

Automatic Seed Word Selection for Unsupervised Sentiment Classification of Chinese Text

Taras Zagibalov John Carroll

University of Sussex

Department of Informatics

Brighton BN1 9QH, UK

{T.Zagibalov, J.A.Carroll}@sussex.ac.uk

Abstract

We describe and evaluate a new method of automatic seed word selection for unsupervised sentiment classification of product reviews in Chinese. The whole method is unsupervised and does not require any annotated training data; it only requires information about commonly occurring negations and adverbials. Unsupervised techniques are promising for this task since they avoid problems of domain-dependency typically associated with supervised methods. The results obtained are close to those of supervised classifiers and sometimes better, up to an F_1 of 92%.

1 Introduction

Automatic classification of document sentiment (and more generally extraction of opinion from text) has recently attracted a lot of interest. One of the main reasons for this is the importance of such information to companies, other organizations, and individuals. Applications include marketing research tools that help a company see market or media reaction towards their brands, products or services, or search engines that help potential purchasers make an informed choice of a product they want to buy. Sentiment classification research has drawn on and contributed to research in text classification, unsupervised machine learning, and cross-domain adaptation.

This paper presents a new, automatic approach to automatic seed word selection as part of sentiment classification of product reviews written in Chinese, which addresses the problem of do-

main-dependency of sentiment classification that has been observed in previous work. It may also facilitate building sentiment classification systems in other languages since the approach assumes a very small amount of linguistic knowledge: the only language-specific information required is a basic description of the most frequent negated adverbial constructions in the language.

The paper is structured as follows. Section 2 surveys related work in sentiment classification, unsupervised machine learning and Chinese language processing. Section 3 motivates our approach, which is described in detail in Section 4. The data used for experiments and baselines, as well as the results of experiments are covered in Section 5. Section 6 discusses the lessons learned and proposes directions for future work.

2 Related Work

2.1 Sentiment Classification

Most work on sentiment classification has used approaches based on supervised machine learning. For example, Pang *et al.* (2002) collected movie reviews that had been annotated with respect to sentiment by the authors of the reviews, and used this data to train supervised classifiers. A number of studies have investigated the impact on classification accuracy of different factors, including choice of feature set, machine learning algorithm, and pre-selection of the segments of text to be classified. For example, Dave *et al.* (2003) experiment with the use of linguistic, statistical and n-gram features and measures for feature selection and weighting. Pang and Lee (2004) use a graph-based technique to identify and analyze only subjective parts of texts. Yu and Hatzivassiloglou (2003) use semantically-oriented words for identification of polarity at the sentence level. Most of this work assumes binary classification (positive and negative), some-

© 2008. Licensed under the *Creative Commons Attribution-NonCommercial-Share Alike 3.0 Unported* license (<http://creativecommons.org/licenses/by-nc-sa/3.0/>). Some rights reserved.

times with the addition of a neutral class (in terms of polarity, representing lack of sentiment).

While supervised systems generally achieve reasonably high accuracy, they do so only on test data that is similar to the training data. To move to another domain one would have to collect annotated data in the new domain and retrain the classifier. Engström (2004) reports decreased accuracy in cross-domain classification since sentiment in different domains is often expressed in different ways. However, it is impossible in practice to have annotated data for all possible domains of interest. Aue and Gamon (2005) attempt to solve the problem of the absence of large amounts of labeled data by customizing sentiment classifiers to new domains using training data from other domains. Blitzer *et al.* (2007) investigate domain adaptation for sentiment classifiers using structural correspondence learning. Read (2005) also observed significant differences between the accuracy of classification of reviews in the same domain but published in different time periods.

Recently, there has been a shift of interest towards more fine-grained approaches to processing of sentiment, in which opinion is extracted at the sentence level, sometimes including information about different features of a product that are commented on and/or the opinion holder (Hu and Liu, 2004; Ku *et al.*, 2006). But even in such approaches, McDonald *et al.* (2007) note that information about the overall sentiment orientation of a document facilitates more accurate extraction of more specific information from the text.

2.2 Unsupervised Approach

One way of tackling the problem of domain dependency could be to use an approach that does not rely on annotated data. Turney (2002) describes a method of sentiment classification using two human-selected seed words (the words *poor* and *excellent*) in conjunction with a very large text corpus; the semantic orientation of phrases is computed as their association with the seed words (as measured by pointwise mutual information). The sentiment of a document is calculated as the average semantic orientation of all such phrases.

