

PerspectiveSpace: Opinion Modeling with Dimensionality Reduction

Jason B. Alonso¹, Catherine Havasi², and Henry Lieberman³

¹ MIT Media Laboratory, Personal Robots
20 Ames St E15-468a, Cambridge MA 02139, USA
jalonso@media.mit.edu

² Brandeis University, Lab for Linguistics and Computation,
415 South Street, Waltham, MA 02454
havasi@cs.brandeis.edu

³ MIT Media Laboratory, Software Agents
20 Ames St E15-384a, Cambridge MA 02139, USA
lieber@media.mit.edu

Abstract. Words mean different things to different people, and capturing these differences is often a subtle art. These differences are often “a matter of perspective”. Perspective can be taken to be the set of beliefs held by a person as a result of their background, culture, tastes, and experience. But how can we represent perspective computationally?

In this paper, we present PerspectiveSpace, a new technique for modeling spaces of users and their beliefs. PerspectiveSpace represents these spaces as a matrix of users, and data on how people agree or disagree on assertions that they themselves have expressed. It uses Principal Component Analysis (PCA) to reduce the dimensionality of that matrix, discovering the most important axes that best characterize the space. It can then express user perspectives and opinions in terms of these axes. For recommender systems, because it discovers patterns in the beliefs about items, rather than similarity of the items or users themselves, it can perform more nuanced categorization and recommendation. It integrates with our more general common sense reasoning technique, AnalogySpace, which can reason over the content of expressed opinions.

An application of PerspectiveSpace to movie recommendation, 2-wit, is presented. A leave-one-out test shows that PerspectiveSpace captures the consistency of users’ opinions very well. The technique also has applications ranging from discovering subcultures in a larger society, to building community-driven web sites.

1 Finding Perspective

The variations in people’s beliefs and personalities lie at the heart of many common problems where people are trying to make use of information subject to opinion. People look at reviews of movies and products, but people don’t always think the same way about them. An online forum of many users is often strewn with many disagreements, and users have difficulty navigating them to find the

useful commentary amongst the noise. This is especially true for newcomers to a forum, where the reputations of the regular contributors are unknown to entrants. What is needed here is a tool that lets people express themselves honestly and then captures even the subtle differences between people and the opinions in a conversation in a meaningful way.

PerspectiveSpace is an analysis of person-to-person interactions that explores the similarities and differences in what people believe by discovering descriptive axes on which people can be arranged. These belief patterns underlie the different “perspectives” that people may have, which can be taken to be the set of beliefs held by a person as a result of his or her background, culture, tastes, and experience. In addition to studying the varying beliefs of different social or cultural groups, PerspectiveSpace has applications in recommender systems in that it utilizes knowledge about how people think about the items being recommended.

2 Related Work

2.1 Common Sense

The roots of the PerspectiveSpace project lie in research into collecting and applying common sense knowledge. Open Mind Common Sense (OMCS)[1] is a project that seeks to collect a large body of common sense knowledge in natural language from volunteer contributors over the Internet. ConceptNet[2] is a semantic network designed to be a machine-usable representation of the corpus of knowledge captured by the OMCS project. The nodes of the semantic network are normalized strings of natural language, called “concepts,” and these concepts are interconnected with labeled directed links.

AnalogySpace [3] is a new reasoning technique designed to work well on common sense knowledge, which is often vague, redundant, or inconsistent. Rather than compute logical truth, it is oriented towards computing less stringent notions such as similarity, plausibility, analogy, or position along a descriptive spectrum. AnalogySpace represents a space of knowledge as a matrix of concepts vs. features of these concepts. It uses the technique of Singular Value Decomposition (SVD) to reduce the dimension of this matrix, and find the most important descriptive axes that best account for variation in the data. AnalogySpace is computationally very efficient, since the SVD can be computed in advance, and semantic notions such as similarity can be easily computed at run time with basic vector mathematics.

