The Semantic Content of the General Classifier in Mandarin

James Myers
Shu-Ping Gong
Zhong-Gang Shen
National Chung Cheng University

Address of first author:

James Myers
Graduate Institute of Linguistics
National Chung Cheng University
Min-Hsiung, Chia-Yi 621
Taiwan, ROC

Tel: 886-5-242-8251

Email address of first author:
Lngmyers@ccu.edu.tw

Presented at the International Association of Chinese Linguistics Eighth Annual Conference
Melbourne, Australia
July 5-7, 1999

Abstract

A basic question in the study of nominal classifier systems concerns the relative roles of lexical semantics and grammatical rules. In this paper we address this issue by studying whether the Mandarin general classifier *ge* can be selected by a lexicon-independent default rule. Evidence comes first from a corpus analysis, which found no semantically coherent set of privileged lexical exemplars for *ge*. This finding was confirmed in an experiment in which native speakers of Taiwan Mandarin were given classifiers and asked to list as many appropriate nouns as they could; the type/token ratio for *ge* was significantly lower than that for most other classifiers tested, again suggesting that *ge* has no privileged exemplars. A second experiment was conducted to determine if the choice of *ge* can be forced by the presence of specific lexical semantic features. The experiment showed that it cannot; no significant difference in rate of *ge* use was found between nonce words given specific semantic features and nonce words given no meaning at all. Together the evidence implies that *ge* does not have semantic content, and therefore is selected by grammatical rule independent of the lexicon.

1. Introduction

One of the most fundamental debates in the study of language concerns the relative importance of the lexicon versus grammar. The debate is of deep importance to all aspects of language structure and use, including phonology (are word pronunciations completely specified in the lexicon, or do speakers apply rules and/or constraints to generate surface forms?), morphology (is morphological structure represented in the memorized lexical items, or by morphological rules in the grammar?) and syntax (is syntactic structure built primarily with the help of syntactic information stored with individual words and morphemes, or is syntax fully independent of the lexicon?) More generally, is the division of the lexicon and grammar into separate modules the defining characteristic of the human linguistic capacity, as claimed by researchers such as Pinker (1994), or is it merely an outdated assumption with no empirical foundation, as claimed by researchers such as Bates and Goodman (1997)?

A form of this debate has also influenced the study of nominal semantic classifier systems, such as those found in many East Asian languages (including the Sinitic languages, Thai, Korean, and Japanese) and in fact all across the world (Allan (1977), Aikhenvald (1999)). Here the central question takes the following form: Do speakers choose classifiers solely on the basis of idiosyncratic lexical semantic properties of individual nouns, or is lexicon-independent grammatical processing involved as well?

A critical piece of evidence in this debate concerns the existence (or lack thereof) of a so-called general or default classifier, a semantically vacuous morpheme that acts as an "elsewhere" or "miscellaneous" classifier for nouns, used merely to fulfill syntactic requirements that a classifier be used in all constructions of some type (e.g. in Mandarin, in NPs with a numeral). If a nominal classifier system shows evidence of such a default, then it may appear that a grammatical approach to classifier systems is necessary, since by definition defaults are chosen in the absence of more specific information, in this case lexical semantics. In Mandarin, for example, the classifier *ge* has been described in the standard literature as just such a default; Li and Thompson (1981, 112), for instance, claim that *ge* is Mandarin's "general classifier".

In recent years, however, such a straightforward conclusion has faced some new challenges. First, based on a cross-linguistic survey, Zubin and Shimojo (1993) have argued that the very concept of a general "default" classifier is ill-defined, as it confuses three distinct semantic functions realized differently in different languages. They argue that a general classifier may serve the Complement Function, whereby it marks all nouns that don't have a semantically more specific classifier, and/or the Default Function, whereby it may replace other classifiers under certain circumstances, and/or the Unspecified Referent Function, whereby it may be used in semantically vague contexts. If correct, this typology poses a challenge for the grammatical view of default classifiers. Not only is it in principle possible for a single language to have three distinct "general" classifiers, each serving a distinct function, but even for languages with a single default, it appears to be unpredictable which of these functions will be served. Thus with three distinct functions available, the behavior of the general classifier in a given language cannot fall out automatically from its role as an empty grammatical morpheme.