Yarowsky (1995) describes a 'semi-supervised' approach to the problem of sense disambiguation of words, also using a set of initial seeds, in this case a few high quality sense annotations. These annotations are used to start an iterative process of learning information about the contexts in which senses of words appear, in

each iteration labeling senses of previously unlabeled word tokens using information from the previous iteration.

2.3 Chinese Language Processing

A major issue in processing Chinese text is the fact that words are not delimited in the written language. In many cases, NLP researchers working with Chinese use an initial segmentation module that is intended to break a text into words. Although this can facilitate the use of subsequent computational techniques, there is no a clear definition of what a 'word' is in the modern Chinese language, so the use of such segmenters is of dubious theoretical status; indeed, good results have been reported from systems which do not assume such pre-processing (Foo and Li, 2004; Xu *et al.*, 2004).

2.4 Seed Word Selection

We are not aware of any sentiment analysis system that uses unsupervised seed word selection. However, Pang *et al.* (2002) showed that it is difficult to get good coverage of a target domain from manually selected words, and even simple corpus frequency counts may produce a better list of features for supervised classification: human-created lists resulted in 64% accuracy on a movie review corpus, while a list of frequent words scored 69%. Pang *et al.* also observed that some words without any significant emotional orientation were quite good features: for example, the word "still" turned out to be a good indicator of positive reviews as it was often used in sentences such as "Still, though, it was worth seeing".

3 Our Approach

Our main goal is to overcome the problem of domain-dependency in sentiment classification. Unsupervised approaches seem promising in this regard, since they do not require annotated training data, just access to sufficient raw text in each domain. We base our approach on a previously described, 'almost-unsupervised' system that starts with only a single, human-selected seed 好 (*good*) and uses an iterative method to extract a training sub-corpus (Zagibalov & Carroll, 2008). The approach does not use a word segmentation module; in this paper we use the term 'lexical item' to denote any sequence of Chinese characters that is treated by the system as a unit, whatever it is linguistically — a morpheme, a word or a phrase.

Our initial aim was to investigate ways of improving the classifier by automatically finding a better seed, because Zagibalov & Carroll indicate that in different domains they could, by manual trial and error, find a seed other than 好 (*good*) which produced better results.

To find such a seed automatically, we make two assumptions:

1. Attitude is often expressed through the negation of vocabulary items with the opposite meaning; for example in Chinese it is more common to say *not good* than *bad* (Tan, 2002). Zagibalov & Carroll's system uses this observation to find negative lexical items while nevertheless starting only from a positive seed. This leads us to believe that it is possible to find candidate seeds themselves by looking for sequences of characters which are used with negation.
2. The polarity of a candidate seed needs to be determined. To do this we assume we can use the lexical item 好 (*good*) as a gold standard for positive lexical items and compare the pattern of contexts a candidate seed occurs in to the pattern exhibited by the gold standard.

Looking at product review corpora, we observed that *good* is always more often used without negation in positive texts, while in negative texts it is more often used with negation (e.g. *not good*). Also, *good* occurs more often in positive texts than negative, and more frequently without negation than with it. We use the latter observation as the basis for identifying seed lexical items, finding those which occur with negation but more frequently occur without it.

As well as detecting negation¹ we also use adverbials² to avoid hypothesizing non-contentful seeds: the characters following the sequence of a negation and an adverbial are in general contentful units, as opposed to parts of words, function words, etc. In what follows we refer to such constructions as negated adverbial constructions.

¹We use only six frequently occurring negations: 不 (*bu*), 不会 (*buhui*), 没有 (*meiyou*), 摆脱 (*baituo*), 免去 (*mianqu*), and 避免 (*bimian*). We are trying to be as language-independent as possible so we take a simplistic approach to detecting negation.

²We use five frequently occurring adverbials: 很 (*hen*), 非常 (*feichang*), 太 (*tai*), 最 (*zui*), and 比较 (*bijiao*). Similarly to negation, we deliberately take a simplistic approach.