The key insight that led us from AnalogySpace to PerspectiveSpace is that users of a system can be represented as an extra dimension in this matrix. Different users may or may not agree with any given assertion, which allows patterns in sets of beliefs or sets of users to emerge from the mathematical analysis. We observed that these patterns could be used to capture some notion of perspective, as it is generally understood in user modeling. Details of the mathematics, and the relationship between AnalogySpace and PerspectiveSpace, are further explained below.

2.2 Recommender Systems

Recommender systems are tools that look at the behavior (like purchasing activity) or direct input (like movie ratings) of a user to make an informed recommendation of content, products, or other entities in which the user would likely take an interest.

Recommender systems generally come in two types. First, those that are based on the direct characteristics of the product or item to be recommended. These represent each item as a feature vector, with features categorized in advance, and look for similarity between the user’s vector of preferences, and individual items.

A second category of recommender systems is collaborative filtering, which is based on similarity between people, according to their history of selected items, and recommend other items selected by similar users. A notable example of collaborative filtering is the Tapestry project from Xerox PARC [4]. A more recent example that uses SVD is given in [5]. A variant, feature-guided collaborative filtering, represents the users and items using predefined feature vectors, as for the case of recommenders based on item characteristics.

PerspectiveSpace also can group users according to their similarity, but with greater sensitivity to subtle differences in beliefs. Rather than being based on their direct selection of items, it uses their beliefs about these items, and can discover emergent patterns of these beliefs. Different users may select the same item for different reasons, or users who have similar beliefs may nonetheless select different items. Thus PerspectiveSpace operates at a finer grain than traditional collaborative filtering. It holds the potential for more nuanced recommendation, and can deliver better insight into why these recommendations were chosen.

2.3 Sentiment Analysis

Opinion mining and sentiment analysis are disciplines very closely related to PerspectiveSpace. Morinaga et al. even used PCA to find associations between products and natural-language opinion terms that describe them [6]. A recent survey of these fields [7], however, makes it very clear that the field focuses on finding a general affinity score for a given product using information extracted (mined) from free text. In particular, emphasis is given to reducing the semantics of free text to either a “positive” or “negative” opinion. The survey discusses treatment of reviewer reliability, but that treatment appears perfunctory and is secondary to establishing the affinity score. Furthermore, there does not appear to be much treatment of characterizing and differentiating reviewers as is typically done with recommender systems.

3 Methodology

3.1 From AnalogySpace to PerspectiveSpace

AnalogySpace, mentioned above, is a transformation of the scored assertions (concepts crossed with features, giving scores) in ConceptNet that yields a

compressed vector for each concept or feature, permitting elementary linear operations to be used to perform calculations on semantic similarity [3]. PerspectiveSpace, which is separate from but related to AnalogySpace, is a transformation of the ratings people assign to statements that gives a compressed vector for each statement and person. The axes of these vectors represent significant variations which can be used to characterize different subcultures.

There are a couple notable properties of the SVD that are important to PerspectiveSpace. One is that the principle components found (termed “axes” in the lingo of AnalogySpace and PerspectiveSpace) are ordered in decreasing significance, and that degree of significance is measured by its singular value. The most significant axis divides the data into the two most divisive sets of items and their properties, summarizing groups of properties with a single varying parameter. In the case of PerspectiveSpace, this varying parameter tends to be either a divisive issue or a set of beliefs neither held nor contradicted by groups of people not falling on the axis. Subsequent axes divide and describe the data along successively less significant but empirically independent parameters. Each of these varying parameters can be taken to describe a group of properties holistically. Another important property is that the discovered axes are orthogonal, which means that each successively less significant axis describes successively more subtle variations in the data. In appropriate circumstances, the most subtle variations can be taken to be noise, which can then be removed to make sensible interpolations of missing data.

3.2 Creating a PerspectiveSpace

As an example for motivating many of the computation steps in calculating PerspectiveSpace, we introduce 2-wit, a recommender system that recommends reviews of products in the consumer market (in this example, movies), which is distinctly different from the traditional approach of recommending products themselves. First, we will discuss the initial data set and matrix for a PerspectiveSpace analysis, then the normalization and decomposition.