Second, some researchers (e.g. Ahrens (1994), Tyan (1996)) have claimed that a classifier system may have secondary "general" classifiers restricted to specific semantic domains; for example, Mandarin is claimed to have a default ANIMAL classifier, in addition to more a specific classifier just for horses.[1] This would seem to imply that the so-called general classifier *ge* in Mandarin does not have any special grammatical status, since the mechanism that accounts for the relationship between default and specific classifiers in some domain (e.g. animals) could also be used to account for the relationship between default *ge* and all other specific classifiers. That is,

ge would merely be quantitatively different from the default ANIMAL classifier, not qualitatively different.

Finally, in Mandarin in particular, researchers have claimed that the so-called general classifier *ge* actually has specific "core" meanings, and thus cannot be a semantically vacuous default that is selected for purely grammatical reasons (e.g. Zubin and Shimojo (1993), Loke (1994), Tyan (1996)).

Nevertheless, we argue that the Mandarin classifier *ge* is indeed a default, selected merely to fulfill grammatical principles rather to obey semantic restrictions stipulated in the lexicon. First, we show that *ge* actually serves all three of Zubin and Shimojo's functions, which is as should be expected if it is truly a default in processing. Second, we show how apparent semantically restricted secondary "general" classifiers in the Mandarin classifier system differ in nature from the true unique default classifier. Third, and most important, we show that *ge* not only does not have the core meanings that have been claimed for it, but that it appears to have no meaning whatsoever. This third point is supported by an analysis of a large corpus of written Mandarin and from two psycholinguistic experiments. Together these results support not only the standard analysis of *ge* but also, more generally, a "grammatical" approach to classifier systems.

This paper is organized around two central predictions of the grammatical approach to the Mandarin classifier system: first, that there can be at most one unique default, and second, that this default classifier does not itself have a core meaning. We (re)introduce the Mandarin classifier system in section 2, and address the first prediction in section 3. In the remaining sections we address the second prediction. Section 4 gives a summary of some of our earlier corpus work (Myers, Chiang, and Gong (1999)). Section 5 describes an experiment in which subjects were asked to list nouns associated with a given classifier, a task that allowed us to assess the degree to which *ge* is associated with specific lexical items. Section 6 describes an experiment in which subjects had to choose classifiers for nonce words that were given various experimentally manipulated "meanings," the goal being to find out if the choice of *ge* could be affected by lexical semantic information. Both experiments confirmed our expectations: *ge* is selected independently of lexical semantics. Finally, in our conclusions in section 7, we discuss implications of the grammatical approach for the study of defaults in other linguistic systems.

2. The Mandarin classifier system

Mandarin noun classifiers (sometimes called measure words, though see Tai (1994) for discussion of this concept) are required in NPs that contain numbers, determiners, or certain quantifiers. There are several different kinds of Mandarin classifiers (see Tai (1994), Kuo (1998), Ahrens and Huang (1996)). These include standard measures (e.g. *yi-bang rou* 'a pound of meat'), container measures (e.g. *yi-bei cha* 'a cup of tea'), partitive measures (e.g. *yi-kuai dangao* 'a piece of cake', *yi-pian tusi* 'a slice of toast'), group measures (e.g. *yi-qun gou* 'a pack of dogs', *yi-shuang kuaizi* 'a pair of chopsticks'), kind classifiers (e.g. *na-zhong ma* 'that kind of horse') and event classifiers (*zhe-chang dianying* 'this (showing of a) movie').

In our studies, we have focused on individual classifiers, morphemes that are selected by individual entities on the basis of their inherent semantics. In the following figure we give examples showing the use of some of the most common individual classifiers.

