

# Issues in Decision Tree Classification of Film Genre Using Plot Features

Joshua Tanenbaum

Simon Fraser University  
School of Interactive Arts + Technology  
250 – 13450 – 102 Avenue  
Surrey, British Columbia, V3T 0A3  
[joshuat@sfu.ca](mailto:joshuat@sfu.ca)

## Abstract

Classification and abstraction of narrative content is a problem that spans several domains and disciplines ranging from Film Studies and Narratology to Adaptive Hypermedia and Computational Preference Modeling. Applications for automatic recognition and classification of narrative content range from commercial recommendation systems to interactive digital storytelling (IDS) systems.

The notion of *genre* is commonly applied to a narrative as a means of classification; however *genre* is imprecise for *taxonomic* classifications, as it suffers from various overlaps and inconsistencies that defy systematic categorization. This paper describes a machine learning approach to genre classification that uses a decision tree to learn the genre associations of films based on a database of “plot keywords”, representing narrative and stylistic features. While the approach is ultimately unsuccessful, the challenges encountered point at some important lessons for future work in computational genre recognition.

## 1. Introduction

The classification and abstraction of narrative content is a problem that spans several domains and disciplines ranging from Film Studies and Narratology to Adaptive Hypermedia and Computational Preference Modeling. Much work has been done on how to categorize a given story, how to describe the potential target audience for that story, and how to classify it in regards to other stories. Perhaps the earliest work of Narrative Formalism is Aristotle’s *Poetics* [1], which separated narrative into epic poetry, tragedy and comedy. Since Aristotle, many have tried to describe narrative in terms of formal structures. In 1928, the Russian formalist Vladimir Propp published his *Morphology of the Folktale*, which attempted to reduce the narrative structures of folklore to a set of elemental functions [18]. His creation of a notation language to describe narrative elements has been adapted by many computer scientists working in Automatic Story Generation [12, 17, 21, 25].

Applications for the automatic recognition and classification of narrative content range from commercial recommendation systems to interactive digital storytelling (IDS) systems. One major application for narrative content classification is in commercial recommendation systems for films, such as those found on Amazon.com and Netflix. These systems attempt to intelligently recommend a film to rent or purchase based on knowledge of a user’s tastes and

viewing preferences. Commercial recommendation systems model a number of user preferences including past viewing habits, actor and director preferences, similarity to other user’s tastes, and genre preference.

IDS systems often attempt to build some model of the reader’s preferences before adapting narrative content to match [22, 24, 26]. The two most common approaches to this problem attempt to model the reader’s emotional response or try to parse the reader into some sort of play-style stereotype. This paper is an exploration of a third valence for adaptivity: *genre preference*. The research contained in this paper represents a preliminary investigation of a machine learning approach to genre preference recognition for the Tangible Ubiquitous Narrative Environment (TUNE), an adaptive IDS system. In TUNE, interactors will explore a physical space which has been computationally augmented to contain a multi-linear interactive narrative. Each narrative “track” in TUNE is associated with a different genre, so that as interactors express a genre preference to the system, the genre of the experience will adapt.

In order to adapt to the interactor’s genre preference, TUNE requires some knowledge about the relationship between specific narrative and stylistic “features” and common narrative genres. In this paper, I describe a machine learning approach to the problem of genre recognition, using a decision tree. I first present an overview of the notion of *genre* from Film Studies, and then discuss previous approaches to film recommendation and classification. Finally, I describe GenreTree, a decision tree built using data gathered from the Internet Movie Database [27] and present some preliminary results, along with a discussion of lessons learned from this research.

## 2. Genre in Film

Genre films work by engaging viewers through an implicit contract. They encourage certain expectations on the part of spectators, which are in turn based on viewer familiarity with the conventions. [11]

The notion of *genre* is commonly applied to a narrative as a means of classification; however, *genre* is imprecise for *taxonomic* classifications as it suffers from various overlaps and inconsistencies that defy systematic

categorization. There are at least three significant issues with genre as a system of classification:

1. Genre is different in and out of “the wild”
2. Genre arises from multiple sources for multiple reasons.
3. Genre is not exclusive.

