Ontology of Robot Theatre

David V. Lu

***Abstract*— There are an increasing number of theatrical productions that involve mobile robots functioning as characters. The underlying systems that drive their performances can vary in numerous ways, including the level of adaptive control and the level of human involvement. However, oftentimes all of the performances as labeled equally as "robot acting," hiding numerous possible complexities and simplifications in the programming. This paper aims to categorize and classify the different algorithms that can exist behind the curtain to create these performances.**

I. Context

The first usage of the word *robot* does not come from a technical paper, but a work of drama, the play "R.U.R. (Rossum's Universal Robots)" [1]. However, mechanical devices have been integrated into theatre starting long before Čapek's play. The term "deus ex machina", referring to a mysterious outside force that enters the play and cleanly wraps up the plots, stems from the Ancient Greek practice of lowering a mechanical god onto stage as a plot device. And in recent years, with the advent of cheaper and more commonplace robotic technologies, placing robots into plays has become increasingly common.

Any time that a piece of interesting technology is placed on stage, a fundamental question arises as to what role the technology plays. In many cases, it is often just a prop, much like a skull, spoon or other object that an actor might interact with on stage. However, once the technology starts moving on the stage, apparently under its own volition, it starts to become something else. It moves closer to the role that the humans on stage have: actor.

It is tempting to say that any such robots are acting much like their human counterparts. This is especially true when the robots are humanoid or recite lines of dialogue. However, making the claim that they are truly acting is not that easy. There is a large gray area in between things that are known not to be acting and people who are undeniably acting. On one side, there could be an immobile robot that occasionally beeps placed on stage, playing the part of "inert robot" in a scientist's lab. On the other side of the spectrum, one could envision a future robot which can read the script for "Hamlet," make decisions about how it wanted to perform its assigned role, and then perform on stage with other actors, altering its own performance to react in real time to the others' performances. Between these two extremes exists a continuum of robots that could be considered "acting robots" bu are clearly not doing the same things.

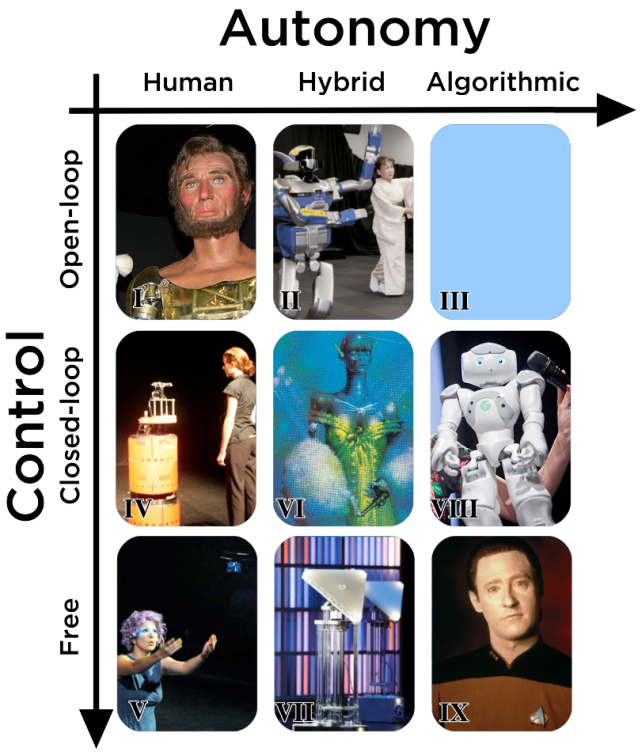

Fig. 1. Examples in The Ontology. (I) Audio-Animatronic Lincoln at Disneyland (II) Dancing HRP-4 from AIST (III) None (IV) B-21r performing at Washington University (V) Quadcopter doing Shakespeare at Texas A&M (VI) Cover for the short story "The Darfstellar" (VII) Operabot from "Death and the Powers" (VIII) Data (Nao) performing comedy at TED (IX) Lt. Data from Star Trek. All images copyright their respective owners.