(1)	ge	*yi-ge ren* 'a person', *yi-ge guojia* 'a country', *yi-ge xigua* 'a watermelon', *yi-ge taiyang* 'a sun'
	wei	*yi-wei laoshi* 'a teacher'
	zhang	*yi-zhang zhi* 'a (piece of) paper', *yi-zhang zhuozi* 'a table'

tiao	*yi-tiao lu* 'a road', *yi-tiao yu* 'a fish'
jian	*yi-jian shiqing* 'a thing', *yi-jian yifu* 'an article of clothing'
pian	*yi-pian yezi* 'a leaf'
zhi[2]	*yi-zhi gou* 'a dog', *yi-zhi xiezi* 'a shoe'
zhi[2]	*yi-zhi yuanzibi* 'a ballpoint pen'
ke	*yi-ke yachi* 'a tooth'
li	*yi-li mi* 'a rice grain'
mian	*yi-mian qiang* 'a wall'
gen	*yi-gen gunzi* 'a stick'
ba	*yi-ba daozi* 'a knife', *yi-ba yizi* 'a chair'

By examining such lists, several important observations can be made. First, as demonstrated by recent work such as Tai and Wang (1990), Tai (1992), Tai and Chao (1994), Tai (1994), Shi (1996), Huang, Chen, and Lai (1996), and many others, individual classifiers are indeed selected on the basis of inherent semantics. For example, all animals can be classified with the morpheme *zhi*, and people that one should be polite to should take *wei*. Second, it is nevertheless not always obvious precisely what semantic properties are being marked by a given classifier. For example, nouns for oblong articles of clothing can take *tiao*, such as *kuzi* 'pants', yet *duanku* 'shorts', which are by definition not oblong, also take *tiao*, perhaps because it shares a head morpheme with *kuzi* (Wiebusch (1995)). Third, some classifiers seem to mark a disparate set of items, and it is sometimes difficult to decide the reasons for this. Such cases often seem to involve polysemy (i.e. they indicate distinct but related meanings), such as *ba*, which classifies both knives and chairs since both are traditionally thought of as capable of being picked up with one hand, or *zhang*, which indicates objects with an extended surface (including tables and chairs, which are not flat overall, but whose functional portion is) and mouths (which can be opened wide). Other cases seem to involve such extremely extended meanings that they may essentially be treated as homonymy, such as *zhi*, which classifies both animals and a member of a pair, and if its homophones are included, also certain kinds oblong objects and rings.

The most extreme example of an individual classifier that can be used for more than one semantic class is *ge*. The above figure shows its use with people (e.g. *ren* 'person'), abstractions (e.g. *guojia* 'country'), and non-flat, non-oblong objects of various sizes (e.g. *xigua* 'watermelon', *taiyang* 'sun'). In addition, it can be used as an optional replacement for other classifiers. As Loke (1994) notes, certain kinds of classifiers appear more resistant to such "neutralization" to *ge* than do others, but the study of classifier use in natural speech (e.g. Erbaugh (1986)) shows that in principle *ge* can replace any classifier at all. For such reasons, Li and Thompson (1981, 112) write that *ge* "is gradually becoming the general classifier and replacing the more specialized ones." In fact, *ge* has probably been a general classifier since the Tang dynasty (618-907 CE), and even before then it was already being used a variety of disparate semantic contexts (Wang (1989)).

3. Classifier systems and defaults

In this section we deal with the concept of a unique default classifier. This concept makes two predictions: first, that a grammatical default classifier will necessarily serve all three of Zubin and Shimojo's (1993) functions, and second, that a classifier system will necessarily have one and only one default. We address each prediction in turn.

3.1 The functions of defaults

In the grammatical view of classifier systems, general classifiers are selected by a single grammatical rule (e.g. in Mandarin, "NPs with numbers, determiners or certain quantifiers must have a classifier"; see Cheng and Sybesma (1998) and Li (1999) for more sophisticated analyses of the syntax of classifiers in Mandarin). As is claimed to be the case with default regular inflection (e.g. Pinker (1991) and subsequent work), there must be only one default classifier rule, which applies if and only if the proper output is not generated through lexical considerations. In the case of Mandarin, this means that if *ge* is a default classifier, then it is the only one.

