

An Active Analysis and Crowd Sourced Approach to Social Training

Dan Feng¹, Elin Carstensdottir¹, Sharon Marie Carnicke², Magy Seif El-Nasr¹,
and Stacy Marsella¹

¹ Northeastern University, Boston, MA 02115, USA,
{danfeng, elin, magy, marsella}@ccs.neu.edu

² University of Southern California, Los Angeles, CA 90089, USA,
carnicke@usc.edu

Abstract. Interactive narrative (IN) has increasingly been used for social skill training. However, extensive content creation is needed to provide learners with flexibility to replay scenarios with sufficient variety to achieve proficiency. Such flexibility requires considerable content creation appropriate for social skills training. The goal of our work is to address these issues through developing a generative narrative approach that reconceptualizes social training IN as an improvisation using Stanislavsky’s Active Analysis (AA), and utilizes AA to create a crowd sourcing content creation method. AA is a director guided rehearsal technique that promotes Theory of Mind skills critical to social interaction and decomposes a script into key events. In this paper, we discuss AA and the iterative crowd sourcing approach we developed to generate rich, coherent content that can be used to develop a generative model for interactive narrative.

Keywords: Intelligent Narrative Technologies, Active Analysis, Theory of Mind, Crowd Sourcing, Social Skills Training

1 Introduction

Effective social interaction is a critical human skill. To train social skills, there has been a rapid growth in narrative-based simulations that allow learners to role-play social interactions with virtual characters, and, in the process, ideally learn social skills necessary to deal with socially complex situations in real-life. Examples include training systems to address doctor-patient interaction [1], cross-cultural interaction within the military [2], and childhood bullying [3].

Although interactive narrative systems represent a major advancement in training, their designs often do not provide learners with flexibility to replay scenarios with sufficient variety in choices to support learning and achievement of proficiency. Attempts to create more generative experiences face various challenges. First, the combinatorial explosion of alternative narrative paths poses an overwhelming challenge to create content for all the paths, especially if there are long narrative arcs. Second, current narrative systems are often brittle and constrained in terms of their generativity and flexibility to adapt to the interaction.

Additionally, training systems are often designed around exercising specific skills in specific situations. However, it is also important for social skill training to teach skills more broadly. Fundamental to effective human social interaction

is the human skill to have and use beliefs about the mental processes and states of others, commonly called Theory of Mind (ToM) [4]. ToM skills are predictive of social cooperation [5] and collective intelligence [6] as well as being implicated in a range of other social interaction constructs, e.g., cognitive empathy [7] and shared mental models [8]. Although children develop ToM at an early age, adults often fail to employ it [9]. On the other hand, people engaging in ToM across multiple situations, including actors, have improved ToM skills [10].

Our goal is to develop generative social training narrative systems that support replay as well as embed ToM training. Achieving this requires: (1) a new generative model for conceptualizing narrative/role-play experiences, (2) new methods to facilitate extensive content creation for those experiences and (3) an approach that embeds ToM training in the experience to support better learning outcomes.

Our approach begins with a paradigm shift that re-conceptualizes social skill simulation as rehearsing and improvising roles instead of performing a role. We have adapted Stanislavsky’s Active Analysis (AA) rehearsal technique [11] as the design basis for social simulation training. AA was developed to help theater actors rehearse a script or text. The overall script is divided into key events (i.e., short scenes) that actors rehearse and improvise under a director’s guidance. AA has two attributes especially relevant to social skills training. First, AA is designed to foster an actor’s conceptualization of the beliefs, motivations and behavior of their own as well as other actors, and thus is developed to engender ToM reasoning. Second, by adopting AA to simulation based social skills training, the emphasis shifts to developing short scenes that allow variability and re-playability. Decomposition into short rehearsal scenes *helps*: (a) break the combinatorial explosion that exacerbates content creation for long narrative arcs, (b) support users replaying scenes, possibly in different roles with different virtual actors, and (c) users to directly experience in subsequent scenes the larger social consequences of behaviors.