Where should the line be drawn between acting and not acting? Who gets artistic credit for robot performance? In order to answer these questions, it would be useful to have a common framework with which to compare the different types of performance. In this paper, I propose an ontology for classifying the different categories of robot performances and explore where some recent robots fit into the proposed classification system.

II. "Robot" "Theatre"

Before enumerating the different kinds of robot theatre, it is helpful to have definitions for what is meant by "robot" and what is meant by "theatre." Most people have their own baseline definition for these concepts. However, there is a need to be fairly precise in order to discuss the interesting edge cases that will be presented.

This work was partially supported by the National Science Foundation under award IIS-0917199.

David V. Lu is with the Department of Computer Science and Engineering at Washington University in St. Louis, St. Louis, MO 63130, USA. davidlu@wustl.edu

Fig. 2. Scene from "I Worker" with Wakamaru robot (right). Source: http://www.seinendan.org/en/special/robot01/RHT-Forest.html

One definition for a robot is a mechanical device with moving parts that is able to make changes in its environment. Some might include sensors or a computer "brain" as essential parts. However, in this analysis, those aspects will not be included since requiring them would eliminate some key types of robots, including certain sorts of teleoperated robots.

There is one one further important quality for distinguishing robots from other mechanical devices that can appear on stage, namely agency. For example, in the play "I Worker" [2], there are two vaguely humanoid Wakamaru robots which function as characters in the piece, moving about stage and interacting with the human characters, as seen in Figure 2. There is also a tape player on stage with which the actors interact. Common sense would likely say that the tape player is not a robot. However, on some level it is the same type of object: it is non-living and mechanical, it responds to human commands, it plays sound when needed. The key difference is that the Wakamaru machines seem agentic. That is, they are under their own control, able to initiate different behaviors and decide to act on their environment, much in the same way human actors do.

As for the definition of theatre, that depends on how we define acting. In previous work [3], we've discussed how an essential quality to acting is that actors can never truly become their characters. Instead they do an action in a way to make the audience think that they are the character. The difference between truly doing something and acting *as if* you were doing that something is in the relationship to the audience. When an actor acts, there is an extra piece of information that they are attempting to convey to the audience. It is not merely that the actor has picked up a skull; they want to convey to the audience the idea that the character Hamlet has picked up the skull. Similarly, if a robot moves from point A to point B, that is just the robot doing an action. However, if it moves from A to B in order to represent to an audience an additional belief about the robot, then it is acting. This means that many robot activities which do not resemble traditional theatre settings can involve robot acting. However, the same holds true for human actors, so the definition remains valid.

III. The Ontology

A. The Axes

The categories in this ontology are differentiated along two different axes. First, there is the automation level, which describes how much of the robot's actions are prescribed by a human, and how much of it is derived by some algorithm running on the robot. The second variable is the control of the system, exploring how reactive the robot is to the immediate situation on stage.

1) Automation: Examining the automation level of a robot actor allows us to differentiate between the complexities of the algorithms needed to run the robot. It allows us to decide who is actually acting. On the one side, there are systems where all of the robot's actions are directly *human-produced*, either through teleoperation of the robot, specifying exactly how the robot should move in code (i.e. hardcoding the motions), or using a human's performance to directly generate the robot's performance. On the other side of the spectrum, there are *algorithmic* systems which generate their behavior via computation without explicit human input. Finally, there is the space in between for *hybrid* systems, which have behavior partially specified by the human which can then be modified by an algorithm to achieve some additional behavior, or vice versa, with the human over-riding some algorithmically specified behavior.

For purposes of discussion here, we are ignoring the question of who wrote the script that the actors build their performance around. Barring further advances in generative writing algorithms, there will be some high-level structure to the play that is specified by a human. In human theatre, the writing and the acting are two distinct creative processes. The writer creates the structure for a play, and the actors create behaviors which fill that structure.

