

Affective Behaviors for Theatrical Agents

Michael Ruberry*[†]
michael.ruberry@yale.edu

Sara Owsley[†]
sowsley@cs.northwestern.edu

David A. Shamma[†]
ayman@cs.northwestern.edu

Kristian Hammond[†]
hammond@cs.northwestern.edu

Jay Budzik[†]
budzik@cs.northwestern.edu

Conrad Albrecht-Buehler[†]
conrad@cs.northwestern.edu

*Department of Computer Science
Yale University
P.O. Box 208285
New Haven, Connecticut 06520

[†]Intelligent Information Laboratory
Northwestern University
1890 Maple Avenue, 3rd Floor
Evanston, Illinois 60201

ABSTRACT

In this article, we explore the role of emotion in digital theater. Unlike heavily scripted productions, we create a theatrical agency which uses the world wide web to determine what content to deliver and how to deliver it with emotional competency. We describe the use of the web as a cultural artifact and how an autonomous agency can externalize these online connections during a performance. We address the problems that arise in bringing affect to an automated theatrical experience, and present our solution.

Keywords

Network Arts, Theater, Emotion, Affect, World Wide Web, Search, Agents.

INTRODUCTION

Computing machinery is no longer a device of pure computation. People have readily grasped the Web and use it to share their personal life experiences through message boards, web logs (blogs), and other community building web sites. We have created an installation which exposes and highlights these shared experiences through improvisational theater; an on stage production autonomously driven by the zeitgeist of the online community [10].

Our work in this area focuses on a unique intersection of art and technology. Not only does the machinery play several actors, but also the director. Our goal is to create a theatrical agency that can exist on stage alongside human players, conveying both content and emotion. To successfully convey emotion, we must enable the agents to assess the affective connotation of words they are speaking. In this paper, we will discuss a method to automate the affective analysis of



Figure 1: Two actors in a warm up game from the Association Engine.

words and how enabling the agents with this ability makes the theatrical experience more compelling.

THE ASSOCIATION ENGINE

The Association Engine is an example of such a theatrical experience. It is comprised of five virtual, autonomous actors. These actors participate in an improvisational warm-up game which is followed by a performance [7]. The performance can take the form of a simple story generation or a performance driven by what people have published on the Web, particularly in the form of blogs. In either embodiment of the Association Engine, portrayal of affect by the actors is crucial to engaging the audience.

Improvisation is about being in the moment, a team of actors on time and on the same contextual page. Human actors use an improvisational warm-up game, called *The Pattern Game*, to arrive at a mutual context. Actors stand in a circle, performing free association given a seed word. The Association Engine initiates this Pattern Game with its virtual actors by

Pattern Game

hate	→ love	→ adore	→ abhor
	→ repugnant	→ loathe	→ detest
	→ despise		

Table 1: A sample pattern from the improvisational warm-up game. ‘Hate’ was the seed, provided by the audience, and the remaining words were contributed by the actors.

Web Stories

I hate not being smart enough to make A’s in school with no effort. I hate going to bed early. I hate getting up early. I hate being broke. I hate not having a perfect body. I hate wearing bras. I hate breaking a nail after you have had long nails for such a long time.

I love walking around Boston with my Dad. Actually, I love walking around Boston with either one of my parents, but they both have completely different perspectives on the city. My dad is 67 years old. He grew up in a poor Irish family in East Boston, one of five kids born to my grandparents, who were immigrants from Newfoundland.

Table 2: Two discovered stories from the Pattern Game context in Table 1. Both stories are direct retrievals from blogs on livejournal.com.

taking a seed word from the viewers/audience, for example ‘hate.’ The first actor free associates from that word, speaking ‘love’ aloud using an animated face with a text-to-speech engine. The next actor in line adds his contribution, ‘adore,’ free associating from the first actor’s word, ‘love.’ The game continues until enough ideas have been supplied to create a context from which to do a performance. A sample completed chain can be found in Table 1.

The resulting word chain serves as a representation of context [3], from which the actors proceed with a performance. In the Association Engine, this performance can take the form of a *One Word Story* which the actors present, by contributing one word at a time. However, making these stories both interesting and coherent is a large problem of knowledge engineering. In response, we explored story discovery, more specifically, in the context of what people are currently thinking and saying online. The most recent embodiment of the Association Engine is a performance in which the actors read from blogs [5], discovered on the Web, that are on point with the Pattern Game. Each actor takes the stance of one particular blog author and performs that blog. Two discovered stories from a Pattern Game context are shown in Table 2.