Collecting the Data. The collection of data for PerspectiveSpace must have four features:

1. contributors must be able to express their beliefs in succinct, natural language assertions;
2. agreement and disagreement between contributors on their assertions must be readily and frequently ascertained;
3. each contribution must be linked to the identity of the contributor; and
4. each contributor may only issue one rating per assertion, though that rating may be altered over time.

It should be noted that assertions are, in general, structured with “tags” and “statements.” A tag is a simple indication of what the contributor thinks about the statement. In the simplest of implementations, the only tag available is “agree,” which indicates whether or not the contributor believes the statement is

Movie: WALL-E (2008)

Thu, 17 July 2008 @ 4:33 p.m. EDT 6 reviews

Blank description

kcarnold thinks...

Fri, 18 July 2008 @ 4:51 p.m. EDT 4 votes
 Thinkerprint: aecI

Asks apt questions about what being human means. Illustrates a dystopia of overconsumption, lack of environmental stewardship, and poor leadership. The ending (a back-to-Eden restoration) is tacked-on and not credible. But overall it's very well-executed: good story-telling, nice visuals, and the combination of kid and adult appeal that Pixar has become known for. The opening short is cute, too.

What do you think?

...or so I think

jalonso thinks...

Thu, 17 July 2008 @ 4:36 p.m. EDT 4 votes
 Thinkerprint: adjE

A cute film about finding one's purpose.

Fig. 1. A screenshot of the 2-wit collection interface

true. In the case of 2-wit, each statement is a review made of a particular movie. The tags in 2-wit are “agree,” for indicating that the contributor either agrees (positive rating) or disagrees (negative rating) with a statement, and “junk” for indicating that the contributor believes the statement is obscene, spam, or otherwise generally useless in the opinion of the contributor.

In the case of 2-wit, the interface (see figure 1) presents logged-in users with a simple box for entering one or more reviews or otherwise short commentaries on a movie. All of the reviews entered by other users are visible for consideration, and there are three icons at the top of each review for users to click to express their opinion of each review: a green thumbs-up (agree), a red thumbs-down (disagree), and a yellow flag (junk).

Constructing the matrix. The construction of PerspectiveSpace begins with the preparation of a ratings matrix, which is denoted as M_R in this paper.

The rows of M_R are person-tag pairs, its columns are assertions or statements, and its values are ratings. M_R is also sparse, which is an important consideration for the normalization methods discussed in section 3.2, and it influences the choice of SVD implementation. The sparse matrix is populated such that $M_R[(i, k), j]$ is positive if and only if person i gave statement j with tag k a rating of agreement, negative with a rating of disagreement, and zero with no rating or a neutral rating.

By coupling tags with people in this manner, an SVD can detect patterns where one group of people tend to make or believe statements that another group tends to believe are junk (or otherwise assign a different tag). As such, the matrix describes the statements (columns) in terms of what people think of them with rows like “Alice agrees” and “Bob thinks is junk”.

Normalization. The matrix M_R is normalized to obtain \hat{M}_R . In the context of PerspectiveSpace, this process reshapes the data set going into the SVD, while altering its semantics as little as possible, so as to maximize the effectiveness of the algorithm.

Sparse mean-shifting. PCA methods generally require that zero-mean vectors be prepared before performing an SVD[8]. Given that PerspectiveSpace data sets are usually sparse, some attention must be paid to the shifting of the terms, as the sparse SVD methods used do not permit the use of nonzero unpopulated entries. The result is that the shifts take only non-zero columns into account.

This approach, unfortunately, places undue emphasis on negative ratings. In the OMCS corpus, for example, 75% to 80% of the assertions are considered true when reviewed by human judges, and so most of the ratings made by contributors are ratings of agreement. This means that very few ratings by a person, if any, are negative. Mean shifting thus increases the magnitude of the negative ratings significantly, while diminishing the magnitude of the many positive ratings. Possibly even more damaging is that people who have only rated things positively will have no non-zero values in their rows of \hat{M}_R .

What is needed is another way to bring the mean rating per person to zero without introducing an imbalance.