The automation level is important because it helps locate the intelligence of the system. If the robot performs a particularly effective action within the context of the play, the action is much more impressive if its computer generated as opposed to if its human generated. Put another way, this variable is a way of attributing authorship to the actions. An algorithmic system that is generating its own actions is more like the robot itself actually acting, whereas in the human-produced case, it is more akin to having the human acting through the mechanical device.

Note that this automation level differs from Sheridan's autonomy scale [4]. Here autonomy is defined by the autonomy of the generation of actions, either requiring human intervention or not. The separate issue of how the robot reacts to its sensor data and chooses which actions to do is covered by the other axis, control.

2) Control: For humans, one of the keys to effective acting is a constant awareness of what others are doing in the scene around them and being able to adjust their own performance to better mesh with them. Failure to do so may yield the unfairly derogatory label of "robotic actor." The more reactive and dynamic a performance becomes, the higher the quality of the performance. The same holds true

for robots.

We divide up this variable space into three categories as well. The first is a *open-loop* performance in which the robot does the same exact thing in every performance regardless of external factors. Second, there is the *closed-loop* performance, which is generally the same performance every time but has a feedback mechanism which allows the robot to react to the given circumstances. Finally, the *free* performances are those in which the system is constructed so that the robot can do nearly anything. The free control systems are capable of inventing or generating new responses to their circumstances. Human actors in live theatre generally fall into this type of performance because they are respond to the other actors and minute changes in timing differently every performance. Robots, as we see in the following section, can fall into any of the three categories.

B. Classifications

We now can examine the different combinations of these variables. As you can see in Figure 1, we divide the area up into nine separate regions, which correspond to to nine different classes of robot actor. The regions are numbered in one possible ordering corresponding to increasing difficulty of the task. The nine classes are further grouped into four larger categories, as seen in Figure 4.

Category 1: Playback: This category contains the three classes that are completely open-loop. Some might argue that these are not quite actors (or robots for that matter) because they do not actually react to their environment at all. All of these classes boil down to playing back some previous set of actions for an audience.

Class I includes robot systems that are open-loop and human-produced. The prototypical example of Class I is most of the Audio-Animatronics at Walt Disney World. The animatronic pirates and presidents repeat the same exact actions time after time. They were also essentially hand animated by the Imagineers to perform the exact movements that they wanted. This level is the baseline because, while it can give the appearance of being interactive, it is just an illusion, since the responses will never change. This also means that this class can easily break the illusion of social interaction if the other actors do not follow the script, making the robot's canned responses seem out of place. Furthermore, there is no intelligence in such a system, as it is merely playing back precisely what the programmer told it to. In this way, it is more like the tape player than a true actor.

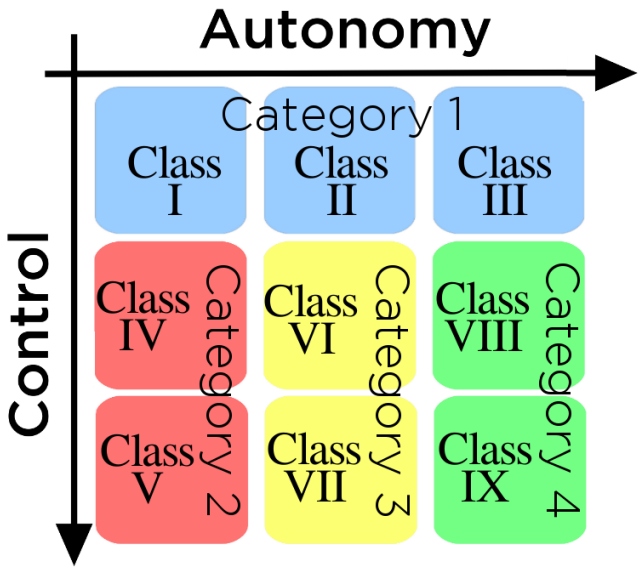

Fig. 3. Breakdown of categories and classes

Compare that with Class II, which is also open-loop, but now has hybrid control. This includes systems where the actions are largely dictated by humans, but some portion is derived by the computer. The dancers created at AIST fall into this category[5]. The robot dancer bases its movement on motion capture of a human dancer, but then modifies the motions in order to maintain stability and stay upright. This is a slightly more advanced algorithm than Class I, but still requires human input and will still result in the same performance each time.