Affect in the Association Engine

One metric of an human actor’s success is their ability to communicate emotion during a performance. The Association Engine’s actors are no exception; we require a notion of affective states to carry the audience through a compelling performance. For the pattern game, this becomes a problem of being able to identify what sort of emotion the actors should portray given each word or phrase spoken. This not only more effectively communicates the warm up game to the audience but also strengthens the association within the resulting context, allowing another metric of association to be traversed. The positive and negative affect could be selected by the agents for the warm up game or the story telling; the latter allowing a larger gamut of stories along the same topic but different emotional stances to be performed.

RELATED WORK

Most work in the field of affect determination is concerned with large corpora of text. It employs context, statistical NLP, hidden Markov models, Bayesian networks, and real world knowledge in an attempt to understand what is being said and what it means affectively. The massive characterization of ‘common sense’ knowledge exhibited by the OpenMind project [11] is one such corpus of real world knowledge that has been used with relative success for affective document characterization [4].

Yet the ability to diagnose the related space of an individual word effectively is still beyond the power of such a contextual database, which is rarely concerned with colloquial word use. Understanding an individual word’s affect is almost a moot point on document classification: keyword spotting is thankfully no longer considered a viable method.

But what about the cases where there is only an individual word, such as the Association Engine’s performance of the pattern game? In these cases keyword spotting is the only method available. Furthermore, an extensive list of such keywords is required. This list can also be valuable as part of a more advanced system for affect characterization of documents. Our work with affect aims to provide this list in a way that no human could.

AN AFFECTIVE CLASSIFICATION OF WORDS

A large corpus of affective words with comprehensive scores is difficult to come by. Traditionally, word valence was determined by focus groups, human studies, surveys, etc. Since our actors deal with a large subset of the English language and are purely autonomous, existing corpora will not scale to our needs. Although several people have worked on an affective lexicon [6] to classify words, we need a substantially larger corpus with precise, verifiable scores.

To create this corpus, we chose to extend an existing set of affectively scored terms, the Affective Norms of English Words (ANEW) [1]. This list contains 1,034 unique words with affective valence (a scale from unpleasant to pleas-

ant/negative to positive), arousal (a scale from calm to excited), and dominance (a scale from submissive to dominated) scores on a scale of 1–9.

Starting with the ANEW list as a seed, our classification expands outward into surrounding synonym sets from Thesaurus.com [12], an online thesaurus, to characterize new words’ affect scores. Given a word’s score and synonym set we determine its colloquial affiliation with other words and use this information to assign new values based on ones already known.

Concept Level Encoding

We realize that the ANEW list’s simple characterization of words is not entirely appropriate for our purposes. For example, the ANEW list represents ‘dog’ as valence = 7.57, arousal = 5.76, and dominance = 6.25. This encoding of ‘dog’ might seem accurate, but is problematic for further synonym expansion. The singular word ‘dog’ has two distinct meanings: a noun (mammal) and a verb (to pursue). While the fluffy pet should score a high valence, pursuing someone (dogging them) definitely should not. In this case, the ANEW list’s naïve encoding is not sufficient.

Since one affective score per word cannot capture the disconnect between different senses of the word, we chose to encode multiple scores per word (one per concept). Since the ANEW list was compiled based on a person’s initial reaction to a word, we assume that the ANEW list score for a word should be used for the most common concept of the word, as determined from Thesaurus.com. Thus, expanding from one word on the ANEW list, say ‘charitable,’ means mapping the ANEW score for ‘charitable’ to the synonyms of its most common concept, ‘generous.’

Affective Synonyms

In our representation, a word’s synonyms define its contextual space. But a synonym set alone offers only a small insight into a word’s affective connotations. It lacks information about colloquial usage, and when representing the entirety of a word’s lexical space it can be confusing and even contradictory.

To address this, we divide each word into its senses and concepts, as discussed above. Instead of an affect score associated with a raw word token, scores are associated with words’ concepts. By dividing a word in this fashion, its synonym set is also divided among its concepts. This new idea of a word’s space clearly represents its multiple meanings, allowing score expansion with no fear of a contradiction in terms. This prevents ‘dog’ the noun’s cute puppy score from being confused with ‘dog’ the verb’s negative pursuing score.

To capture the colloquial space of a word we utilize web document frequency [9]. A simple algorithm, based on the cultural associations discovered on the Web, can determine how closely related terms are. Based on these relations, the syn-

Word	Valence	Arousal	Dominance
Charity	7.82	4.4	5.85
Unselfish	6.89	4.5	5.7
Pointed	3.95	4.24	4.78
Embrace	7.72	4.4	5.85
Uneducated	2.5	6.31	4.91
Utopia	7.41	5.14	6.13
Demented	2.85	5.83	4.12
Sarcastic	3.85	4.24	4.78
Hospitality	7.82	4.30	5.67
Submissive	3.86	4.11	3.09
Terror	2.76	7.21	4.63
Gruesome	1.93	6.27	3.58

Table 3: Using 100 words from the ANEW list to seed of our algorithm, we expanded the list to 2,908. Here is a random sampling of word scores that were determined using our algorithm.

onyms are ranked and used to calculate the affect score for the specified concept of a word.

