Extraction of User Opinions by Adjective-Context Co-clustering for Game Review Texts

Kevin Raison, Noriko Tomuro, Steve Lytinen, and Jose P. Zagal

DePaul University, College of Digital Media,
243 S. Wabash, Chicago, IL 60604 USA
raison@chatsubo.net
tomuro,lytinen@cs.depaul.edu
jzagal@cdm.depaul.edu

Abstract. We present our preliminary work on extracting fine-grained user opinions from game review texts. In sentiment analysis, user-generated texts such as blogs, comments and reviews are usually represented by the words which appeared in the texts. However for complex multi-faceted objects such as games, single words are not sufficient to represent opinions on individual aspects of the object. We propose to represent such an object by pairs of aspect and its quality/value, for example "great-graphics". We used a large adjective-context co-occurrence matrix extracted from user reviews posted at a game site, and applied co-clustering to reduce the dimensions of the matrix. The derived co-clusters are pairs of row clusters × column clusters. By examining the derived co-clusters, we were able to discover the aspects and their qualities which the users care about strongly in games.

Keywords: adjective-context relation, game reviews, co-clustering, sentiment analysis, opinion mining

1 Introduction

Sentiment analysis has been receiving attention in recent years in Natural Language Processing (NLP) [7]. With the proliferation of user posts on online social media, such as weblogs, product reviews and message boards, NLP techniques have been applied to automatically extract and analyze the user opinions and sentiment expressed in those texts. Sentiment analysis has been incorporated in a variety of applications, for example to obtain product marketing information [8], to track political opinions [10], and to search for "buzz" (hot topics) in social networks. Game, especially videogame, is one of the domains which would benefit from automated sentiment analysis. Not only is there an enormous amount of data (reviews, comments, etc.) available already (for instance at game sites such as Gamespot (`www.gamespot.com`), IGN (`www.ign.com`) and Giantbomb (`www.giantbomb`)), new games are created continuously and rapidly. Moreover, game users/players are generally quite passionate and vocal about the games they like or dislike, therefore reviews tend to be long.

In sentiment analysis (or opinion mining), user-written texts are typically represented by the words which appeared in them, then categorized for polarity orientation (positive/negative/neutral) or emotions (e.g. joy, sadness). However, for complex objects that have many facets/aspects such as games, single words are not sufficient to represent opinions on individual aspects of the object. For example, reviews such as "graphics are great, but gameplay is horrible" and "graphics are horrible, but gameplay is great" are critically indistinguishable if only single words are used (represented by the same four single words: "graphics", "great", "gameplay", "horrible"). A better approach is to represent each aspect and its quality/value together, for instance "great-graphics" and "horrible-graphics". This representation scheme can also express more fine-grained as well as accurate opinions expressed in the texts.

In this paper, we present the preliminary results of our work on extracting fine-grained opinions from game review texts. We started from the adjective co-occurrence dataset used in [11, 12], which contained the adjective bigrams (i.e., bigrams in which either word is an adjective) extracted from the user reviews posted at a game site (Gamespot). In this dataset, the adjectives and the *context words* (i.e., words which co-occurred with the adjectives in a bigram) were represented separately, by a matrix of adjectives (on the rows) $\times$ context words (on the columns). Then we applied co-clustering [3] to reduce the matrix. Reducing the size of the matrix was necessary in order to make the computation feasible, but it also served as a way to group similar words (e.g. "graphic", "look") or typographical errors (e.g. "graphic", "grafic"). Co-clustering is a technique which cluster rows and columns of a matrix simultaneously while preserving the dependencies between them. Then we examined the derived *co-clusters* (manually) to discover the aspects and their qualities of games which the users care about strongly. We also used those co-clusters to cluster games and investigated the results for any interesting patterns that might have emerged.

2 Related Work

Co-clustering has been used in several previous work to capture the dependency between two variables (or objects and features), which are typically represented by the rows and columns of a matrix. For example, [1, 2] applied co-clustering in the task of document categorization, and showed improved results obtained by utilizing the set of document clusters produced by co-clustering as compared to the clusters generated by clustering on a single dimension. Another work in [4] applied co-clustering in generating product recommendations. This work represented the item rating scores posted by the users by a matrix of items on the rows and users on the columns, and applied co-clustering to discover the patterns of user preferences. Then they used the derived co-clusters to predict rating scores and recommendations for new users.