AA provides a basis for experience design that uses ToM constructs and enables crowd sourcing to generate content for rich, coherent interactive experiences. Several researchers have proposed crowd sourcing techniques for narrative creation (e.g., [12, 13]). The work presented here differs in that we focus on an iterative crowd sourcing method designed to create content for crafting a space of rich social interactions in which players explore a wide range of social gambits, from ethical persuasion and personal appeals to even deception; the content is created through the crowd using carefully designed tasks and interfaces that use AA and ToM as theoretical foundation.

In this paper, we present our approach in detail. Specifically, we will start by discussing the theoretical foundation for our work: Active Analysis. We then outline our overall approach to interactive narrative system design. Following this, we discuss the 3-step iterative crowd sourcing approach outlining the method, as well as the results for each step of the process. We then discuss overall results of the current work and projections for future work. Following this discussion, we outline previous related work and then conclude the paper.

2 Active Analysis Primer

AA serves several purposes in our project. It defines the overall structure of the learning experience as rehearsals of key scenes. Also, it helps define ways to manipulate the challenges the learner faces in those scenes as well as the feedback the learner gets. Most critically, AA provides a novel and potentially very powerful tool for ToM and more generally social skills training. Through AA, the human role-player/learner is called upon to conceive their behavior in terms of how other participants as well as observers understand and react to it. Further, they may have flawed mental models of others. Additionally, re-conceptualizing interactive narrative as performance rehearsal brings in elements of game play, such as repeat play as a form of rehearsal/mastery.

Stanislavsky developed AA in the latter part of his career to transform actor training and performance rehearsal techniques by emphasizing the importance of cognitive-social-affective mental states. As prelude to actors engaging in active analysis, the script is broken into small key scenes that the actors improvise and analyze. For our purposes in automating AA, this breakdown of a script would be part of the design, and therefore a larger training scenario is assumed to be broken into these brief key events/scenes. Consider the following simple multi-scene scenario that is currently the basis of our work:

Scenario: 3 Scenes from Mistaken Guilt on a Train

1. Scene 1: A novice reporter (*R*) witnesses a passenger attacking other passengers and a black good Samaritan tackles the rampaging passenger to protect the other passengers. A policeman enters and arrests the good Samaritan.
2. Scene 2: *R* goes to interview police chief (*PC*) at his house but a guard (*G*) at the house may have orders to block reporters.
3. Scene 3: *R* meets *PC* to persuade him to let the good Samaritan go.

AA of a scene consists of three phases: Framing, Improvisation and Performance Analysis. These phases are repeated, the rehearsal director often changing the motivations and tactics of the actors as well as the roles they play.

Framing: The actors and director determine what the overall context of the scene is and what event (or goals) should occur. This guides the actors as they improvise. Some actors will play ‘impelling actions’ which move purposefully toward the event. Others play ‘counter-actions’ which resist, delay or block the targeted event/goal. Using Scene 2 from the above scenario, impelling actions for *R* might be to explain or justify her purpose to meet *PC*, to elicit empathy from *G* or even threaten or attack *G*. The rehearsal director helps the actors explore different motivations and approaches to their role. In particular, he/she can selectively provide information about goals and actions to an actor without the other actors knowing that information. Thus, an actor may have a flawed mental model of the other actors leading into the improvisation. This mental model manipulation provides a powerful tool in social skills training.

Improvisation: During improvisations, actors explore different tactics to achieve their goals. Sometimes this improvisation fails without the goal being achieved and the actors not knowing how to proceed. This also can be a key

point for the design of social training simulations, as sometimes the interaction may just simply break down and a different approach must be taken.

Performance Analysis: After completing the improvisation, the director and actors evaluate the work at three levels of analysis: (1) the individual, the actor thinking of their own actions, (2) the social, how those actions are countered by other actors, and (3) the audience, how the actions expressed information to others observing the performance. The actors then revise their preliminary analysis and repeat the improvisation. This process of framing, improvisation and analysis is repeated until the group has achieved what they wish.