Let $W(s)$ represent a function that returns the web document frequency for a given query string s :

$$\frac{W(\text{Synonym} + \text{Word})}{W(\text{Word}) + W(\text{Synonym})} \quad (1)$$

determines the weight that the synonym’s affect score will have on its associated concept’s affect score: synonyms that are more closely related will have a larger weight on the concept’s score.

From equation (1), we gather that the highest weight possible is $\frac{1}{2}$ in the case of total co-occurrence. We have found this to be an effective measure of synonym proximity.

As an example of this, consider the word ‘church,’ and two of its synonyms, ‘cathedral’ and ‘chantry.’ To determine the weight that ‘cathedral’s’ affect score has on church’s affect score, we use equation (1). As seen in equation (2), the weight for ‘cathedral’ was 0.0283. Similarly, as seen in equation (3), the weight for ‘chantry’ was 0.0009. By our measure, ‘cathedral’ is a much closer synonym of ‘church’ than ‘chantry’ is; meaning that ‘cathedral’ will be much more influential than ‘chantry’ in calculating the affect score for ‘church.’

$$\frac{W(\text{church cathedral})}{W(\text{church}) + W(\text{cathedral})} = 0.0283 \quad (2)$$

$$\frac{W(\text{church chantry})}{W(\text{church}) + W(\text{chantry})} = 0.0009 \quad (3)$$

Finding The Most Common Concept

Characterizing ‘dog’ based on synonyms retrieved from Thesaurus.com, our encoding of the noun ‘dog’ as the concept ‘mammal’ is valence = 7.57, arousal = 5.76, dominance

Word		Valence	Arousal	Dominance
Dog	ANEW	7.57	5.76	6.25
Burn	ANEW	2.73	6.22	4.22
Dog	Noun	7.57	5.76	6.25
Dog	Verb	3.60	5.75	4.56
Burn	Fire	6.16	4.67	5.26
Burn	Pain	2.73	6.22	4.22
Burn	Yearn	5.57	7.57	6.00

Table 4: Two examples of expanding affect scores from the ANEW list. The most common concept of the word inherits the ANEW score. Affect scores for other concepts of the same term are calculated via the scores of their synonym sets.

= 6.25, carried over from the ANEW list. Our encoding of the verb ‘dog,’ as the concept ‘to pursue,’ is valence = 3.60, arousal = 5.75, and dominance = 4.56. These scores are more meaningful and accurate as the verb ‘dog’ is now represented with a more negative connotation than the noun.

While the most common concepts are typically representative of the ANEW list’s affect score for words, this is not always so. As an example of this, consider the word ‘burn.’ The ANEW list gives ‘burn’ a very negative valence of 2.73. The most common concept for ‘burn’ (from Thesaurus.com) is the verb meaning ‘fire,’ and its synonym set contains words such as ‘bake,’ ‘blaze,’ and ‘light.’ The ANEW list’s scores for these words are neutral, indicating that the subjects in the ANEW studies probably did not think of the concept of ‘fire’ when they first heard the word ‘burn.’ In this case, propagating the negative valence onto the synonyms of ‘burn’ under the concept of ‘fire’ would not be ideal. Table 4 shows a comparison of the ANEW list scores with our method.

As a solution to this problem, we consider the next most common concept and assign the ANEW score to this concept if such a disconnect does not exist between the ANEW score and the ANEW scores of synonyms within this concept. Moving on to the next most common concept of ‘burn,’ a better match is found. The concept of ‘pain,’ with synonyms ‘sting,’ ‘hurt,’ and ‘bite’ represents a much more negative valence. As such, the scores for ‘burn’ from the ANEW list were assigned to the less common concept of ‘pain.’

Since later expansion relies heavily on the integrity of our concept-based representation, ensuring that the initial encoding of ANEW words correctly maps to concepts is vital. As can be seen from the final mapping of burn, each concept carries a distinctive score. These scores are only meaningful if the original ANEW words have been correctly represented. Table 3 shows a sample of our expanded list.

SCORES TO EMOTIONS TO FACIAL EXPRESSIONS

When integrating emotions into our theatrical agents, we want the agents to show Ekman’s six emotional states: sad, happy, surprise, fear, anger, and disgust [2]. In addition, we add a neutral state for words that do not fit well into any of

these categories. With our expanded list, the resulting classification is comprised of concepts and an associated VAD (Valence Arousal Dominance) scores. We follow a simple mapping of the VAD scores to the seven emotional states.