Class III moves even further away from human input to actions that are computer/robot generated, creating a open-loop and algorithmic system. Such a system is not known to exist at this time. One might see this class of robot actor when an acting algorithm generates the behavior, but is too slow to run in real time (or multiple times), requiring that the algorithm be run once beforehand.

Category 2: Teleoperated: The next category of acting systems integrate human input in a much more direct way. By having a human doing the control during the actual performance (either in person or remotely), the performance becomes more interactive, providing a feedback mechanism and ergo becomes more closed-loop. Category 2 contains two classes (IV and V) that are both human-based, but have variable degrees of interactivity.

Class IV contains the closed-loop systems with human input, meaning that the performance can change based on some conditions on stage, but not in arbitrary ways, which would push the classification to Class V. The difference is best examined by looking at the number of different choices that are available for the human operator to choose from. Class V control systems allow the robot to perform anything that the human wants, giving unlimited possibilities. If there are only a finite number of different possible actions, even though the human can make the system react to other actors, they are still constrained to certain behaviors, making it a class IV system. This covers cases with structures like dialog trees, where there are a number of pre-programmed responses that the human operator can select between. This was the motivating use case for Polonius [6], a Wizard-Of-Oz platform that allowed the human operator to choose between a limited subset of possible outcomes and control the timing of the robot's actions. The system was used on the B-21r, Lewis, in a performance at Washington University in 2010 [7].

One notable example of a character that transitioned from

Class IV to Class V is Craig Ferguson's robot sidekick Geoff Peterson on "The Late Late Show."[8] Designed by Grant Imahara, the robot originally had a number of canned phrases that it could say, triggered by someone offstage. However, once the character caught on, the voice artist was kept in the studio to operate Geoff, allowing him to say anything he saw fit. There are also numerous other examples of teleoperated actors with the humans in complete control, including the production of Heddatron in New York [9] and the production of A Midsummer Night's Dream performed at Texas A&M with the assistance of teleoperated quadcopters [10].

It is worth noting that there are situations where both classes of teleoperated systems are appropriate. Class V requires a very talented operator who has the benefit of being able to react in the most appropriate way, but also has many poorer options that can be performed. Class IV limits the appropriateness of the response, but does not require control that is as finely grained.

Category 3: Collaborative: The third category contains all interactive hybrid systems. The performance combines elements of human performance with some degree of machine intelligence. The human performance can either be done beforehand (e.g. an algorithm that is based on motion captured human motion) or done live and processed in real time (e.g. augmented performance).

Class VI is defined for all closed-loop hybrid systems, where the performance combines human generated elements with algorithmically generated elements, although only a finite number of ways. The robots in Walter Miller's short story "The Darfstellar" [11] are a good example. The plot develops a future where human actors are obsoleted by autonomous acting systems which base their performances on the "personality matrices" of human actors recorded long ago. When a human actor inserts himself into the play, "Maestro", the main computer controlling the actors, is able to alter the performance around him and make changes based on the audience reaction; however, the system is limited in its responses to the new stimuli by changing line readings. However the system can't do something not in the script, like escort him from the stage.

Another example of a class VI system is the play "Roboscopie" created at LAAS [12]. While portions of the PR2 robots behaviors were controlled by the human or preprogrammed, the inclusion of autonomous navigation and advanced perception techniques qualify it as a hybrid system.