3 Overview of the AA Social Skill Training System




Fig. 1. The AA Social Skill Training System

Figure 1 presents the current overall design of our system. In line with AA rehearsals, we assume the overall scene is broken into highly variable interaction scenes that constitute key events. These scenes will be performed by the human and virtual actors through the director agent’s guidance. To represent this structure, we use a high level state graph adapted from StoryNet [14]. The nodes in each scene represent the AA’s key events/scenes. They are free-play areas in which the human participant can rehearse their role with virtual characters, exploring different approaches to their role within a scene. The edges in the graph connect the scenes. Traversal along an edge constitutes a completion of the rehearsal of a scene and a move onto a subsequent scene to rehearse. This traversal is dependent on whether the rehearsal has achieved its goal or some minimal number of improv sessions have occurred. *Framing* sets up the learner’s role in the scene. *Performance Analysis*, as in AA, provides feedback on the characters’ beliefs and motivations as well as cues on how to achieve scene goals. This feedback would leverage knowledge acquired during the crowd sourcing process.

In the following sections we discuss our progress in automating the collection of AA structured content within scenes, using crowd sourcing techniques. Other parts of the architecture, including the director agent, are left to future work.

4 Within Scene Content

To craft scenes, we implemented a multi-step crowd method, as seen in Fig. 2, that intentionally separates the creative, analytical and assessment processes. During the creative step, crowd workers generate a scene based on an initial

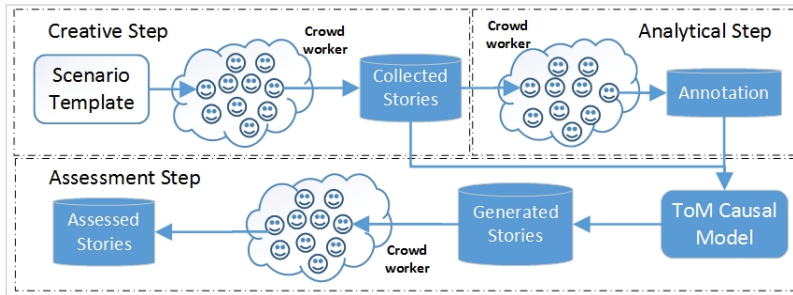


Fig. 2. A Iterative Crowd Sourcing Method Based on AA and ToM

situation, goals and beliefs of the characters, resulting in a richly varied and creative set of alternative scenes. In the analytical step, a separate set of crowd workers are tasked to annotate the scenes generated from the creative step using a novel ToM inspired annotation scheme that identifies how the actions seek to alter the beliefs, goals and attitudes of the characters. These annotations are critical to the pedagogy and also help constrain how elements across stories can be intermixed to provide variation in both learner’s action choices as well as the non-player actions. In the assessment step, the results from these efforts provide a set of distinct actions that are then recombined using constraints generated from the annotation step and fed to a third set of crowd workers to further evaluate coherence. The result is a large space of coherent stories that support rich, meaningful variations in both learner’s and non-player characters’ choices and thus allowing for rehearsal and replay. These stories can then be used as a base for the interactive narrative system described above. Details on these steps are discussed below.

4.1 Creative Step: Crowd Sourcing Scenes

The task for the scene creation is to create very different variants of how a particular scene plays out from which a space of possible actions can be isolated and recombined to create a rich interactive experience. We designed and deployed multiple interfaces. Due to space constraints we only discuss the most promising design, the AA Interface, which was based on AA techniques for eliciting variations from actors improvising a scene. The AA interface seeks to limit the worker to craft relatively simple sentences where a sentence describes one action that a character performs (in line with [15]). The same scenario description was used for both tasks: *R, an aspiring reporter, wants to interview PC at his house. PC employs G to make sure that nobody is able to bother him. G guards the door. R is in front of PC’s house and is planning to attempt to interview him.*

The design of the AA interface is segmented into three steps. This is to mimic the repeated rehearsal element of AA. In every step the worker is asked to write a new story (one sentence in a line as shown in Fig. 3), where instructions differ slightly from previous steps. Step 1 offers a basic description of the