4 Method

We use a similar sentiment classifier and iterative retraining technique to the almost-unsupervised system of Zagibalov & Carroll (2008), summarized below in Sections 4.2 and 4.3. The main new contributions of this paper are techniques for automatically finding the seeds from raw text in a particular domain (Section 4.1), and for detecting when the process should stop (Section 4.4). This new system therefore differs from that of Zagibalov & Carroll (2008) in being completely unsupervised and not depending on arbitrary iteration limits. (The evaluation also differs since we focus in this paper on the effects of domain on sentiment classification accuracy).

4.1 Seed Lexical Item Identification

The first step is to identify suitable positive seeds for the given corpus. The intuition behind the way this is done is outlined above in Section 3. The algorithm is as follows:

1. find all sequences of characters between non-character symbols (i.e. punctuation marks, digits and so on) that contain negation and an adverbial, split the sequence at the negation, and store the character sequence that follows the negated adverbial construction;
2. count the number of occurrences of each distinct sequence that follows a negated adverbial construction (X);
3. count the number of occurrences of each distinct sequence without the construction (Y);
4. find all sequences with $Y - X > 0$.

4.2 Sentiment Classification

This approach to Chinese language processing does not use pre-segmentation (in the sense discussed in Section 2.3) or grammatical analysis: the basic unit of processing is the 'lexical item', each of which is a sequence of one or more Chinese characters excluding punctuation marks (so a lexical item may actually form part of a word, a whole word or a sequence of words), and 'zones', each of which is a sequence of characters delimited by punctuation marks.

Each zone is classified as either positive or negative based whether positive or negative vocabulary items predominate. As there are two parts of the vocabulary (positive and negative), we correspondingly calculate two scores (S_i ,

where i is either *positive* or *negative*) using Equation (1), where L_d is the length in characters of a matching lexical item (raised to the power of two to increase the significance of longer items which capture more context), L_{phrase} is the length of the current zone in characters, S_d is the current sentiment score of the matching lexical item (initially 1.0), and N_d is a negation check coefficient.

$$S_i = \frac{L_d^2}{L_{phrase}} S_d N_d \quad (1)$$

The negation check is a regular expression which determines if the lexical item is preceded by a negation within its enclosing zone. If a negation is found then N_d is set to -1 .

The sentiment score of a zone is the sum of sentiment of all the items found in it.

To determine the sentiment orientation of the whole document, the classifier computes the difference between the number of positive and negative zones. If the result is greater than zero the document is classified as positive, and vice versa.

4.3 Iterative Retraining

Iterative retraining is used to enlarge the initial seed vocabulary into a comprehensive vocabulary list of sentiment-bearing lexical items. In each iteration, the current version of the classifier is run on the input corpus to classify each document, resulting in a training subcorpus of positive and a negative documents. The subcorpus is used to adjust the scores of existing positive and negative vocabulary items and to find new items to be included in the vocabulary.

Each lexical item that occurs at least twice in the corpus is a candidate for inclusion in the vocabulary list. After candidate items are found, the system calculates their relative frequencies in both the positive and negative parts of the current training subcorpus. The system also checks for negation while counting occurrences: if a lexical item is preceded by a negation, its count is reduced by one.

For all candidate items we compare their relative frequencies in the positive and negative documents in the subcorpus using Equation (2).

$$difference = \frac{|F_p - F_n|}{(F_p + F_n)/2} \quad (2)$$

If $difference < I$, then the frequencies are similar and the item does not have enough distinguishing power, so it is not included in the vocabulary. Otherwise the sentiment score of the item is (re-)calculated – according to Equation (3) for positive items, and analogously for negative items.

$$F_p - F_n \quad (3)$$

Finally, the adjusted vocabulary list with the new scores is ready for the next iteration³.

4.4 Iteration Control

To maximize the number of productive iterations while avoiding unnecessary processing and arbitrary iteration limits, iterative retraining is stopped when there is no change to the classification of any document over the previous two iterations.

5 Experiments

5.1 Data

As our approach is unsupervised, we do not use an annotated training corpus, but run our iterative procedure on the raw data extracted from an annotated test corpus, and evaluate the final accuracy of the system with respect to the annotations in that corpus.