As for the focus on adjectives as a part-of-speech (POS), from early work on sentiment analysis in NLP, adjectives have been effectively used to extract user opinions or polarity from texts [6, 7]. Recently, several work have used adjectives combined with nouns which co-occurred with the adjectives (e.g. "wonderful ideas", "horrible taste") for the purpose of extracting more accurate or fine-grained opinions. For example, [9] extracted adjective-noun pairs, which are in the dependency/modifying relation, from the opinions posted at an online forum on eGovernment for the purpose of mining the public opinions on various government decisions; while [1] used a similar approach to identify consumer preferences from product reviews (posted at Amazon).

The work which is closest to ours in this paper would be [13]. They extracted adjective-noun dependency pairs from parsed user review texts, and applied a model based on Latent Dirichlet Allocation (LDA) to derive clusters of adjective-noun pairs (where each cluster is a pair of an adjective set and a noun set). Whereas in our work in this paper, we used bigrams surrounding adjectives, and applied the *information-theoretic* co-clustering algorithm from [1] to derive co-clusters (of adjective-context pairs).

3 Game Review Dataset

We conducted experiments using the data used in [11, 12]. This dataset contained 723 adjectives and 5,000 context words (which appeared in a bigram surrounding an adjective, i.e., one word before and one word after the adjective). The data was extracted from the corpus of user reviews posted at Gamespot (as of April, 2009). The corpus covered 8,279 game titles, and the entries in the data were raw frequencies of the bigrams/co-occurrences in the corpus.

Now, to construct adjective-context co-occurrences from this dataset, we will end up with a total of 3,615,000 (= 723 $\times$ 5,000) pairs − prohibitive to use in any computational task. Although around a half of the entries (48%) had a value of zero, the frequency distribution had a rather fat tail: the percentages of the entries with values ≤ 2, ≤ 3 and ≤ 4 (all including 0) were 85%, 89% and 91% respectively. That means we will still be left with over 300,000 ($\approx$ 9%) features, even if we were to keep the ones with frequencies ≥ 5. Dimensionality reduction was obviously necessary.

4 Co-clustering

To reduce dimensions of the data, we first applied a co-clustering algorithm. In general, a co-clustering algorithm works such that, given a matrix of m rows and n columns (i.e., $m \times n$), it generates p row clusters and q column clusters (where $p \leq m$ and $q \leq n$) by exploiting the mutual dependency (or *duality*) between the rows and the columns. The particular algorithm we used in this work is the *information-theoretic* co-clustering from [1], which reduces the dimensions of rows and columns simultaneously while minimizing the loss of the amount of Mutual Information (MI) contained in the matrix. Formally, the MI between two random variables x and y, denoted $I(x, y)$, is:

$$I(x, y) = log \frac{p(x, y)}{p(x)p(y)}$$

where $p(x, y)$ is the joint probability of x and y, and $p(x)$ and $p(y)$ are the (marginal) probability of x and y respectively. MI is symmetric, and indicates the mutual dependence between two random variables. In our case, $I(x, y)$ essentially indicates how well a context word is correlated with a given adjective, and vice versa. Thus by applying co-clustering to our data, we can identify the aspects and their qualities of games which the users care about, which might be different from other products or domains.

5 Results

In the current work, we specified to generate 100 row/context clusters and 30 column/adjective clusters (thus a total of 3,000 co-clusters). We chose those numbers of clusters after experimenting with various configurations and choosing the one which we thought gave the least tradeoff between the loss in MI and the reduction of the dimensions.

Table 1 shows some examples of notable co-clusters. Although the results yet have to be examined closely, we have observed that many high frequency co-clusters had context clusters which refer to the overall look and feel of games, for instance ⟨ graphics, look, sound, music,.. ⟩ and ⟨ game, overall, line, conclusion, qualities, .. ⟩ – which suggests that users care strongly about the aesthetics of the game and gameplay. On the other hand, co-clusters with context clusters containing concrete objects such as ⟨ map, gun, city, step, explosion, planet,.. ⟩ had relatively low frequencies – implying that whatever specific objects/props used in the game might not be so important.