In other words, our definition of a "default" is in processing terms: a default is chosen by speakers when, for whatever reason, a specific classifier is not accessed from the lexicon. In Myers et al. (1999), we demonstrated that *ge* is used in precisely this way, appearing in a wide variety of contexts that have nothing in common except the inability to form analogies with exemplars in the mental lexicon, e.g. when nouns are too dissimilar from lexical exemplars, when nouns are derived from other syntactic categories, when nouns cooccur with classifiers too infrequently, and when speakers have memory access problems due to brain damage or inexperience with the language (as an L1 or L2 speaker).

In this context, the three functions ascribed to general classifiers by Zubin and Shimojo (1993) can be seen as side-effects of this fundamental processing fact. A default is used in the Complement Function if an appropriate specific classifier is unavailable because it simply doesn't exist, in the Default Function if the speaker didn't happen to think of an existent specific classifier at the moment (due to various transient memory-access problems), and in the Unspecified Referent Function if there is insufficient information to decide which specific classifier would be most appropriate.

Zubin and Shimojo (1993) list Mandarin *ge* as an example of a general classifier serving the Default Function due to its well-known ability to replace other classifiers, but it actually serves all three functions. The fact that it also serves the Complement Function is demonstrated by the disjointedness of the semantic categories that it marks (e.g. HUMAN, ABSTRACTION, 3D-OBJECT). The reason these categories are disjoint is that they represent the negative space left by removing the more coherent categories marked by the specific classifiers. Zubin and Shimojo (1993, 493) admit as much when they allude to a study now published as Ahrens (1994), which found that speakers are more likely to replace a specific classifier with *ge* for less prototypical members of a category, thus indicating that "*ge* does show some tendency toward *complement* function" (italics in original).

It is important not to mistake this Complement Function of *ge* for evidence that *ge* itself has meaning. Loke (1994) observes that *ge* is more likely to replace function-based classifiers (e.g. the "vehicle" classifier *liang*) than shape-based classifiers (e.g. *zhang* or *tiao*), but his conclusion, that *ge* therefore itself implies something about shape (e.g. 3D-OBJECT), does not follow. This fact actually concerns, not *ge*, but rather the differences between function-based and shape-based specific classifiers. Function-based classifiers are less frequently used than shape-based classifiers, and research on language development has found that there is a strong preference to use shape rather than other characteristics to classify objects (Pinker (1989)). The same point can be made for the animal classifier *zhi*, which speakers almost never replace with *ge* (though Erbaugh (1986) found that some speakers do); Pinker (1989) also notes that animacy is a highly salient property in word learning. Thus what Loke (1994) has observed is the fact that *ge* is more often used for nouns that have weak classifiers, that is, classifiers that are more difficult to access and process. Hence this behavior is actually an instance of what Zubin and Shimojo (1993) would label the Complement Function (and what we would simply call default rule processing).

Finally, *ge* serves not just these two functions, but the Unspecified Referent Function as well. This can be seen from common phrases such as *zhege* 'this (one)' and *nage* 'that (one)', which like their English glosses, are all-purpose deictic NPs (see Myers et al. (1999) for further discussion).

The fact that *ge* serves all three functions suggests that it truly is processed as a default. The issue of differing behavior by general classifiers in other languages will be discussed in the concluding section.