Our test corpus is derived from product reviews harvested from the website IT168⁴. All the reviews were tagged by their authors as either positive or negative overall. Most reviews consist of two or three distinct parts: positive opinions, negative opinions, and comments ('other') – although some reviews have only one part. We removed duplicate reviews automatically using approximate matching, giving a corpus of 29531 reviews of which 23122 are positive (78%) and 6409 are negative (22%). The total number of different products in the corpus is 10631, the number of product categories is 255, and most of the reviewed products are either software products or consumer electronics. Unfortunately, it appears that some users misuse the sentiment

³An alternative approach might be to use point-wise mutual information instead of relative frequencies of newly found features in a subcorpus produced in the previous iteration. However, in preliminary experiments, SO-PMI did not produce good corpora from the first iteration. Also, it is not clear how to manage subsequent iterations since PMI would have to be calculated between thousands of new vocabulary items and every newly found sequence of characters, which would be computationally intractable.

⁴<http://product.it168.com>

tagging facility on the website so quite a lot of reviews have incorrect tags. However, the parts of the reviews are much more reliably identified as being positive or negative so we used these as the items of the test corpus. In the experiments described below we use 10 subcorpora containing a total of 7982 reviews, distributed between product types as shown in Table 1.

<i>Corpus/product type</i>	<i>Reviews</i>
Monitors	683
Mobile phones	2317
Digital cameras	1705
MP3 players	779
Computer parts (CD-drives, motherboards)	308
Video cameras and lenses	361
Networking (routers, network cards)	350
Office equipment (copiers, multifunction devices, scanners)	611
Printers (laser, inkjet)	569
Computer peripherals (mice, keyboards, speakers)	457

Table 1. Product types and sizes of the test corpora.

We constructed five of the corpora by combining smaller ones of 100–250 reviews each (as indicated in parentheses in Table 1) in order to have reasonable amounts of data.

Each corpus has equal numbers of positive and negative reviews so we can derive upper bounds from the corpora (Section 5.2) by applying supervised classifiers. We balance the corpora since (at least on this data) these classifiers perform less well with skewed class distributions⁵.

5.2 Baseline and Upper Bound

Since the corpora are balanced with respect to sentiment orientation the naïve (unsupervised) baseline is 50%. We also produced an upper bound using Naive Bayes multinomial (NBm) and Support Vector Machine (SVM)⁶ classifiers with the NTU Sentiment Dictionary (Ku *et al.*, 2006) vocabulary items as the feature set. The dictionary contains 2809 items in the 'positive' part and 8273 items in the 'negative'. We ran

⁵We have made this corpus publicly available at <http://www.informatics.sussex.ac.uk/users/tz21/coling08.zip>

⁶We used WEKA 3.4.11 (<http://www.cs.waikato.ac.nz/~ml/weka>)

both classifiers in 10-fold stratified cross-validation mode, resulting in the accuracies shown in Table 2. The macroaveraged accuracies across all 10 corpora are 82.78% (NBm) and 80.89% (SVM).

<i>Corpus</i>	<i>Nbm (%)</i>	<i>SVM (%)</i>
Monitors	86.21	83.87
Mobile phones	86.52	84.49
Digital cameras	82.27	82.04
MP3 players	82.64	79.43
Computer parts	81.10	79.47
Video cameras and lenses	83.05	84.16
Networking	77.65	75.35
Office equipment	82.13	80.00
Printers	81.33	79.57
Computer peripherals	84.86	80.48

Table 2. Upper bound accuracies.

We also tried adding the negations and adverbials specified in Section 3 to the feature set, and this resulted in slightly improved accuracies, of 83.90% (NBm) and 82.49% (SVM).

An alternative approach would have been to automatically segment the reviews and then derive a feature set of a manageable size by setting a threshold on word frequencies; however the extra processing means that this is a less valid upper bound.

Another possible comparison could be with a version of Turney's (2002) sentiment classification method applied to Chinese. However, the results would not be comparable since Turney's method would require the additional use of very large text corpus and the manual selection of positive and negative seed words.

5.3 Experiment 1

To be able to compare to the accuracy of the almost-unsupervised approach of Zagibalov & Carroll (2008), we ran our system using the seed 好 (*good*) for each corpus. The results are shown in Table 3. We compute precision, recall and F_1 measure rather than just accuracy, since our classifier can omit some reviews whereas the supervised classifiers attempt to classify all reviews. The macroaveraged F_1 measure is 80.55, which beats the naïve baseline by over 30 percentage points, and approaches the two upper bounds.