Unity magnitude. Magnitude normalization methods in PerspectiveSpace were inherited from the methods used in constructing AnalogySpace. Though they do not accomplish the zero-mean property typically used with PCA, they improve the quality of the discovered axes by preventing the most significant axes from being dominated by the most populated rows and columns of the input matrix. When AnalogySpace is prepared without magnitude normalization, the most significant axes simply described the concepts for which OMCS had the most data[3]. In terms of 2-wit, a user who rates 100 reviews would overpower the user who rates 10 reviews in the data set, such that the user who rated 100 reviews would likely establish an axis unto himself/herself. Unity magnitude normalization rescales the input values, so users can readily establish their positions in PerspectiveSpace without simply granting the most prolific users their own axes. This is important, as the data collection process places no constraints on the relative number of ratings obtained from each person.

An immediately apparent limitation of unity magnitude normalization is that rows of the input matrix with very little content, which would describe a person (in PerspectiveSpace) who contributed a single rating, would have just as much influence in the formation of axes as rows with a lot of content. The solution adopted in AnalogySpace, and subsequently adopted in PerspectiveSpace, was to add a constant term in each row to make the magnitude of the normalized row vary with the magnitude of the initial row.

This normalization model works as follows, with b as the “base parameter” to be added to each row to diminish insignificant rows and \mathbb{P}_A as the set of assertions (that is, all of the columns of M_R):

$$\hat{M}_R[i, j] = \frac{M_R[i, j]}{\sqrt{b + \sum_{j'' \in \mathbb{P}_A} M_R[i, j'']^2}} \quad (1)$$

Complementary ratings. Under the complementary (or “mirrored”) ratings model, every assertion was matched with a complementary, opposite assertion. That is, for every rating a person gives a normal assertion, the person is modeled to give the opposite rating to the complementary assertion. This ensures that the average rating given by a user is always zero without placing undue bias on any particular assertion. As such, the artifacts that appear as a result of using unity magnitude normalization methods disappears.

This approach can be taken to the logical extreme by having complementary users in addition to complementary assertions.

Combined magnitude normalization with mirrored ratings. The final normalization method chosen for 2-wit, and recommended for any PerspectiveSpace preparation, was a combination of mirroring and unity magnitude normalization. Mirroring is applied first to ensure that the average value in each row of the resultant matrix is zero. Applying unity magnitude normalization (with $b = 2.0$) does not alter the zero-mean property of the matrix.

Performing the decomposition. Once a normalized matrix is obtained, a singular value decomposition (SVD) is taken using the Lanczos algorithm, yielding U , Σ , and V^T

$$\hat{M}_R \approx U \Sigma V^T \quad (2)$$

U gives the coordinates of each user in PerspectiveSpace, while V gives the coordinates of each assertion in the same space.

4 2-Wit

2-wit, whose name is a contraction of “What Will I Think,” is a web-based system that collects opinions people have about movies and suggests, to each user, opinions from other users that he or she might agree with.

The 2-wit implementation, used as an example above, was designed to give people a social space where they could share their opinions on the movies they’ve seen, with the general expectation that they could also rely on each other to help make decisions on movies to see in the future. Though the interface was designed to emphasize the social nature of the system, it is not a social networking system that allows people to declare “friends” or “groups” as contemporary social networking sites like Facebook permit.

The movie review domain was chosen as for any two people in the same society, there are likely to be at least a few movies that they have both seen. At the same time, people’s tastes and opinions of movies are sufficiently varied to make a range of trends readily detectable. As such, even a small number of contributors was expected to produce usable results.

4.1 Using 2-Wit to Evaluate PerspectiveSpace

A variation of the leave-one-out test was performed to test the interpolative abilities of PerspectiveSpace in the 2-wit implementation. The presumed usage scenario for the 2-wit movie system is that a user considering movies that he or she has not seen would search for reviews that he or she would agree with. With this in mind, the traditional leave-one-out test was extended to leave out one user-movie pair per trial. As such, all ratings a particular user gave for reviews of a particular movie were removed for the duration of the trial, and recall was defined in terms of the number of reviews for which the user’s rating could be properly estimated. The test, accordingly, is dubbed “leave-some-out.”

