Empirical and Computational Investigation. *Journal of Experimental Psychology: General*, 130:273–298, 2001.

[24] M. Si, S. C. S. Marsella, D. V. Pynadath, and M. Rey. Evaluating Directorial Control in a Character-Centric Interactive Narrative Framework. *Proceedings of 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 1289–1296, 2010.

[25] R. Sommer. The Distance for Comfortable Conversation: A Further Study. *Sociometry*, 25(1):111–116, 1962.

[26] R. L. Sterne. *John Gielgud Directs Richard Burton in Hamlet by Richard L. Sterne*. Random House, 5th edition, 1967.

[27] E. Sundstrom and I. Altman. Interpersonal Relationships and Personal Space: Research Review and Theoretical Model. *Human Ecology*, 4(1):47–67, 1976.

[28] C. Talbot. Virtual Companions and Friends. In *Proceedings of the 49th Annual Southeast Regional Conference*, pages 356–357, New York, NY, USA, 2011. ACM.

[29] C. Talbot and G. M. Youngblood. Spatial Cues in Hamlet. In Y. Nakano, M. Neff, A. Paiva, and M. Walker, editors, *Intelligent Virtual Agents*, pages 252–259. Springer, 2012.

[30] S. Tellex and Massachusetts Institute of Technology. Dept. of Architecture. Program in Media Arts and Sciences. *Natural language and spatial reasoning*. PhD thesis, 2010.

[31] M. Theune, S. Faas, and D. Heylen. The virtual storyteller: Story creation by intelligent agents. *Proceedings of the Technologies for Interactive Digital Storytelling and Entertainment TIDSE Conference (2003)*, 2003.

[32] M. Thiebaux, M. Rey, A. N. A. N. Marshall, S. Marsella, and M. Kallmann. SmartBody : Behavior Realization for Embodied Conversational Agents. *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 179(AAMAS):151–158, 2008.

[33] K. R. Thorisson and H. H. Vilhjalmsson. Functional Description of Multimodal Acts: A Proposal. *AAMAS 2009 Workshop Towards a Standard Markup Language for Embodied Dialogue Acts*, pages 1–4, Mar. 2009.

[34] H. Vilhjalmsson, N. Cantelmo, J. Cassell, N. E. Chafai, M. Kipp, S. Kopp, M. Mancini, S. Marsella, A. Marshall, C. Pelachaud, and Others. The Behavior Markup Language: Recent Developments and Challenges. *Intelligent Virtual Agents*, pages 99–111, 2007.