Next, we represented all games using the derived 3,000 co-clusters as features and clustered (all 8,279) games. Games which were grouped in the same cluster by this process should share a similar pattern of characteristics on the aspects which the users strongly care about, for example games which are "overall-innovative", and have "great-graphics" but "difficult-control". To cluster games, we applied the standard K-means algorithm with K=30.

Evaluation of clustering are in general difficult because, unlike classification, there is no category to which the result can be measured for accuracy. Our work in this paper is particularly difficult because the clustering is based on users' subjective opinions rather than the objective item similarity.[1]

But for the present purpose, as a preliminary evaluation we manually inspected the clusters to see if any interesting patterns have emerged. Table 2

[1] Our ultimate goal, in the future work, is to use such game clusters in a game recommender system – to recommend similar games to the one(s) the user has already or is considering purchasing.

Table 1. Some Notable Adjective-Context Co-clusters

Frequency		Words
50,845	Adj	great, amazing, excellent, fantastic, incredible, decent, outstanding,
	Context	graphics, look, sound, music, voice, idea, job, soundtrack, acting,
42,502	Adj	basic, simple, smooth, innovative, interesting, linear, immersive, engaging,
	Context	game, overall, line, conclusion, qualities, summary, theory,
27,382	Adj	new, original, total, disappointing, additional, updated, expanded,
	Context	bit, whole, unlock, content, brand, introduce, dimension, disaster,
17,189	Adj	bad, hard, difficult, stupid, tough, mean, smart, sad, tricky,
	Context	control, learn, task, pull, achieve, navigate, maneuver,
5,796	Adj	unique, realistic, arcade, cinematic, sandbox, polished, enhanced,
	Context	effect, animation, visual, design, texture, presentation, cut, art,

shows examples of derived game clusters. A few interesting points were observed. First, it seems sports games are quite salient in some clusters (e.g. Cluster 3 and 5). The same for driving games, and they are often grouped together with sports games. This was somewhat surprising because the combination of those two genres/categories seems to cut across all sports, even with the ones which have different gameplay, such as traditional sports vs. other physical activities (e.g. wrestling, skateboarding, snowboarding), even including music games (of the kind played with peripherals). Second, strategy games seem to appear together often. Interestingly, they also seem to appear with First Person Shooter games as well (e.g. Cluster 2). Third, there were some clusters which grouped kids/teen games using established properties (Disney, Harry Potter, etc.). Also overall, sequels seem to be grouped together. This itself makes some sense, but we still have to look closer to find possible cases for sequels that didn't appear together (and investigate why).

6 Conclusions and Future Work

In this paper, we presented the results of extracting user opinions by applying co-clustering to an adjective-context co-occurrence matrix. From the derived co-clusters, we discovered that game users tend to care about the overall look and feel aspect of the game more than concrete elements used in the game. However, this result is only preliminary, obtained from just one experiment. In future work, we plan to do more experiments using other techniques to obtain results which could provide more, and different insights and discoveries on the user preferences in the game domain.

We must also do an analysis of the linguistic characteristics of the game domain, for example which adjectives modified which nouns more often, in conjunction with the sentiment. For example, the word "addictive" is generally used negatively, referring to the persistent and compulsive use of a substance known to