3.2 Semantically specific "defaults"

The second prediction of the grammatical view is that a classifier system will have at most one default. This seems to be challenged by the common claim that such-and-such a specific classifier is a "default" within some specific semantic domain. For example, in a psycholinguistic experiment, Tyan (1996) found that speakers' judgments of the acceptability of *ge* for a given noun were inversely correlated with the acceptability of a specific classifier for that noun (e.g. *ge* N was judged better when *zhang* N was judged worse). Yet within the category of small objects, *li* and *ke* stood in a similar inverse relation with other classifiers, implying that they too are a sort of default, though just within this semantic category. Similarly, within the category of animals, some researchers (e.g. Hu (1993), Ahrens (1994)) suggest that the "default" is *zhi*, not *ge*, so that when adults or children fail to use the specific HORSE classifier *pi*, they say *yi-zhi ma* instead, almost never *yi-ge ma*. Zubin and Shimojo (1993) also assume this concept of semantically restricted "defaults" in their cross-linguistic survey of classifier systems.

Although this view is common, we think it is mistaken. Such cases of "neutralization" to a classifier other than *ge* simply show that default processing is not the only way to select a classifier. As noted in the previous section, *ge* rarely replaces the animal classifier *zhi*, not because *zhi* is the processing default for the animal domain, but because animacy is a highly salient semantic property, which overrides the default *ge* rule. In our view, the replacement of the HORSE classifier *pi* by the ANIMAL classifier *zhi* is a result of a quite different process from that involved in the replacement of the VEHICLE classifier *liang* by *ge*, as illustrated schematically in (2).

(2)	a.	ma 'horse'	HORSE	→	pi
			ANIMAL	→	**zhi**
	b.	chezi 'car'	VEHICLE	→	liang

In the case of *ma* 'horse', different semantic features are lexically linked to different classifiers. Through some principle, akin perhaps both to the Elsewhere Condition (Kiparsky (1973)) and the conversational maxim of quantity (Grice (1975)), speakers are expected to select the most specific feature associated with a classifier, in this case HORSE. There is a competing feature, however, and this one, ANIMAL, is cognitively more salient (as indicated by bolding), and so its associated classifier is often used instead. By contrast, in the case of *chezi* 'car', there is only one relevant semantic feature, VEHICLE. Like the feature HORSE, however, it is not highly salient, and yet unlike the previous case, there is no highly salient feature to compete with it. Thus if a speaker fails to access *liang* in response to the feature VEHICLE, there is no choice but to fall back on the default classifier *ge* in order to fulfull syntactic obligations.

Independent evidence for this dual-mechanism view comes from the observation that in spite of variations in probability of neutralization across classifiers, in practice all classifiers can be replaced by *ge*, even highly salient ones like *zhi*. As already noted, in casual speech, some adults use *ge* in almost all environments where syntactic principles require an individual classifier

(Erbaugh (1986)).

It is worthwhile to note that such non-default "neutralization" also occurs outside of classifier systems. Children acquiring English past tense inflection also overuse patterns other than the default *-ed*. For instance, children may say *brung* rather than *brought* as the past tense of *bring* (Xu and Pinker 1995). This doesn't mean that English has an additional default past-tense rule, but only that *sting-stung*, *dig-dug*, and so on, allow for particularly robust analogies that can override the default rule.

4. Distributional evidence for a meaningless *ge*

Some scholars acknowledge that *ge* is a "general" classifier in some sense, but insist that it nevertheless has a specific "core" meaning (i.e. semantic function). Zubin and Shimojo (1993) state that this meaning is HUMAN, while Loke (1994) notes that *ge* also indicates ABSTRACTION and 3D-OBJECT. By contrast, if *ge* is a true default, then it should not be allowed to have any special meaning of its own. In this section we review one piece of evidence for the meaningless of *ge* from our earlier corpus analysis (reported in Myers et al. (1999)). This study used the Academia Sinica Balanced Corpus, which is composed primarily of articles from newspapers in Tawian (for further details about this corpus, see Chen, Huang, Chang and Hsu (1996); World Wide Web access is at http://www.sinica.edu.tw/ftms-bin/kiwi.sh).