In the following sections I will explore these concerns in greater detail.

## 2.1 Genre “in the wild” vs. formal genre

The relationship between “genre as commonly understood” and “genre as a formal critical notion in film-studies” is a problematic one. Film studies has developed critical definitions of genre that, while appropriate for a survey of film as a medium, do not always align with the average viewer’s genre categories. Langford compares the genre categories of film theorists with the taxonomic systems commonly found in video stores.

For example, while some of these genres – action, thriller horror, science fiction, comedy – match up fairly well with standard genre headings, the video store omits several categories widely regarded as of central importance in the history of genre production, such as Westerns, gangster films and musicals...let alone more controversial yet (in academic discussion) ubiquitous classifications as *film noir* or melodrama. Other categories are uncanonical by any standard: ‘latest releases’ is self-evidently a time-dated cross-generic category; ‘classics’ is generically problematic in a different way, since it apparently combines both an evaluative term (‘all-time classic’, ‘landmark’, etc.) with a temporal one... [15]

In order to build a classification system that will accurately reflect the genre understanding of an average film viewer, it is necessary to deal with genre as a messy and socially constructed phenomenon, rather than as a formal system. A computational model of genre built on the work of classical film scholarship would place an inordinate emphasis on genres such as the Western and the Musical, which would correctly reflect their importance in the history of film. This model would not reflect the genre models of today’s audiences, where westerns and musicals are outlying specimens.

## 2.2 Multiple sources of genre.

As a scheme for classifying narrative content, genre emerges from a diverse collection of narrative features including aesthetic and stylistic elements of the story, thematic and cultural associations, and specific aspects of the narrative subject matter. For example, *Film-Noir* is a genre term originally coined by French film critics in the 1950s to describe a particular type of film that had begun

to emerge from Hollywood in the wake of World War II [11]. Noir is a canonical film genre in film studies; however, as has been argued by David Bordwell, Noir is less of a broad genre category and more of a film *style* which can “seep into any genre” [4] as cited in [11]. It is therefore possible to find examples of Noir Westerns, Noir Science Fiction, and even Noir Comedy.

Genre distinctions can arise around film elements that are not specific to the content of the narrative, as in the case of the Musical where the presence of singing and dancing is the dominant genre determinant, and other narrative and stylistic concerns are subsidiary. Alternatively, some genre determinants are tied directly to the presence of specific narrative elements. Films with time travel, space travel, or alien races, are almost always classified as Science Fiction, while films with Magic invariably are categorized as Fantasy. In some cases, the genre distinction is harder to make: two films might both contain a story of a new relationship, with romantic elements and humorous elements. The ratio of these elements to each-other is a significant genre determinant, causing the one with more humorous elements to be classified as a Comedy, while the other is considered to be a Drama or a Romance. As discussed in [15], some genre distinctions are the result of the film’s age, or perceived cultural impact; hence, categories such as Classic, New-Release, Foreign Film, and World Cinema. Some genre distinctions arise from the circumstances of the film’s production, as in the Bollywood Musical, the Hong Kong Martial Arts Film, or the Independent Film, all of which share a common quality of having been produced outside of the Hollywood mainstream.

## 2.3 Genre is not Exclusive

As should be evident from the above discussion, genre is not a simple container into which a narrative can be placed, but is instead a property which a film can demonstrate in varying degrees and combinations. While the presence of aliens may mark a film as Science Fiction, it does not discount it from also being a Comedy (as in *Men in Black*), a Survival Horror film (as in *Alien*), an Action Adventure (as in *The Fifth Element*), or even a Retro-Futurist Film Noir (as in *Dark City*). Many modern Hollywood films mix and match genres freely, often as a strategy for appealing to as wide an audience as possible. Mixing genres is an important source of innovation and change in film narratives. [5]

As we will see in our discussion of the GenreTree system, the potential for one film to fit into multiple genre categories simultaneously in varying amounts is a not insignificant problem for any computational representation of genre.