We chose facial expressions to match these emotions. Using Ekman’s studies of facial expressions, we chose mouth, lip, eye, and eyebrow positions for each state. Each emotion has an associated action set, built on Ken Perlin’s Responsive Face framework [8] to transform the face from any emotion to any emotion.

AFFECT IN THE PATTERN GAME

Enabling the actors with an affective response to their participation in the Association Engine’s Pattern Game results in a more compelling performance. The agents are no longer dry and unresponsive, but show an affective response during the game. Based on each one word contribution and the mood of the previous contributions, an affect score is realized which is then translated into an action set and passed back to the actor as they speak their contribution. The action set is a set of commands that tells the actor how to communicate the desired emotion. Not only is an action set sent to the speaking actor, but action sets are sent to all of the listening actors to tell them how to affectively respond to the word that the speaker contributes. This results in the team of actors displaying emotions based on each word they speak and exhibiting an emotional response to the others’ contributions.

FUTURE WORK

Our affective classification opens the doors to new possibilities for bringing affective behaviors to our theatrical agents. Not only can they emotionally respond to words in an improvisational performance, but they can actually retrieve stories from the Web based on emotional stance, enabling the theatrical agents to juxtapose happy and angry stories on the same topic. Empowered by these new capabilities, we hope to produce an engaging theatrical experience driven by emotion and cultural understanding.

Bringing the Association Engine to the improvisational stage involves a crucial relationship between the actors and the audience. We have discussed the necessity of affect to keeping the audience engaged. In addition, the virtual actor to human actor interaction becomes crucial. The dynamic of the teamwork of a good improv troupe is like no other. The integration of affective states could prove to be very influential to the interaction of our theatrical agents and human actors.

Finally, we hope to improve our representation of word space through the incorporation of data from more diverse sources, such as from free association databases. This will further our understanding of a word’s colloquial connotations and more accurately represent its affective meaning. It also opens the possibility for more advanced manipulation in the future.





Figure 2: Top: Three actors in a warm up game speaking terms with no attributed valence. Bottom: Actors use the attributed valence scores to adjust their displayed emotion: happy, scared, angry.

ACKNOWLEDGMENTS

The authors would like to thank Andrew Ortony at Northwestern University for his affective know-how and advice. Hugo Liu and Henry Lieberman at the MIT Media Lab for their assistance with this work. In addition, continuing thanks to the guiding comments of Larry Birnbaum and other members of the Intelligent Information Laboratory at Northwestern University.

REFERENCES

1. M. M. Bradley and P. J. Lang. Affective norms for English words (ANEW): Stimuli, instruction manual, and affective ratings. Technical Report C-1, Center for Research in Psychophysiology, University of Florida, Gainesville, Florida, 1999.
2. P. Ekman. *Emotions Revealed: Recognizing Faces and Feelings to Improve Communication and Emotional Life*. Henry Holt and Company, New York, NY, 2003.
3. B. Laurel. *Computers as Theater*. Addison-Wesley, 1993.
4. H. Liu, H. Lieberman, and T. Selker. A model of textual affect sensing using real-world knowledge. In *Proceedings of the 8th international conference on Intelligent user interfaces*, pages 125–132. ACM Press, 2003.
5. LiveJournal. <http://www.livejournal.com>, 2004.
6. A. Ortony, G. L. Clore, and M. A. Foss. The referential structure of the affective lexicon. *Cognitive Science*, 11:341–362, 1987.
7. S. Owsley, D. A. Shamma, K. J. Hammond, S. Bradshaw, and S. Sood. The association engine: a free associative digital improviser. In *Proceedings of the 12th International Conference on Multi-Media*. ACM Press, 2004.
8. Responsive Face Project, NYU Media Research Lab. <http://www.mrl.nyu.edu/perlin/facedemo/>, 2000.
9. D. A. Shamma, S. Owsley, S. Bradshaw, and K. J. Hammond. Using web frequency within multi-media exhibitions. In *Proceedings of the 12th International Conference on Multi-Media*. ACM Press, 2004.
10. D. A. Shamma, S. Owsley, K. J. Hammond, S. Bradshaw, and J. Budzik. Network Arts: Exposing cultural reality. In *Alternate track papers & posters of the 13th international conference on World Wide Web*, pages 41–47. ACM Press, 2004.
11. P. Singh, T. Lin, E. T. Mueller, G. Lim, T. Perkins, and W. L. Zhu. Open mind common sense: Knowledge acquisition from the general public. In *Proceedings of the First International Conference on Ontologies, Databases, and Applications of Semantics for Large Scale Information Systems*. Heidelberg: Springer-Verlag, 2002.
12. Thesaurus.com. <http://www.thesaurus.com>, 2004.