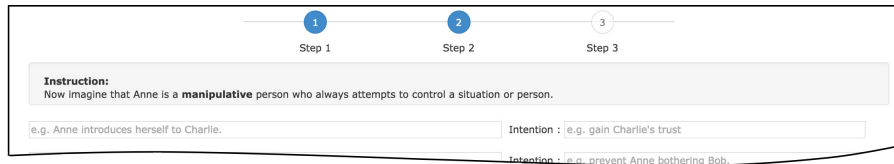


Fig. 3. A section of the AA Interface. Each worker has 6 lines presented to them initially as a default but is allowed to add as many additional lines as they see fit.

scene. Steps 2 and 3 give workers one piece of new information respectively, in addition to the scenario description. This information relates to one of the character’s motivations, relationships or traits. This design is to encourage the worker to re-evaluate the same scenario but from different perspectives, much like AA rehearsal directors do with their actors. Such process will ideally a) spur the creativity of the worker by mimicking the AA rehearsal experience and b) guide the worker to provide the varied content needed to support replay and pedagogical goals in the training system built from this content.

Story Collection Results: We collected 108 stories using the same AA interface. Most stories consisted of adjacency and act/react pairs, where R would take an action and G would then respond. The average length of a story was 6 lines.

The stories are very rich in terms of character actions and the intention of each action. This applies to R in particular, while G ’s behaviors varied in terms of responses to actions. The intentions for the actions however, remain relatively static throughout for most of the stories until the conclusion is reached where G either changes his goal of blocking R or successfully blocks R . The stories collected show an impressive range in complexity, richness and variety in tactics that R employs, resulting in a sizable action set for our characters. To give a better sense of the variety of actions collected, we provide two examples, a relatively simple bribe story and a more complex manipulation story:

Bribe: R asks G to see PC . G declines her request. R tries to bribe G . G declines R ’s bribe. R adds \$50.00 more to the bribe. G accepts the bribe.

Manipulation: R flirts with G . G tries to ignore R . R compliments G a lot. G begins to flirt with R . R tells G that she just needs a few teeny minutes with PC . G becomes wary and tells her to leave.

For each sentence, the intention was described by the worker as previously mentioned. It’s worthwhile to note that despite some of the stories having similar or even identical actions, the actions in question can have very different intentions assigned to them. One example of this, ‘ R throws a stone (a pebble) to the window’ in one case intended to get PC ’s attention directly. In another, the same action has the intention to distract G so R can sneak into the house. This was in line with our original goal of not only gathering enough information to map the potential action space of the characters but also the variety of the possible intentions that motivated the actions the characters employed. Reaching this level of richness in the collected stories is key in the context of our work,

Interaction	Attempt	Was the attempt successful?
Anne dresses up as a postwoman.	Anne wants to change Charlie's belief about Anne is likely a reporter	No
Charlie replies that Bob is not home.	Charlie wants to change Anne's belief about Bob is in the house	No
Anne informs Charlie that Bob must sign to acknowledge delivery.	Anne wants to change Charlie's belief about Anne is likely a reporter	No
Charlie informs Anne that this is simply not possible.	Charlie wants to change Anne's goal to meet with Bob	No
Anne begins to cry and says her job is on the line.	Anne wants to change Charlie's goal to block reporters from entering the house	yes
Charlie finally relents and let her in.	Charlie wants to change his own goal to <input checked="" type="checkbox"/> block reporters from entering the house please Bob	yes

Fig. 4. The ToM annotation interface. Left column shows actions. Middle column is the ToM annotation choices. Right column allows the crowd to report whether the intended effect was successful. Workers selected the description from a drop-down list which is generated based on the intention of each action collected in the creative step. The list changes dynamically depending on other choices.

as our goal is to collect as many actions and intentions as possible, to allow our players as much variety as possible in exploring social situations from a ToM standpoint. At a more detailed level, we found that the AA interface could successfully influence the type of stories being generated. For instance, as shown in Fig. 3, by framing R as ‘manipulative’, we collected actions like ‘tell G his mother is in the hospital’. By framing R as G ’s ‘ex-girlfriend’, we got actions like ‘ask G to get back together’.