One objective test of such claimed core meanings would be to examine the distribution of the different semantic classes that collocate with *ge* (e.g. humans, abstractions, 3D objects) to determine which has the highest proportion of privileged exemplars. For example, if the core meaning of *ge* were really HUMAN, then we would expect that of all the nouns that collocate with *ge*, the most common collocations would involve nouns naming humans. Nouns that fall into the *ge* class merely by default would tend to collocate with *ge* less frequently. The result would be that *ge*-HUMAN collocations should cluster at the top of a list of collocations ranked by collocation frequency. This should be true even if all HUMAN nouns automatically fell into the *ge* class regardless of collocation frequency, since the default non-HUMAN items would crowd the HUMAN items out at the bottom of the list.

Such an ordering by collocation frequency can be made using the measure of mutual information (MI), whose formula is given in (3) (see Church and Hanks (1990)). Essentially, MI describes how common a collocation is when the lexical frequencies of each word have been factored out. If two words *x* and *y* are distributed randomly, MI(x,y) ≤ 0; if they form meaningful collocations, MI(x,y) $>> 0$; and if they are in complementary distribution, MI(x,y) $<< 0$.

(3) Mutual information value (for two words *x* and *y*)

$$MI(x,y) = \log\left(\frac{prob(x,y)}{prob(x)\cdot prob(y)}\right)$$

We used the MI calculations automatically provided by the public WWW interface to the Sinica Corpus, given a window size of five words (i.e. all instances where a classifier appeared within five words before a given noun). All examples were screened to make sure that the classifier and noun were indeed in an agreement relation within an NP. The result was a list for each classifier we examined showing all collocating nouns with positive MI values.

To deal with the issue of *ge*'s core meaning, we compared the number of collocating nouns of

different semantic classes with an MI value above versus below the median. None of the proportions for nouns with various proposed core meanings, shown in Table 1, reached significance by chi-squared tests. Thus contrary to what is standardly thought, nouns for humans, abstractions and 3D objects do not have any special status in the *ge* category. Even the most frequent HUMAN noun, *ren* 'person', doesn't collocate with *ge* unusually often, appearing roughly halfway through the list.

TABLE 1

humans	abstractions*	deverbal nouns**	3D objects
27:37	121:104	8:15	27:29

*such as *shehui* ("society") **such as *xiwang* ("wish")

One relevant observation that can be made about the information in Table 1 is that if one were forced to posit a core meaning for *ge* based on the number of tokens, it would have to be ABSTRACTION rather than HUMAN. In our view, this provides further evidence for the meaninglessness of *ge*, since, as is well known, there is nothing so difficult to define as an abstraction.

As for the proportions in Table 1, what we are observing here is either that all nouns fall into the *ge* class via default selection, which is what we claim, or that no nouns ever fall into the *ge* class by default, but instead *ge* specifically selects all nouns for humans, abstractions and 3D objects. Both hypotheses would predict an even distribution of nouns of various types across the whole MI range. However, the second hypothesis makes the absurd claim that *ge* not only has core meanings, but that it is in fact a completely semantically consistent classifier: every noun that collocates with it has a specific semantic reason for doing so. This would make *ge* more semantically consistent than most so-called specific classifiers, which as we have already seen, often seem to collocate with nouns in a somewhat haphazard way.

5. Experiment 1

Linguistic descriptions and corpus analyses may provide strong hints that *ge* does not have semantic content and is therefore selected by a default rule, but since our claims concern processing, a more satisfying method is to conduct behavioral experiments. In this section and in section 6 we describe two such experiments.

In the first, our interest was simply to find out something about the nature of the categories named by various common classifiers. We examined this by performing a simple listing task, such as that used by Rosch (1973): we gave classifiers to subjects and asked them to name associated nouns that came to mind. By doing this, we wanted to discover if the category named by *ge* had any privileged exemplars, that is, nouns that are particularly closely associated with this classifier. If so, then this would be tantamount to saying that the core meaning of *ge* is the semantic class represented by these exemplars. If by contrast ge were found to have no privileged exemplars, this would support our claim that it is essentially meaningless. More generally, this method enabled us to calculate a type/token ratio: the number of types (particular words) divided by the total number of words listed by all the subjects, including words named by more than one person. If *ge* is truly a meaningless classifier, it should have a very high type/token ratio; people should name many different nouns rather than the same nouns over and over.