Corpus	Iter	P	R	F ₁
Monitors	12	86.62	86.24	86.43
Mobile phones	11	90.15	89.68	89.91
Digital cameras	13	81.33	80.23	80.78
MP3 players	13	86.10	85.10	85.60
Computer parts	10	69.10	67.53	68.31
Video cameras and lenses	10	82.81	81.44	82.12
Networking	11	69.28	68.29	68.78
Office equipment	12	81.83	80.36	81.09
Printers	12	81.04	79.61	80.32
Computer peripherals	10	82.20	81.84	82.02
Macroaverage		81.05	80.03	80.54

Table 3. Results with the single, manually chosen seed 好 (*good*) for each corpus.

5.4 Experiment 2

We then ran our full system, including the seed identifier. Appendix A shows that for most of the corpora the algorithm found different (highly domain-salient) seeds. Table 4 shows the results achieved.

Corpus	Iter	P	R	F ₁
Monitors	11	85.57	85.07	85.32
Mobile phones	10	92.63	92.19	92.41
Digital cameras	13	84.92	83.58	84.24
MP3 players	13	88.69	87.55	88.11
Computer parts	12	77.78	77.27	77.52
Video cameras and lenses	11	83.62	81.99	82.8
Networking	13	72.83	72.00	72.41
Office equipment	10	82.42	81.34	81.88
Printers	12	81.04	79.61	80.32
Computer peripherals	10	82.24	82.06	82.15
Macroaverage		83.17	82.27	82.72

Table 4. Results with the seeds automatically identified for each corpus.

Across all 10 subcorpora, the improvement using automatically identified seed words compared with just using the seed *good* is significant

(paired t-test, $P < 0.0001$), and the F_1 measure lies between the two upper bounds.

6 Conclusions and Future Work

The unsupervised approach to seed words selection for sentiment classification presented in this paper produces results which in most cases are close to the results of supervised classifiers and to the previous almost-unsupervised approach: eight out of ten results showed improvement over the human selected seed word and three results outperformed the supervised approach, while three other results were less than 1% inferior to the supervised ones.

How does it happen that the chosen seed is usually (in our dataset – always) positive? We think that this happens due to the socially accepted norm of behaviour: as a rule one needs to be friendly to communicate with others. This in turn defines linguistic means of expressing ideas – they will be at least slightly positive overall. The higher prevalence of positive reviews has been observed previously: for example, in our corpus before we balanced it almost 80% of reviews were positive; Pang *et al.* (2002) constructed their movie review corpus from an original dataset of 1301 positive and 752 negative reviews (63% positive). Ghose *et al.* (2007) quote typical examples of highly positive language used in the online marketplace. We can make a preliminary conclusion that a relatively high frequency of positive words is determined by the usage of language that reflects the social behaviour of people.

In future work we intend to explore these issues of positivity of language use. We will also apply our approach to other genres containing some quantity of evaluative language (for example newspaper articles), and see if it works equally well for languages other than Chinese. It is also likely we can use a smaller set of negation words and adverbials to produce the seed lists.

Acknowledgements

The first author is supported by the Ford Foundation International Fellowships Program.

References

- Aue, Anthony, and Michael Gamon. 2005. Customizing Sentiment Classifiers to New Domains: a Case Study. In *Proceedings of the International Conference RANLP-2005 Recent Advances in Natural Language Processing*.