4.2 Analytical Step: Crowd ToM Annotation

As noted, ToM annotations are critical for: the pedagogy provided during Performance Analysis and constraining how actions are intermixed to create variation for the learners’ experience in the interactive narrative. To get this information, we conducted a second crowd task using the annotation interface seen in Fig. 4.

The challenge in employing non-expert workers for this annotation task is to help them correctly parse character actions into intention descriptions in ToM terms. To set up the initial goals and beliefs for each character as a director would in an AA rehearsal, we provided workers with an initial list of goals and beliefs for each character as shown in Table 1. Instructions to workers also included the scenario description. Each sentence in a story represents one action. For each sentence the worker is asked to stipulate how the character taking the action intends to change another character’s goal, belief, liking or trust, the subject of the change and, whether the attempt is successful. Since each action may have multiple intended effects, we combined the output of multiple workers to get the multiple intentions, instead of requiring each worker to provide all the intentions.

Annotation Result and Discussion: We randomly selected 3 stories to annotate from the story corpus obtained through the first crowd stage, described in the previous section. Data from 120 workers was collected; each story was annotated by 40 workers. For each action, the 2 most agreed upon combinations were selected as that action’s potential effects.

Table 1. Initial goals and beliefs given to the workers

Character	Object	Content
The Reporter (<i>R</i>)	goal	meet with the police chief (<i>PC</i>)
	belief	<i>PC</i> is in the house <i>G</i> is likely to block her
The Guard (<i>G</i>)	goal	block <i>R</i> from entering the house please <i>PC</i>
	belief	<i>PC</i> is in the house <i>PC</i> doesn't want to meet any reporters <i>R</i> is likely a reporter

Examples of the annotation results are shown in Table 2. It lists the sentence/action, the two most popular annotations of the intention of the action described in the sentence and the corresponding ratio of workers who chose that annotation option. Possible choices for one annotation can be as many as 20 (3 characters, 1-2 goals, beliefs, liking and trust for each character). Inter-rater reliability was determined with Chance-corrected agreement (Krippendorff's alpha) for all annotated stories. The average Krippendorff's alpha value across all stories is 0.887 with a standard derivation of 0.01. This relatively high agreement implies the design of the annotation scheme is not overwhelming and easy for non-technical personal to understand.

4.3 Assessment Step: ToM Causal Model

To develop an interactive system with a rich and generative action space, we need to reorganize and recombine collected actions but maintain the coherence of the story. As discussed above, we decided to use the crowd to assess stories developed through the crowd sourcing process. However, the combinatorial explosion caused by random recombination makes it a daunting task for workers to assess the result. For instance, in the 3 stories we annotated, there are just 20 sentences including 10 actions for each character. If we randomly combine 6 sentences to generate a new story, the number of possible combinations is $10 * 10 * 9 * 9 * 8 * 8 = 518400$. Most of these narratives are nonsensical or incoherent. Thus in order to eliminate them, the crowd provided ToM annotation data was used to create a causal model, namely the ToM Causal Model (TCM),

Table 2. Examples of the annotation result

Sentence	Intentions (wants to change)	%
<i>R</i> dresses up like a postman.	<i>G</i> 's belief about she is likely a reporter	55
	<i>G</i> 's goal to block her from entering the house	25
<i>G</i> accepts the bribe.	<i>G</i> 's own goal to block <i>R</i> from entering the house	43
	<i>G</i> 's own goal to please <i>PC</i>	28
<i>R</i> tells <i>G</i> the interview has already been scheduled.	<i>G</i> 's belief about <i>PC</i> does not want see reporter.	60
	<i>G</i> 's goal to block <i>R</i>	30

that establishes a precedence relation between sentences. The basic assumption for the TCM is that the intention of an action selected by a character should be consistent with the intention of previous actions in the story, which can be used to determine what actions follow a particular prefix of actions in a generated story.