In determining viable user-movie pairs for the leave-some-out test, the test skipped the degenerate case of reviews where only one user provided a rating in trials testing the applicable user and the relevant movie. Finally, the case of users who commented on only one movie was also avoided, as there were no grounds for an SVD to interpolate ratings for such users.

Given these restrictions, it was determined that exactly 61 user-movie pairs were viable as trials for the evaluation, and so the evaluation covers all possible user-movie pairs rather than a random subset. This test was also repeated with varying values of k , which determines the number of axes that should be computed and used for the analysis.

4.2 Results

The results are shown in table 1. For the purposes of this evaluation, “hits” are defined as the number of trials in which the user’s rating of a movie review was successfully interpolated by taking the dot product of the PerspectiveSpace vectors for the person and the review (a process called perspective projection) with a threshold $\tau = 0.01$. “Misses” are defined similarly where the rating was predicted opposite to the the correct value. “Undecided” is defined as the number of reviews for which the magnitude of the predictive rating was smaller than τ , where there is insufficient confidence to make a reasonable estimate of a user’s acceptance or rejection of the movie review. The “% decided” represents the fraction of reviews for which a confident estimation was made, while “% correct” represents the portion of *the decided reviews* for which the correct estimation was made.

The results are positive in that 2-wit performed substantially better than chance in estimating the opinions that users would have about movies. Given any user-movie pair for testing, the leave-some-out test appropriately excluded all reviews the user made of any assertion about the movie. It is further interesting to note that the best performance was obtained with $k = 4$ —a reasonable interpretation of this phenomenon is that the omission of lower-ranking axes removed noise from the data set, particularly excluding “information” that specifically worked to diminish properly-interpolated ratings for which there was no direct measurement.

Table 1. Results of the 2-wit evaluation with $\tau = 0.01$

Tags	k	Cumulative				Average		
		Hits	Misses	Undecided	% decided	% correct	% decided	% correct
all	4	215	74	88	77	74	76	78
	8	215	74	88	77	74	78	74
	16	150	81	146	61	65	63	66
agree	4	93	46	43	76	67	76	71
	8	92	42	48	74	69	75	69
	16	58	45	79	57	56	60	59
junk	4	122	28	44	77	81	76	84
	8	123	32	39	80	79	81	79
	16	92	36	66	66	72	65	71

5 Other Applications Using PerspectiveSpace

PerspectiveSpace has also been used in other applications. It was used by the Common Sense Computing Initiative to identify and combat malicious or disruptive users. When contributors rate statements for truth, such as in OMCS, one can find the class or classes of users which tend to be at odds with more trusted users. This is also useful for discovering when a user has created multiple accounts which he uses only to agree with his primary account or bolster his ratings[9].

SlantExplorer, an application prototype, is a web-based interface for navigating the conflicting opinions that underlie or are otherwise applicable to a document. SlantExplorer is designed as a tool for composing expository documents or for assisting users trying to draft document summaries. By considering the perspectives that people can take on a document, SlantExplorer can help identify portions of a document of interest to people belonging to particular groups or backgrounds [9].

6 Perspectives for Understanding an Increasingly Connected Society

In addition to recommender systems, the potential application domains for PerspectiveSpace fall into four major categories: characterizing societies, detecting microtheories, and community-driven content.

6.1 Characterizing Societies

PerspectiveSpace offers many opportunities for the identification and study of behavior pattern within a larger community. PerspectiveSpace can tie into the AnalogySpace family of language analysis tools and together they offers a means for the systematic study of jargon usage, dialects, and belief patterns related to culture or subculture. In more general terms, PerspectiveSpace is a tool for opinion analysis, which has direct applications in marketing and political settings.

With care, these “opinions” can be applied toward understanding the language of particular groups of people (or characterizing people by the language they use) in natural language processing applications. A simple example of this application would be useful in common sense reasoning. Since people of different countries and cultures speak the same language, it would be useful to isolate which users were representative of which cultures.