- Blitzer, John, Mark Dredze, and Fernando Pereira. 2007. Biographies, Bollywood, Boom-boxes and Blenders: Domain Adaptation for Sentiment Classification. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. 440–447.
- Dave, Kushal, Steve Lawrence, and David M. Pennock. 2003. Mining the Peanut Gallery: Opinion Extraction and Semantic Classification of Product Reviews. In *Proceedings of the Twelfth International World Wide Web Conference*. 519–528.
- Engström, Charlotte. 2004. *Topic Dependence in Sentiment Classification*. Unpublished MPhil Dissertation. University of Cambridge.
- Foo, Schubert, and Hui Li. 2004. Chinese Word Segmentation and Its Effects on Information Retrieval. *Information Processing and Management*, 40(1). 161–190.
- Ghose, Anindya, Panagiotis Ipeirotis, and Arun Sundararajan. 2007. Opinion Mining using Econometrics: A Case Study on Reputation Systems. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. 416–423.
- Hu, Minqing, and Bing Liu. 2004. Mining and Summarizing Customer Reviews. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 168–177.
- Ku, Lun-Wei, Yu-Ting Liang, and Hsin-Hsi Chen. 2006. Opinion Extraction, Summarization and Tracking in News and Blog Corpora. In *Proceedings of the AAAI-2006 Spring Symposium on Computational Approaches to Analyzing Weblogs*. AAAI Technical Report.
- McDonald, Ryan, Kerry Hannan, Tyler Neylon, Mike Wells, and Jeff Reynar. 2007. Structured Models for Fine-to-Coarse Sentiment Analysis. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. 432–439.
- Pang, Bo, and Lillian Lee. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*. 271–278.
- Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*. 79–86.
- Read, Jonathon. 2005. Using Emoticons to Reduce Dependency in Machine Learning Techniques for Sentiment Classification. In *Proceedings of the ACL Student Research Workshop at ACL-05*. 43–48.
- Tan, Aoshuang. 2002. Problemy skrytoj grammatiki. Sintaksis, semantika i pragmatika jazyka izolirujuščego stroja na primere kitajskogo jazyka [Problems of a hidden grammar. Syntax, semantics and pragmatics of a language of the isolating type, taking the Chinese language as an example]. *Jazyki Slavjanskoj Kultury*.
- Turney, Peter D. 2002. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. 417–424.
- Xu, Jia, Richard Zens, and Hermann Ney. 2004. Do We Need Chinese Word Segmentation for Statistical Machine Translation? In *Proceedings of the Third SIGHAN Workshop on Chinese Language Learning*. 122–128.
- Yarowsky, David. 1995. Unsupervised Word Sense Disambiguation Rivaling Supervised Methods. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*. 189–196.
- Yu, Hong, and Vasileios Hatzivassiloglou. 2003. Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinion Sentences. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*. 129–136.
- Zagibalov, Taras, and John Carroll. 2008. Unsupervised Classification of Sentiment and Objectivity in Chinese Text. In *Proceedings of the Third International Joint Conference on Natural Language Processing*. 304–311.

Appendix A. Seeds Automatically Identified for each Corpus.

<i>Corpus</i>	<i>Seed</i>	<i>Corpus</i>	<i>Seed</i>
Monitors	好 (good) 便 (convenient; cheap) 清晰 (clear) 直 (straight) 方便 (comfortable) 满 (fill, fulfill) 锐利 (sharp) 舒服 (comfortable) 爽 (cool)	Video cameras and lenses	清晰 (clear – of sound or image) 方便 (comfortable) 实用 (practical) 理想 (perfect) 爽 (cool)
Mobile phones	好 (good) 支持 (support) 便 (convenient; cheap) 方便 (comfortable) 清晰 (clear –of sound or image) 足 (sufficient) 好用 (easy to use) 舒服 (comfortable) 人性化 (user friendly) 流畅 (smooth and easy) 清楚 (distinct) 爽 (cool) 好了 (has become better) 耐用 (durable) 方便的 (comfortable) 满意的 (satisfied) 适应 (fit, suit) 方便了 (has become comfortable) 适用 (applicable) 顺手 (handy) 科学 (science, scientific)	Digital cameras	好 (good) 便 (convenient; cheap) 方便 (comfortable) 清晰 (clear–of sound or image) 专业 (special) 爽 (cool) 满意 (satisfied) 耐用 (durable) 舒服 (comfortable) 理想 (perfect) 真实 (straight) 稳定 (stable) 方便了 (has become comfortable) 客气 (polite) 详细 (detailed)
Networking	稳定 (stable)	Printers	好 (good)
MP3 players	好 (good) 便 (convenient; cheap) 方便 (comfortable) 实用 (practical) 灵敏 (sensitive) 舒服 (comfortable) 爽 (cool) 方便了 (has become comfortable)	Computer peripherals	好 (good) 便 (convenient; cheap) 方便 (comfortable) 准 (precise) 舒服 (comfortable) 习惯 (habitual) 流畅 (smooth and easy) 稳定 (stable)
Computer parts	好 (good) 稳定 (stable)	Office equipment	好 (good) 方便 (comfortable) 稳定 (stable) 实用 (practical)