### 3. Existing Classification Systems

Due to the potential commercial applications of computational film classification there is already a sizeable body of related work. We can separate research in this domain into two broad categories: automatic classification and intelligent recommendation.

#### 3.1 Automatic Classification

The first category relies on the computational processing of video data in conjunction with machine learning techniques to train a classification system. This approach uses techniques from AI research including Bayesian networks [16] and hidden Markov models [13], as well as less intelligent techniques such as audio-video feature analysis [19, 20] and color histogram analysis [8]. In all of these systems, classification systems are provided with video data that is intended to represent a specific single genre. Analysis of this data is performed to isolate identifying visual characteristics of the genre that can be used to categorize other video footage.

#### 3.2 Intelligent Recommendation

The second category relies on knowledge about the tastes of a user or group of users to make a film recommendation. With the recent surge in content creation and distribution technologies on the web, film viewers are often burdened with what Good et al. describe as Information Overload [10]. Information Overload can be mitigated through the application of recommendation systems of varying levels of “intelligence”. Good et al. describe three popular technologies for sorting information: *Information Retrieval (IR)*, *Information Filtering (IF)*, and *Collaborative Filtering (CF)* [10]. These techniques have seen much use in the creation of film recommendation systems

*Information retrieval (IR)* is a query-based information parsing technique that allows users to filter databases of information via explicit searches of keywords, cast/crew credits, etc. While this technique is commonly accepted for retrieving information it does not support any intelligent *recommendation* of content.

*Information filtering (IF)* is a technique that automatically filters lists of information based on either user specified or mechanically learned criteria. Information Filtering is a form of User Modeling, an AI technique that builds a computational model of what a user knows or wants, often used in the field of Adaptive Hypermedia [6, 7]. Several existing recommender systems take a user modeling approach [10, 23]. Information filtering is capable of making recommendations based on the content features of the media being recommended, but it suffers from not supporting “serendipitous discovery”; the technique works best when used to learn a specific preference and seek out material that is congruent with that preference [10].

*Collaborative Filtering* is a technique that builds databases of user opinions on available items, and correlates them in order to make recommendations. Unlike IF, which uses details about the content domain to make recommendations, Collaborative Filtering does not use any knowledge about the media and instead relies entirely on “nearest neighbor” similarity for its recommendations. One drawback to this technique is that it relies on the participation of multiple users in order to make recommendations. This approach is one of the most common, and can be seen in use for commercial recommendation systems like Amazon.com and in research systems such as [3, 10, 14].

#### 3.3 Limitations of these systems.

As we have seen in section 2, genre does not lend itself to formal computational representation. For this reason it is interesting to explore how the above systems choose to represent genre.

Some of the earliest work in automatic categorization emphasized “genres” that were visually distinct, and thus lent themselves to image analysis based approaches. Fischer et al’s system was able to recognize five genres: News Cast, Sports (car race), Sports (tennis), Commercials, and Animated Cartoon. [8] Iyengar and Lippman described a Hidden Markov Model based approach to automatic TV sequence classification, however their model of genre only contained two items: News and Sports.[13]

Rasheed et al. initially separated films into two categories: action and non-action. These were further separated into Comedy, Horror, and Drama (under non-action), and Fire/Explosions and Other (under Action). [19] Later research from the same group simplified this list further, to just four genres—Action, Comedy, Horror, and Drama—but recognized that films often exhibit multiple genres, and so also considered several combinations including Action-Drama, Comedy-Drama, and Action-Comedy. [20]

Symeonidis et al. combined two databases of film information: The Internet Movie Database (the same one used in this study), and an academic database known as MovieLens consisting of 100,000 film ratings elicited from 943 users on 1682 films. [23] The resulting database contained 23 distinct genres, as well as keyword information and information about cast and crew. MovieLens was also used by [14] and [3], although in both cases the system retained no representation of genre preference, and instead relied on correlating viewer tastes as represented to the system through a ratings scale.