The TCM is based on the idea of an intention being in play. An intention is in play at some point in a seed story, if in that story a preceding action has that intended effect. For example, from the annotation data we know the intended effects of ‘bribe’ are ‘changing G ’s goal of blocking her’ ($I1$) and ‘increasing G ’s liking of her’ ($I2$). Thus, after ‘bribe’ happens, $I1$ and $I2$ are defined as ‘in-play’.

From this information, we can construct an abstract model of actions that describe their in play intentions as well as intended effects, as a heuristic model of preconditions and effects. Specifically, an action is defined as a tuple $A = \langle PI, EI, \Phi \rangle$ where PI is a set of ‘in-play’ intentions that existed when action A happened in the seed story and EI is the set of intended effects by taking this action. Φ defines the phase where the action occurs: beginning, middle or end.

When generating new stories, an action can only be selected if all the intentions in its PI are in play. For example, in the ‘bribe’ story, R bribes G . After the bribe is declined, she bribes \$50 more. Thus, ‘bribe-more’ can only happen when intentions $I1$ and $I2$ are in play. Given our TCM model, these constraints: intentions $I1$ and $I2$, are set as PI for the action ‘bribe-more’. After an action is selected, its intended effects EI will added to the set of ‘in-play’ intentions. Additionally, we constrained the action selection based on the phase Φ . We assume the initial/ending actions in the seed stories like ‘greet’, ‘introduce herself’, ‘let R in’ or ‘arrest R ’ are more likely to happen at the beginning/end of the story. All intermediate actions can happen at any place in between.

The TCM Result and Discussion: Obviously, the TCM is not a perfect coherency filter, in particular it over prunes. Still, pruning makes it feasible to further test coherence using the crowd. So, we asked another group of workers to evaluate the results using 5-point Likert items adapted from [16] to assess stories’ coherence and consistency: I understand the story; The story makes sense; Every sentence in the story fits into the overall plot; The characters behaviors are consistent with their goals and beliefs; The characters’ interaction is believable.

We collected evaluations from 40 workers for each story. We mapped the 5-point Likert items to a score ranging from $[-2, 2]$. We eliminated all stories that have at least one non-positive score leaving us with 12 stories³. To validate the workers’ assessment, we asked an expert who is familiar with ToM and interactive narrative to evaluate all generated stories. The crowd result was generally more selective than the expert’s. We found one story with an illicit action pair (decline bribe before bribe is offered) excluded by the expert but not by the crowd. This suggests we may need to consider alternative criteria. For example, eliminating stories based on total score across 5 questions less than a threshold $T_e = 2.5$ leads to better agreement with the expert.

³ All generated stories can be found at <http://www.northeastern.edu/cesar/project/IxN>

5 Discussion and Future Directions

Our results show the promise of adapting AA and crowd sourcing to generate rich and varied stories from a simple initial scene. The results indicate the utility of designing crowd tasks using AA principles to inspire creative output. In particular, we found directorial guidance from the AA interface can directly influence the stories. For example, the AA interface successfully primed the workers to focus on manipulative and deceptive behavior. This suggests that going forward we will be able to influence the worker to get the needed pedagogical content.

We also found in the ToM annotation task, the crowd was able to annotate these stories with sophisticated ToM concepts and with high agreement. We showed how these annotations can provide critical pedagogical content but we also demonstrate how they can be incorporated into the TCM to help assess generated stories. This led to a significant reduction in the number of stories and enabled us to successfully go to the crowd to do a final validation check.

While giving us significant amount of ToM information, the ToM annotation scheme is still simple. Adding more layers in subsequent schemes might get to more detailed information such as tactics. The high agreement suggests the crowd might be capable of performing more detailed annotation. Using finer-grained annotation schemes would help categorize content to a greater degree, benefiting both the design of the pedagogy and the performance of the TCM.