This method was inspired by a study of Japanese described by Zubin and Shimojo (1993). They found that the classifier *tsu* had a type/token ratio of 0.89, which is quite high, consistent

with the fact that it serves the Complement, Default and Unspecified Referent Functions of a general classifier.

5.1 Methods

Participants: Thirty-five native speakers of Mandarin participated in this study, all students at National Chung Cheng University in Chia-Yi County, central Taiwan. Data from 11 subjects were discarded (3 for not completing the task, 8 for mechanical failures), leaving 24 sets of data to analyze.

Materials: Fourteen common individual classifiers were selected: *ge*, *tiao*, *gen*, *zhi* (ANIMAL), *zhi* (OBLONG), *zhang*, *pian*, *mian*, *li*, *ke*, *ba*, *jian*, *wei*, *chang*. Each was written on the top of separate sheets of paper, and these fourteen sheets were arranged in different random orders for each subject.

Procedure: Subjects were instructed to write down five nouns for each classifier. They were also asked to write down the time at which they started and completed listing nouns for each classifier (a digital clock was visible at the front of the room).

5.2 Results and discussion

For each classifier, two measures were calculated: (1) the average response time (i.e. the average duration it took subjects to write down five nouns for a given classifier); (2) the type/token ratio (i.e. the total number of different words listed by all subjects divided by the total number of words, including repetitions, listed by all subjects). Response time did not yield coherent results and were not further analyzed (they are presumably affected more by lexical frequencies of the nouns than by anything interesting about the processing of classifiers).

As predicted, the type/token ratio was highest for *ge* (0.61). However, chi-squared tests found that the ratio was not significantly different from that of most of the other classifiers. These results are summarized in Table 2 below, where classifiers are ordered by type/token ratio.

TABLE 2

Classifier	Type/token ratio	Most named noun (tokens)
ge	0.61	ren 'person' (20)
pian	0.59	yezi 'leaf' (8)
ba	0.57	yizi 'chair' (6), jian 'sword' (6)
li	0.54	mi 'rice grain' (20)
mian	0.53	jingzi 'mirror' (20)
tiao	0.5	pidai 'belt' (7), shengzi 'rope' (7)
wei	0.5	laoshi 'teacher' (12)
ke	0.5	shu 'tree' (21)
gen	0.47*	shuzhi 'branch' (9)
chang	0.46*	dianying 'film' (10)
zhi (OBLONG)	0.44*	bi 'pen/brush' (15)
zhang	0.43*	zhi 'paper' (22)
zhi (ANIMAL)	0.42*	gou 'dog' (14)
jian	0.33*	yifu 'clothing' (21)

*significantly different from *ge*

In addition to the fact that *ge* did not give dramatically different results from the other classifiers, another problem for our claim that *ge* has no privileged exemplars is the fact that the type/token ratio for *ge*, 0.61, is much lower for that of *tsu*, the apparent default classifier in Japanese, which Zubin and Shimojo (1993) found to have a type/token ratio of 0.89. Both of these problems seem to be caused by the fact that many subjects listed the high-frequency word *ren* 'person' for *ge*. In fact, it is apparent that this task suffers from a confound with word frequency: a noun can be frequently named for some classifier either because the classifier has the associated core meaning, or simply because that noun happens to be of high frequency.

If the word *ren* is removed from the analysis, along with the most common response for each of the other classifiers (listed in the rightmost column of Table 2), the type/token ratio for *ge* is recalculated as 0.72. This is still higher than all other classifiers tested, and is significantly higher than all except for *pian* (0.63), *li* (0.64), *mian* (0.62), and *ke* (0.61). We already saw that *pian* is less semantically restricted than most individual classifiers since it can also be used as a partitive measure The reasons for the high type/token ratios for the other three specific classifiers are less obvious.