We can see from this brief overview that systems for automatic categorization of genre have tended towards




Figure 1: A Portion of one of GenreTree's Graphs

models of genre that are overly simplistic. In the case of [8] and [13], the genres listed are far removed from the average viewer's notion of genre. The systems that rely primarily on user input and modeling are able to support a more diverse list of genres, however they are ultimately "content agnostic" and do not create any computational model of narrative elements.

#### 4. Building the GenreTree

GenreTree attempts to create a computational model of genre that is built on an analysis of the conceptual and aesthetic features of specific genre films. Unlike the automatic categorization approaches described above, GenreTree does not work directly with video image data, and unlike the intelligent recommendation systems described above, GenreTree is not concerned with the interests of a group of viewers, nor does it attempt to make specific film recommendations. The goal of GenreTree is to guess a viewer's genre preference in the moment, based on expressed interest in *specific narrative and stylistic features*. In effect, by building a *global model* of genre and feature relationships, GenreTree is able to create a *local model* of user preference at interaction time; a model which can be used to adapt narrative content in TUNE to the reader's tastes.

GenreTree explores the viability of a decision tree based machine learning approach to the task of genre recognition. I have implemented an open-source decision tree library in Java called JaDTi [9], which I have trained using two sets of genre data. The first set is hand-authored, and comprised of 60 films, in 6 genres, with 96 stylistic and narrative feature tags. The second set is derived from the Internet Movie Database, and is comprised of 14,467 Films in 19 Genres, and 2,227 "plot keywords" which represent stylistic and narrative features. Each of these sets has been split into two subsets: one for training the decision tree, and the other for testing what the tree has learned. These trees are visualized using an open source graphing tool called graphvis [2].

### 5. GenreTree Results and Lessons Learned

The creation of GenreTree was undertaken to determine if it was possible to train a computational system to learn the relationships between narrative features and film genres. As discussed above, the notion of genre is a slippery one, with at least three significant sources of confusion. These three concerns - that formal genre schemes differ from those in common usage; that genre arises from multiple sources; and that genres are non-exclusive - all deal with issues of categorizing films into generic categories. In the creation of a decision tree for genre identification, this research needed to contend with an even larger unresolved issue in genre studies: *whether specific narrative and stylistic features are predictive of genre at all*.

This decision tree operated on the hypothesis that it was possible to observe patterns in the relationships between features and genres. The results however, seem to confirm the knowledge of film scholars: there is no hard-fast rule for determining why a film is designated as one genre, and not another: at least no rule that is readily intuited by a human mind. In the following sections I will discuss the results of the GenreTree experiment and the difficulties encountered, and consider some possible future directions for this research.

#### 5.1 The Hand-Authored Data Set

The first tests of the GenreTree used the hand-authored dataset, in order to determine if the software's general principles were sound. The hand-authored set was comprised of 60 films that I was familiar with and that I considered representative of their respective genres. I tagged these with features that I considered salient about them, including specific narrative elements and more general stylistic features. The sample size of films in this set was far too small to generate significant results with only 54 films in the training set and 6 films - one of each genre - in the testing set. Even so, the GenreTree was able to consistently categorize 4 of the 6 films correctly. In experimenting with this set I was able to draw several

conclusions:

**Limiting the feature set resulted in significantly better results.** With such a small testing set and training set, 96 features was too large a number to make any coherent predictions. Through trial and error I found that by removing those features that recurred most frequently in the data, and those features that only appeared once or twice, I was able to increase the accuracy to about 66%. There appears to be a “sweet-spot” in the feature distribution between those which are too common (such as Sex and Violence) and those which are too idiosyncratic (such as Nazis, or Polar Bears).

**Too much of my own knowledge about genre was encoded in the data.** It was evident in this early stage that I was training a system to recognize and represent my own understanding of genre distinctions rather than a generally consensed upon notion of genre.