In addition, we can improve the TCM in other ways to ensure its scalability. With the assessment result, we can construct a stronger TCM by abstracting the conditional relation between 2 actions based on whether they 1) have high mutual information in good stories and 2) are only found in reverse order in eliminated stories. With these 2 criteria, we can derive 2 conditional constraints, ‘bribe more’ and ‘decline bribe’, that only happen if ‘bribe’ already occurred. These causal relations, combined with TCM can be used to filter out incoherent stories. As an alternative to using the TCM, we will also explore machine learning techniques to learn a probabilistic model that can evaluate the generated stories.

To date, we implemented a prototype, using the crowd results reported here, that allows a player to play the role of R or G . This prototype does not include the AA director agent however, depicted in Fig. 1. To that end, our next step is to collect data from the rehearsal practice of leading expert on AA, which will be used to inform the director agent design.

6 Previous Work

Successful examples of systems using crowdsourced narratives to generate new stories include SayAnything [17], the Scheherazade System [18] and, its interactive counterpart [16]. Crowdsourcing has also been used to model behavior such as in the Restaurant Game [19]. A key challenge in our work lies in populating the action space with complex social actions, especially since a sufficient number can mean thousands of actions. For example, the social game PromWeek [20] utilized over 5000 human-authored social rules, a complex and time consuming

task and, [19] uses a set of known actions. Our approach has actions extracted from crowdsourced narratives but additionally, uses those actions to generate new stories to explore the possible causal relationships between them.

Annotating potentially large narrative corpora, like those generated in our work is also a non-trivial task. The Scheherazade annotation tool [21] utilizes Story Intention Graphs (SIG) which capture ToM concepts. However, SIG requires expertise. Annotation frameworks often have a detailed structure requiring expertise to achieve adequate levels of inter-rater reliability, giving rise to much simpler annotation schemes, such as the two compared in [22]. Neither of those two however focuses explicitly on annotating detailed social interactions nor ToM reasoning. Any annotation scheme we use needs to reflect our ToM focus, yet be simple enough not to require training or expertise of the worker.

[23] also uses a multi-step iterative crowd sourcing approach to content creation, similar to the work presented here. However, their work focuses on activity oriented non-interactive narratives and actions, whereas we focus on interactive narratives involving social interaction and ToM reasoning.

In terms of focus on ToM, our work is similar to the interactive drama system Thespian[24]. However, Thespian employs ToM reasoning to generate behavior by modeling the player and characters as autonomous agents, while we use ToM to describe and recompose character behavior without an explicit agent model.

7 Conclusion

In this paper, we introduced an iterative crowd sourcing method for interactive narrative, based on AA and ToM. Designed to realize a space of rich social interactions, it will allow players to explore a wide range of social gambits, from ethical persuasion and personal appeals to deception. The result of using AA and ToM as theoretical foundations show the promise of using such a framework to collect, annotate and generate interactive narratives for social skills training. Going forward, we plan to use a data-driven technique to create a director agent that will incorporate more aspects of AA into the experience.

Acknowledgment: Funding for this research was provided by the National Science Foundation Cyber-Human Systems under Grant No. 1526275.

References

1. Lok, B., Ferdig, R.E., Rajj, A., Johnsen, K., Dickerson, R., Coutts, J., Stevens, A., Lind, D.S.: Applying virtual reality in medical communication education: current findings and potential teaching and learning benefits of immersive virtual patients. *Virtual Reality* **10**(3-4) (2006) 185–195
2. Kim, J.M., Hill, J., Durlach, P.J., Lane, H.C., Forbell, E., Core, M., Marsella, S., Pynadath, D., Hart, J.: BiLAT: A Game-Based Environment for Practicing Negotiation in a Cultural Context. *IJAIED* **19**(3) (2009) 289–308
3. Zoll, C., Enz, S., Schaub, H., Aylett, R., Paiva, A.: Fighting bullying with the help of autonomous agents in a virtual school environment. In: 7th Int. Conf. on Cognitive Modelling. (2006)