Class VII is free of the restrictions of Class VI. One prime example are the performances in MIT's Robot Opera "Death and the Powers." [13] In the premiere productions, the story is told by simple robots, or "operabots," which are primarily autonomously controlled, using reactive navigation algorithms. However, their commands can also be overridden to allow direct teleoperation. The story of the opera revolves around a wealthy man, Simon Powers, who uploads himself into the a computer called "the system," which is embodied by several different parts of the set. The actor playing Simon is offstage most of the show, using wearable sensors and his voice to modulate the displays, sounds and lighting of the embodied system. Having the actor (and occasional teleoperators) remain in the loop qualifies this as a free control system since they have the power to drastically change the action of the play.

Another performance technology that falls into this category is robotcowboy [14], which integrates a human's live performance with algorithmically generated displays and sounds.

Category 4: Autonomous Acting: Finally, we have the category which features the most autonomous of all the robot actors. These robots' behaviors are generated primarily by algorithms, with little human intervention with exception of perhaps feeding in the source material. The human will still write the algorithm[1], but the performance will be generated by that algorithm, not explicitly programmed in or inspired by the performance of the human.

One of the closest examples of Class VIII is Data, the robot stand-up comedian [15][16], which is algorithmic and closed-loop. Although the library of jokes is written by humans, the selection of the jokes is completely autonomous, based on the optimization of a function that estimates audience enjoyment. The robot selects what jokes to perform based on audience feedback. The selection of jokes is determined not by what the human thinks will constitute a good performance, but what the robot selects. However, the robot comedian is limited in its ability to only perform a finite number of jokes, lacking the ability to invent new material based on audience response.

The final class, Class IX has free control of its algorithmic performance. To the author's knowledge, no autonomous

[1]For the time being...

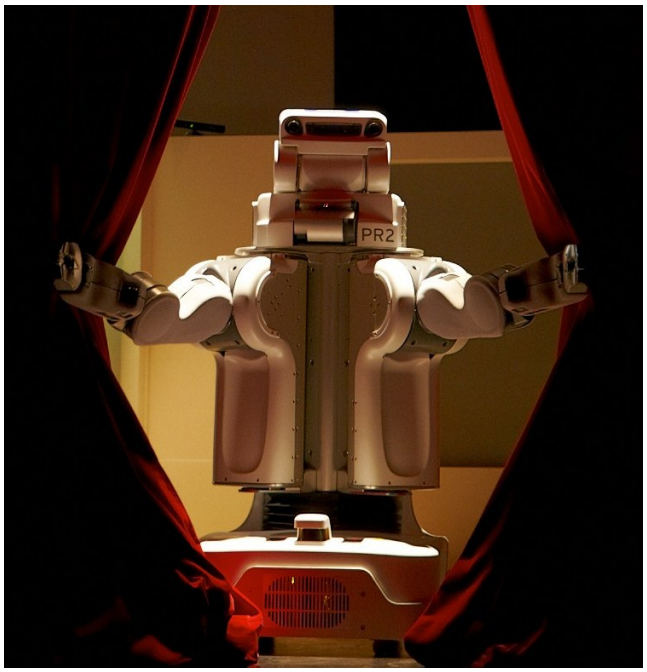

Fig. 4. The PR2 from "Roboscopie" play at LAAS. Source: http://www.openrobots.org/wiki/roboscopie

systems of this type exist currently. However, in the realm of fiction, the examples are plentiful, most notably another Data, Lieutenant Data from "Star Trek: The Next Generation."[17], [18], [19] Lt. Data demonstrates strong artificial intelligence, and is thus able to algorithmically make his own choices. He demonstrates this capacity in frequent simulations in the holodeck, where he hones his acting skill. In some cases, he bases his performance on historical human performances, pushing him more into Class VII. However, as he learns, he starts to invent his own choices and decide for himself how best to present the role. His ability to invent the performance, on his own save for the occasional tutlage from Picard, makes Data the most advanced class of robot actor.