**More data was needed in order to draw any real conclusions.** It was also clear that a training set of 54 films and a test set of 6 were inadequate to the task of training a decision tree.

## 5.2 The IMDb Data Set

In order to address these concerns, I sought out a larger source of film data. The Internet Movie Database (IMDb) provides plaintext files of its content. These contained far more data than I could practically process. The keyword database held 314,958 films, tagged with 2,110,162 unique keywords representing features of the film. The genre database contained 516,980 films, each tagged with multiple genres (28 distinct genres) for a total of 826,162 film/genre pairs.

There were several significant issues with this initial dataset. Most problematic was that IMDb provided multiple genre tags for each film; however it provided no indication of the hierarchy of these identifiers. In other words, a film such as Star Wars was tagged as Action, Adventure, Family, Fantasy, and Sci-Fi, with no indication that the film was *primarily* Science Fiction, with strong aspects of the other genres. Of lesser concern was the fact that the database included a large collection of titles that were not feature films, but were instead television shows, DVD special features, documentaries of film production, video games, direct to video titles, and a large collection of pornographic and adult titles. Many of these were tagged with additional information such as (TV) for television shows and (VG) for video games and were easily filtered from the database, however the largest category of non-film content – pornography – had no clear indicator for automatic removal.

In order to transform this information into useful data for the decision tree, I used a series of Python scripts to correlate, prune, and reformat the raw text. I first created

a dictionary of features and their frequencies, which allowed me to prune out all features appearing less than 150 times throughout the database. I used this feature dictionary to select only those films which had 15 or more features associated with them, and at least one genre tag.<sup>1</sup> I also removed as much of the pornographic and non-feature film content that I could, pruning off 9 of the 28 genres including Short, Documentary, Adult, Music, Sport, Reality-TV, Talk-Show, Game-Show, and News. The initial distribution following this data cleanup is shown in Table 1:

**Table 1 Initial Genre Distribution**

Genre	Frequency
Action	2,638
Adventure	2,178
Animation	666
Biography	1,255
Comedy	4,619
Crime	2,734
Drama	7,392
Family	780
Fantasy	975
FilmNoir	166
History	483
Horror	1,708
Musical	652
Mystery	1,468
Romance	4,579
SciFi	877
Thriller	2,802
War	794
Western	1,199

In this initial version of the dataset, I made no effort to address the issue of multiple genres per film, hoping that the decision tree would be able to learn genre associations, even if a single constellation of features was mapped to multiple different genres. However, the decision tree constructed from this dataset was able to correctly categorize films less than 1% of the time: an outcome that was worse than if it were to randomly guess a genre. Having no other option, I returned to the data and rebuilt the set, this time selecting one genre at random from a given film’s set of possible genres, yielding a dataset comprised of 14,467 films, each tagged with 1 genre. [Table 2] Had the IMDb dataset contained information about the weighting of each genre, I would have been able to select the genre that best categorized the film in question, but sadly this data was not available and I had to

---

<sup>1</sup> I attempted a version of this code on a much larger dataset, with looser cutoffs for number of keywords per film, and frequency of keyword appearance, however this exceeded the memory capacity of the machine running the calculations, necessitating a more modest sample.

rely on random selection for movies with multiple genre tags.

**Table 2 Final Genre Distribution**

Genre	Frequency
Action	764
Adventure	637
Animation	293
Biography	182
Comedy	2,154
Crime	897
Drama	3,571
Family	252
Fantasy	280
FilmNoir	62
History	179
Horror	771
Musical	221
Mystery	420
Romance	1,718
SciFi	263
Thriller	967
War	284
Western	552
<b>Total:</b>	<b>14,467</b>

Finally, I split this data into a testing set and a training set. After performing these operations I was left with a training set of 9,643 films, and a test set of 4,824 films, both with 2,227 unique features.