Table 2. Example Game Clusters

1	ds:TheLegendofZeldaPhantomHourglass, ds:NewSuperMarioBros, Xbox360:NinjaGaidenII, Xbox360:MassEffect, Wii:SuperMarioGalaxy, PS3:MetalGearSolid4GunsofthePatriots, PS2:FinalFantasyXII, PS2:DragonQuestVIII
2	ds:MetroidPrimeHunters, Xbox360:Quake4, Xbox360:Halo3, Xbox360:GearsofWar, PS3:UnrealTournament3, PS3:TomClancy'sRainbowSixVegas, PS2:Killzone, PC:StarWarsEmpireatWar, PC:AgeofEmpiresIII
3	ds:MarioHoops3on3, Xbox360:WWESmackDownvs.Raw2009, Xbox360:NHL08, Xbox:FIFASoccer06, Xbox360:MaddenNFL09, Wii:RockBand, PSP:RidgeRacer, PS3:TonyHawk'sProject8, PS3:NBA2K8, PS2:TonyHawk'sProSkater3
4	ds:PrinceofPersiaTheFallenKing, Xbox:FindingNemo, Wii:LooneyTunesAcmeArsenal, PS2:CuriousGeorge, PS2:Disney'sTarzanUntamed, PS2:HarryPotterandtheSorcerer'sStone, GameBoyAdvance:TheChroniclesofNarniaTheLion,TheWitchandTheWardrobe
5	ds:MajorLeagueBaseball2K7, Xbox360:NCAAMarchMadness08, Xbox360:NHL2K9 Xbox:UEFAEuro2004, Xbox:TonyHawk'sProSkater4, Xbox:TigerWoodsPGATour07 Wii:ProEvolutionSoccer2008, PSP:MaddenNFL08, PS2:NASCARThunder2004

be harmful. But in the domain of games, it is often used positively. We are planning on examining the adjective-context word pairs to focus on specific words, and investigate the language of the game domain.

Game clustering also needs much more work. In addition to further experimentation with clustering (by using other algorithms, also with various parameter values), we plan to conduct rigorous evaluation of both qualitative and quantitative results. Finally, as mentioned earlier, we are planning to develop a game recommender system using the derived game clusters, and test the system with real users.

References

1. Archakn, N. Ghosea, A., Ipeirotis, P.: Show me the Money! Deriving the Pricing Power of Product Features by Mining Consumer Reviews. In: Proceedings of the 13th ACM SIGKDD international conference on Knowledge Discovery and Data Mining (KDD-07), 56–65. (2007)
2. Bisson, G., Hussain, F.: Chi-Sim: A New Similarity Measure for the Co-clustering Task. In: Proceeding of the 7th International Conference on Machine Learning Applications (ICMLA-08), 211–217. (2008)
3. Dhillon, I. S., Mallela, S., Modha, D. S.: Information-Theoretic Co-clustering. In: Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), 89–98. (2003)
4. George, T., Merugu, S.: A Scalable Collaborative Filtering Framework Based on Co-clustering. In: Proceedings of the 5th IEEE International Conference on Data Mining (ICDM-05), 625–628. (2005)
5. Hartigan, J. A.: Direct Clustering of a Data Matrix. In: Journal of the American Statistical Association, 67 (337), 123-129. (1972)

6. Hatzivassiloglou, V., McKeown, K.: Predicting the semantic orientation of adjectives. In: Proceedings of the 8th conference on European chapter of the Association for Computational Linguistics (EACL-97), 174–181. (1997)
7. Pang, B., Lee, L.: Opinion mining and sentiment analysis. In: Foundations and Trends in Information Retrieval, 2 (1-2), 1-135. (2008)
8. Popescu, A., Etzioni, O.: Extracting Product Features and Opinions from Reviews. In: Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP-05), 339–346. (2005)
9. Stylios, G., Christodoulakis, D., Besharat, J., Kotrotsos, I., Koumpouri,A., Stamou, S.: Public Opinion Mining for Governmental Decisions. In: Electronic Journal of Electronic Government, 8 (2), 202–213. (2011)
10. Thomas, M., Pang, B., Lee, L.: Get out of the vote: Determining support or opposition from Congressional Floor-Debate Transcripts. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-06), 327–335. (2006)
11. Zagal, J. P., Tomuro, N.: The Aesthetics of Gameplay: A Lexical Approach. In: Proceedings of the 14th International Academic Mindtrek Conference, 9–16. (2010)
12. Zagal, J. P., Tomuro, N., Shepitsen, A.: Natural Language Processing for Games Studies Research. In: Journal of Simulation & Gaming (S&G), Special Issue on Games Research Methods, 43 (3), 353 - 370. (2011)
13. Zhan, T-J., Li, C-H.: Semantic dependent word pairs generative model for fine-grained product feature mining. In: Proceedings of the 15th Pacific-Asia conference on Advances in knowledge discovery and data mining (PAKDD-11), 460–475. (2011)