4. Whiten, A.: Natural theories of mind: Evolution, development and simulation of everyday mindreading. Basil Blackwell Oxford (1991)
5. Artinger, F., Exadaktylos, F., Koppel, H., Sääksvuori, L.: In others' shoes: do individual differences in empathy and theory of mind shape social preferences? *PloS one* **9**(4) (2014) e92844
6. Engel, D., Woolley, A.W., Jing, L.X., Chabris, C.F., Malone, T.W.: Reading the mind in the eyes or reading between the lines? theory of mind predicts collective intelligence equally well online and face-to-face. *PloS one* **9**(12) (2014) e115212
7. Davis, M.H.: Measuring individual differences in empathy: Evidence for a multi-dimensional approach. *Journal of personality and social psychology* **44**(1) (1983)
8. Converse, S.: Shared mental models in expert team decision making. *Individual and group decision making: Current* (1993)
9. Lin, S., Keysar, B., Epley, N.: Reflexively mindblind: Using theory of mind to interpret behavior requires effortful attention. *Journal of Experimental Social Psychology* **46**(3) (2010) 551–556
10. Goldstein, T.R., Wu, K., Winner, E.: Actors are skilled in theory of mind but not empathy. *Imagination, Cognition and Personality* **29**(2) (2009) 115–133
11. Carnicke, S.M.: *Stanislavsky in Focus: An Acting Master for the Twenty-first Century*. Taylor & Francis (2009)
12. Li, B., Lee-Urban, S., Appling, D.S., Riedl, M.O.: Crowdsourcing narrative intelligence. *Advances in Cognitive Systems* **2** (2012) 25–42
13. Orkin, J., Roy, D.K.: Understanding Speech in Interactive Narratives with Crowd Sourced Data. In: *Proc. of the 8th AAAI Conf. on Artificial Intell. and Interactive Digital Entertainment*, The AAAI Press (2012)
14. Swartout, W., et. al: Toward the holodeck: integrating graphics, sound, character and story. In: *Proc. of 5th Int. Conf. on Autonomous Agents*, ACM (2001) 409–416
15. Li, B., Appling, D.S., Lee-Urban, S., Riedl, M.O.: Learning sociocultural knowledge via crowdsourced examples. In: *Proc. of the 4th AAAI Workshop on Human Computation*. (2012)
16. Guzdial, M., Harrison, B., Li, B., Riedl, M.O.: Crowdsourcing Open Interactive Narrative. In: *The 10th Int. Conf. on the Foundations of Digital Games*. (2015)
17. Swanson, R., Gordon, A.S.: Say Anything: A Massively Collaborative Open Domain Story Writing Companion. In Spierling, U., Szilas, N., eds.: *Interactive Storytelling*. Number 5334 in LNCS. Springer Berlin Heidelberg (2008) 32–40
18. Li, B., Lee-Urban, S., Johnston, G., Riedl, M.: Story Generation with Crowdsourced Plot Graphs. In: *Proc. of the 27th AAAI Conf. on Artificial Intell.* (2013)
19. Orkin, J., Roy, D.: The Restaurant Game: Learning Social Behavior and Language From Thousands of Players Online. *Journal of Game Development* **3**(1) (2007)
20. McCoy, J., Treanor, M., Samuel, B., Reed, A.A., Mateas, M., Wardrip-Fruin, N.: Prom Week: Designing past the game/story dilemma. In: *FDG*. (2013) 94–101
21. Elson, D.: Modeling narrative discourse. PhD thesis, Columbia University (2012)
22. Rahimtoroghi, E., Corcoran, T., Swanson, R., Walker, M.A., Sagae, K., Gordon, A.: Minimal Narrative Annotation Schemes and Their Applications. In: *7th Intell. Narrative Technologies Workshop*. (2014)
23. Sina, S., Rosenfeld, A., Kraus, S.: Generating Content for Scenario-Based Serious-Games Using CrowdSourcing. In: *AAAI*. (2014) 522–529
24. Si, M., Marsella, S.C., Pynadath, D.V.: Thespian: An Architecture for Interactive Pedagogical Drama. In: *Proc. of the 2005 conf. on Artificial Intell. in Ed.: Supporting Learning through Intell. and Socially Informed Technology*. (2005) 595–602