The first attempt at training the tree considered the entire 2,227 feature list. To calculate accuracy, a script was written to evaluate the tree's predictions for the testing set, and compare them to the list of possible genres that each film was initially tagged with. If the output matched any of the possible genres, it was considered correct. Using the complete set of features, GenreTree categorized 569 out of 4,824 correctly (about 12% accuracy). In the tests with the smaller dataset, however, it had been clear that a smaller, more carefully selected set of features could achieve a higher accuracy. Bearing this in mind, I selected a very small collection of features (168 in total), once again using my own understanding of genre to guide what I included. This set yielded significantly better results: 1492 correct out of 4824 (about 31% accuracy). These results, while not particularly successful, do allow me to draw some practical conclusions for the future of this approach.

### 5.3 Conclusions

The fact that careful pruning of the feature set yielded a substantial improvement in precision leads to two conclusions:

**There is a correlation between certain narrative features and film genre.** Different combinations of

features result in different outcomes. These outcomes are significantly higher than can be explained by random chance. With 19 possible genres in the system, we would expect a minimum accuracy of 1 in 19, or 5.3%. 31% accuracy indicates that while the learning in the system is far from perfect, there is still some learning occurring.

**A human may not be able to apprehend this relationship, but an intelligent system might.** Tweaking the feature set to obtain a higher degree of accuracy than this is a matter of exploring a range of combinatorial possibilities in an intelligent manner. Doing this systematically is outside the ability of a human programmer, however this type of recombinant programming is ideal for a genetic algorithm, especially since the fitness function is already built into the system (in the form of the accuracy calculation).

We can also safely assume that one of the most significant sources of error in the system in its current state is the need to *randomly* select one genre for each film in the training set. Ideally, a future version of GenreTree will be an iterative series of trees, each constructed from a different level of a hierarchical list of genres. To do this, it would be necessary to first obtain some source of data that would allow for intelligent ranking of genres within a film. One potential source is the tagging API from Amazon.com, which can provide genre tags and frequency weighting for ranking them. GenreTree could then build individual trees for the primary, secondary, and tertiary (etc.) genres in the training set. If a film in the testing set is miscategorized by the primary tree it could be run through the subsequent trees until a correct category is found

While the GenreTree in its current form does not perfectly perform the task for which it has been built – namely the correct categorization of a cluster of features into an associated genre – it does point toward a series of techniques and approaches that might be able to succeed where it has failed.

### References

1. Aristotle *Aristotle's Theory of Poetry and Fine Art*. Dover Publications, Inc., New York, New York, USA, 1951.
2. Bilgin, A., Ellson, J., Gansner, E., Hu, Y., Koren, Y. and North, S. Graphviz - Graph Visualization Software, 2008. <http://www.graphviz.org/>
3. Bonnefoy, D., Bouzid, M., Lhuillier, N. and Mercer, K. "More Like This" or "Not for Me": Delivering Personalised Recommendations in Multi-user Environments. *UM 2007*, 4511. 87-96.
4. Bordwell, D., Staiger, J. and Thompson, K. *The Classic Hollywood Cinema: Film Style and Mode of*