Natural language processing applications can use PerspectiveSpace to help model the language patterns of people dynamically by generating perspectives representative of a speaker from the language he or she uses, including jargon and idioms. If the PerspectiveSpace were constructed from a semantic resource, these perspectives could be used in later processing steps to aid with understanding the semantics underlying a statement. Given, for example sentences like “Soda is a carbonated beverage” and “Pop is a carbonated beverage,” a system can learn that a particular group of people using the word “pop” might be significantly more likely than other groups to mean a carbonated beverage rather than a small explosive sound.

6.2 Pseudo-microtheories

Microtheories are consistent subsets of larger bodies of assertions, usually descriptive of a specific domain of knowledge. They are usually understood and discussed in the context of formal reasoning systems: the Cyc project, for example, is a body of common sense knowledge professionally-crafted into the form of formal, logical assertions[10].

Principle perspectives, by virtue of their discovery process, tend to identify assertions that are believed or disbelieved in tandem by people. This correlation between assertions appears to imply that the assertions are consistent and related to each other, which make these assertions like microtheories. We hesitate to call these assertions “microtheories” as this discovery process does not represent assertions in predicate logic or show how one assertion can be derived from others, and these are expected in contemporary notions of “microtheories.”

6.3 Community-Driven Content

PerspectiveSpace offers a variety of interesting possibilities for applications that use community-driven content, which share the feature that people can contribute to and navigate a body of knowledge on which not every user would agree. PerspectiveSpace has been tested using the Open Mind Common Sense project[9] and could be utilized in other websites such as Slashdot[11] and Wikipedia[12].

The use of perspective projection can allow contributors and browsers of a community-driven content project to work in a space of assertions that are compatible with their particular belief patterns. For a contributor, this would allow him or her to build upon existing statements in greater detail and comfort. For a browser, this would allow him or her to explore content in a self-consistent form.

Subculture detection, on the other hand, can help a contributor or browser identify and understand the major sides of an argument as well as the prevailing agreements underlying a discussion.

Using perspective projection and subculture detection, both contributors and browsers can easily explore different perspectives. PerspectiveSpace gives them the tools they need to consider counter-arguments and differing opinions.

References

1. Singh, P.: The public acquisition of commonsense knowledge. In: Proceedings of AAAI Spring Symposium: Acquiring (and Using) Linguistic (and World) Knowledge for Information Access. AAAI, Menlo Park (2002)
2. Havasi, C., Speer, R., Alonso, J.: Conceptnet 3: a flexible, multilingual semantic network for common sense knowledge. In: Recent Advances in Natural Language Processing (2007)
3. Speer, R., Havasi, C., Lieberman, H.: AnalogySpace: Reducing the dimensionality of common sense knowledge. In: Proceedings of AAAI 2008 (July 2008)
4. Goldberg, D., Nichols, D., Oki, B.M., Terry, D.: Using collaborative filtering to weave an information tapestry. *Commun. ACM* 35(12), 61–70 (1992)
5. Sarwar, B.M., Karypis, G., Konstan, J.A., Reidl, J.T.: Application of dimensionality reduction in recommender system—a case study. In: WebKDD 2000 Workshop (2000)
6. Morinaga, S., Yamanishi, K., Tateishi, K., Fukushima, T.: Mining product reputations on the web. In: KDD 2002: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 341–349. ACM, New York (2002)
7. Pang, B., Lee, L.: Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* 2(1-2), 1–135 (2008)
8. Smith, L.: A tutorial on Principal Components Analysis, vol. 51, p. 52. Cornell University, USA (2002)
9. Alonso, J.B.: PerspectiveSpace. S.M. thesis, Massachusetts Institute of Technology. Program in Media Arts and Sciences (2008)
10. Lenat, D.: Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM* 11, 33–38 (1995)
11. SourceForge, Inc.: Slashdot: News for nerds, stuff that matters. Slashdot web site (2008), <http://slashdot.org/>
12. Wikimedia Foundation, Inc.: Wikipedia, the free encyclopedia. English language Wikipedia web site (2008), <http://en.wikipedia.org/>