Calculon, the robotic star of the fictional soap opera "All My Circuits" in the show "Futurama,"[20] also fits into this category.

IV. Conclusion and Analysis

Over the past few decades of robotics, in addition to the traditional means of sharing one's work with others, i.e. academic publishing, demos are increasingly common as ways to prove the worth of the research. Whether in the form of live demonstrations or videos, these demos can often conceal both the complexities and the shortcomings of a particular technology. A working demo can be indicative of an incredibly complicated recognition task working with the latest adaptive articulation platforms, or it could be the result of carefully controlled conditions. Nowadays, most researchers attempt to make it clear where their deceptions are, by acknowledging them clearly. This means labeling videos when they are the result of teleoperation, or when they are running open loop, or when they are sped up considerably.

However, in the current environment, robot theatre seems to be an exception to this trend. Describing the work (especially in the press), all robot theatre is referred to nearly equally as "robot acting." In the artistic sense, this might make sense. For human and robot actors alike, as an audience member, who cares what process was used to get to the result? In the end, it is the performance that matters and not what happens in the black box head of the actor.[2] One could imagine a performance by the Lincoln Animatronic from Disneyland. Using a Class I system, the robot could conceivably give the same exact performance if it were a Class IX robot. If the performance is effective and affects the audience in the same intended way, then the process is irrelevant.

However, for academics and developers working in the field, there is a world of difference between Class I and Class IX. And yet previously, we have had no words to differentiate the two. Both are just robot actors. This paper is intended to give structure to the conversations about future robot theatre by providing an ontology of the different kinds of robot theatre. This way, future productions with robots will be able to contextualize the mechanisms behind their systems within the ontological framework.

Having this system is also beneficial in that it allows us to give credit where credit is due for impressive systems. This applies to non-theatrical demonstrations as well. Consider the difference between the systems behind Willow Garage's PR2 in the study by Takayama et al [21], and Georgia Tech's SIMON programmed by Gielniak et al [22]. Both systems perform animated expressive gestures. However, the PR2 is a Class I system with the motions developed by a Pixar animator, and SIMON is a Class VI system that uses an algorithm to autonomously exaggerate pre-programmed motions and maintain eye contact with the study participants. The different class rankings do not necessarily mean one system is superior to the other. However, in the former, the credit for an effective performance should go to the animator, whereas in the latter, the credit should go to the algorithm and its creator.

Proper attribution of credit is important, given that the field of robot theatre does at times face hostility from some in the theatre establishment who worry that the addition of robots takes away from the essential humanity of the theatre. However, even in the autonomous category of robot actors, there is still a human behind the algorithm that programmed the system to accomplish that task, much in the same way that the people who create costumes, sets and lights are still able to bring humanity to their creations. All people working on theatre are working toward the same goal of creating artifacts, be they performances and physical objects, which convey the message and emotion that the play requires. Much like the electric light and sound amplification, the addition of robots to theatre signifies a new tool that can expand the realm of possibilities of how stories are portrayed on the stage.

[2]For further elaboration on the "process does not matter" approach to acting and robots, see [3]

V. Thanks

Thanks to Bill Smart for this insight and encouragement on this paper. Thanks to Laurel Riek for her insightful question about the role of the tape player in I, Worker. Thanks to Dan Lazewatsky and Agata Kargol for their help exploring the corner cases of the ontology. Thanks to Heather Knight for her helpful discussions of the role of the audience.

References

[1] K. Čapek, *R.U.R. (Rossum's Universal Robots)*, English ed. Samuel French, 1923, translated by Paul Selver and Nigel Playfair.

[2] BBC News, "Actor robots take japanese stage," 26 Nov. 2008, http://news.bbc.co.uk/2/hi/asia-pacific/7749932.stm.