- Production to 1960*. Columbia University Press, New York, 1985.
5. Bordwell, D. and Thompson, K. *Film Art: An Introduction*. The McGraw Hill Companies, Inc., New York, 1997.
  6. Brusilovsky, P. and Millan, E. User Models for Adaptive Hypermedia and Adaptive Educational Systems. in Brusilovsky, P., Kobsa, A. and Nejdl, W. eds. *The Adaptive Web*, Springer-Verlag, Berlin Heidelberg, 2007, 3-53.
  7. Fischer, G. User Modeling in Human-Computer Interaction. *User Modeling and User-Adapted Interaction*, 11. 65-86.
  8. Fischer, S., Lienhart, R. and Effelsberg, W., Automatic recognition of film genres. in *3rd ACM International Conference on Multimedia*, (San Francisco, California, 1995), ACM, 295-304.
  9. Francois, J.-M. jaDTi Decision Trees: A Java Implementation, 2004. <http://www.run.montefiore.ulg.ac.be/~francois/software/jaDTi/>
  10. Good, N., Schafer, J.B., Konstan, J.A., Borchers, A., Sarwar, B., Herlocker, J. and Riedl, J. Combining Collaborative Filtering with Personal Agents for Better Recommendations *Innovative Applications of Artificial Intelligence*, American Association for Artificial Intelligence, Orlando, Florida, United States, 1999.
  11. Grant, B.K. *Film Genre: From Iconography to Ideology*. Wallflower, London, 2007.
  12. Grasbon, D. and Braun, N., A Morphological Approach to Interactive Storytelling. in *Artificial Intelligence and Interactive Entertainment, Cast'01, Living in Mixed Realities*, (Sankt Augustin, Germany, 2001), 337-340.
  13. Iyengar, G. and Lippman, A., Models for automatic classification of video sequences. in *Storage and Retrieval for Image and Video Database VI*, (San Jose, CA, USA, 1998), SPIE, 216-227.
  14. Kim, H.-N., Ji, A.-T., Yeon, C. and Jo, G.-S. A User-Item Predictive Model for Collaborative Filtering Recommendation. *UM 2007, 4511*. 324 - 328.
  15. Langford, B. *Film Genre: Hollywood and Beyond*. Edinburgh University Press, Edinburgh, 2005.
  16. Ono, C., Kurokawa, M., Motomura, Y. and Asoh, H. A Context-Aware Movie Preference Model Using a Bayesian Network for Recommendation and Promotion. *UM 2007, 4511*. 247 - 257.
  17. Peinado, F., Gervas, P. and Dias-Agudo, B., A Description Logic Ontology for Fairy Tale Generation. in *Workshop on Language Resources for Linguistic Creativity*, (Lisbon, Portugal, 2004), LREA, 55-61.
  18. Propp, V. *The Morphology of the Folktale*. University of Texas Press, Austin, Texas, 1968.
  19. Rasheed, Z. and Shah, M., Movie Genre Classification By Exploting Audio-Visual Features of Previews. in *16th International Conference on Pattern Recognition (ICPR'02)*, (2002), IEEE.
  20. Rasheed, Z., Sheikj, Y. and Shah, M. On the Use of Computable Features for Film Classification. *IEEE Transaction on Circuit and Systems for Video Technology*, 1 (11).
  21. Schneider, O., Braun, N. and Habinger, G., Storylining Suspense: An Authoring Environment for Structuring Non-Linear Interactive Narratives. in *European Association for Computer Graphics (Eurographics): Journal of WSCG'2003*, (Plzen: University of West Bohemia, 2003), 411-417.
  22. Seif El-Nasr, M., A User Centric Adaptive Story Architecture - Borrowing from Acting Theories. in *International Conference on Advances in Computer Entertainment Technology*, (2004), ACE, 109-116.
  23. Symeonidis, P., Nanopoulos, A. and Manolopoulos, Y. Feature-Weighted User Model for Recommender Systems. *UM 2007, 4511*. 97-106.
  24. Tanenbaum, J. and Tomizu, A. Affective Interaction Design and Narrative Presentation *AAAI Fall Symposium on Intelligent Narrative Technology*, AAAI Press, Arlington, VA, USA, 2007.
  25. Theune, M., Faas, S., Nijholt, A. and Heylen, D., The Virtual Storyteller: Story Creation by Intelligent Agents. in *TIDSE 03: Technologies for Interactive Digital Storytelling and Entertainment*, (Fraunhofer, 2003), 205-215.
  26. Thue, D., Bulitko, V., Spetch, M. and Wasylishen, E. Interactive Storytelling: A Player Modelling Approach *3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2007)*, AAAI Press, Palo Alto, California, 2007.
  27. Unknown. Internet Movie Database, Amazon.com, 2008. <http://www.imdb.com>