[3] D. V. Lu and W. D. Smart, "Human-robot interactions as theatre," in *RO-MAN 2011*. IEEE, 2011, pp. 473–478.

[4] T. Sheridan and W. Verplank, "Human and computer control of undersea teleoperators," DTIC Document, Tech. Rep., 1978.

[5] S. Nakaoka, A. Nakazawa, F. Kanehiro, K. Kaneko, M. Morisawa, H. Hirukawa, and K. Ikeuchi, "Learning from observation paradigm: Leg task models for enabling a biped humanoid robot to imitate human dances," *The International Journal of Robotics Research*, vol. 26, no. 8, pp. 829–844, 2007.

[6] D. V. Lu and W. D. Smart, "Polonius: A wizard of oz interface for hri experiments," in *Proceeding of the 6th ACM/IEEE international conference on Human-robot interaction*, ser. HRI '11. ACM, 2011, pp. 197–198.

[7] D. V. Lu, A. Pileggi, J. Rincker, A. M. Mohr, and W. D. Smart, "Colloquium on performing arts and robotics (video)," in *ICRA Workshop on Robots and Art*. Shanghai, China: ICRA, May 2011.

[8] J. Keller. (2011, 18 Feb.) Craig Ferguson on Stand-up, Theme Shows and Sitting Out the Late-Night Wars. [Online]. Available: http://www.aoltv.com/2011/02/18/craig-ferguson-interview/

[9] B. Brantley, "In robot world she turns more Hedda than Hedda," *The New York Times*, 18 Feb. 2006, http://theater2.nytimes.com/2006/02/18/theater/reviews/18hedd.html.

[10] B. Duncan, R. Murphy, D. Shell, and A. Hopper, "A midsummer night's dream: social proof in HRI," in *Proceeding of the 5th ACM/IEEE international conference on Human-robot interaction*. ACM, 2010, pp. 91–92.

[11] W. M. Miller, Jr, *The Darfstellar*. Ballantine Books, 1955.

[12] S. Lemaignan, M. Gharbi, J. Mainprice, M. Herrb, R. Alami *et al.*, "Roboscopie: A theatre performance for a human and a robot," in *Proceeding of the 7th ACM/IEEE international conference on Human-robot interaction*, ser. HRI 2012 (video). ACM, 2012.

[13] E. Jessop, P. Torpey, and B. Bloomberg, "Music and Technology in Death and the Powers," *New Interfaces for Musical Expression*, pp. 349–354, 30 Jun. 2011.

[14] D. Wilcox and D. Wilcox, "robotcowboy: A one-man band musical cyborg," Ph.D. dissertation, Citeseer, 2007.

[15] H. Knight, S. Satkin, V. Ramakrishna, and S. Divvala, "A savvy robot standup comic: Online learning through audience tracking," in *International Conference on Tangible and Embedded Interaction*, Jan 2011.

[16] H. Knight, "TEDWomen: Silicon-based comedy," www.ted.com, Dec. 2010. [Online]. Available: http://www.ted.com/talks/heather_knight_silicon_based_comedy.html

[17] R. D. Moore (Writer) and R. Scheerer (Director), "The Defector," Star Trek: The Next Generation, Jan. 1990, 158.

[18] P. Lazebnik (Teleplay) and T. Benko (Director), "Devil's Due," Star Trek: The Next Generation, Feb. 1991, 187.

[19] J. Menosky (Teleplay) and C. Bole (Director), "Emergence," Star Trek: The Next Generation, May 1994, 275.

[20] P. M. Verrone (Writer) and B. Haaland (Director), "That's Lobstertainment!" Futurama, Feb. 2001, 3ACV08.

[21] L. Takayama, D. Dooley, and W. Ju, "Expressing thought: improving robot readability with animation principles," in *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 2011, pp. 69–76.

[22] M. Gielniak and A. Thomaz, "Enhancing interaction through exaggerated motion synthesis," in *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 2012, pp. 375–382.