The fact that *ren* was the most commonly listed noun for *ge* should not be interpreted as evidence that *ge* has a core meaning of HUMAN, as claimed by Zubin and Shimojo (1993) and Loke (1994). First, the corpus study had already confirmed that *ge* does not cooccur with *ren* particularly often, since much of the time *ren* appears in the corpus without any classifier at all. Second, in the present experiment only 12% of noun types listed for *ge* referred to humans. Thus we conclude that *ren* was chosen so often merely because of this word's high lexical frequency. Our results therefore not only support the claim that *ge* is meaningless, but also help explain how other researchers could have come to believe that it has a core meaning of HUMAN.

6. Experiment 2

The goal of our second experiment was to see if the choice of *ge* can be forced for nonce words by defining them in such a way that they include the semantic feature HUMAN. In other words, we wanted to know if speakers can select *ge* on the basis of similarity with lexical exemplars. If *ge* is processed via a default grammatical rule, as we claim, this should not be possible.

Here, our method was adapted from research on regular inflection by Prasada and Pinker (1993) and Marcus, Brinkmann, Clahsen, Wiese, and Pinker (1995). These earlier experiments tested the claim that speakers process regular inflection independently of exemplars in memory. Prasada and Pinker (1993), for example, asked whether subjects could produce regular past tense forms as easily for nonce words quite different from all real regular verbs (e.g. *ploamph* → *ploamphed*) as for nonce words that are very similar to real words (e.g. *plip* → *plipped*). As expected by the grammatical view of regular inflection, no difference between these two kinds of nonce forms was found.

In our case, we gave subjects nonce forms that varied in experimentally manipulated meaning, with some quite similar to words that typically take *ge* (i.e. nouns for humans) and others not similar to any specific lexical items (i.e. nouns that were given no semantics at all). If the feature HUMAN automatically selects for *ge*, contrary to what we hypothesize, we would expect there to be a dramatic difference in responses for these two types of nonce forms.

6.1 Methods

Participants: Twenty-one native speakers of Mandarin participated in this study, all students at National Chung Cheng University. 10 participated in condition A, and 11 in condition B (these conditions will be described below). An additional 30 subjects were used to prescreen materials.

Materials: The target items were 40 two-syllable nonce (invented) forms, presented auditorily. A typical nonce form was *da1nan3*, which is composed of legal syllables but is not itself a real word. Auditory presentation was of course necessary given the fact that Chinese orthography makes it impossible to write truly semantically vacuous nonce forms. Nonce forms were prescreened by presenting 60 possible items to 10 native speakers, who judged whether they suggested any connotations; the 40 "least meaningful" nonce forms were chosen as target items. Each nonce item was described to participants as being "similar" (*xiang*) to one of three classes of real words: (a) 10 nouns that do not select for *ge* (e.g. *gou* 'dog'); (b) 10 HUMAN nouns that typically collocate with *ge* (e.g. *huairen* 'bad guy'); (c) 10 semantically coherent sets of nouns which do not take a single classifier (e.g. *yusan* 'umbrella', which takes *ba*; *yuyi* 'raincoat', which takes *jian*; *yuxie* 'galoshes', which takes *shuang*); thus nonce items in (c) were ambiguous as to which classifier was most appropriate.. A fourth class of 10 items (d) were not compared to any real nouns at all; the meanings of the nonce items in this class were thus entirely vague.

The real nouns used in the exemplar-based definitions of the nonce items were carefully selected. All of the real nouns used in the experiment were prescreened by 10 native speakers to ensure that they indeed typically took the classifiers we expected. Of the real nouns used to define the nonce items in class (a), each co-occur with a different specific classifier. Only one (*laoshi* 'teacher', which typically takes *wei*) named a human. The real nouns used in class (b) were all human, but typically take *ge* rather than *wei*, as confirmed both by our prescreening and by corpus analysis. In principle, however, all human nouns that prefer *ge* can also appear with the polite human classifier *wei*, though this choice may be pragmatically very odd. The classes of nouns used in (c) were also prescreened by a further group of 10 native speakers to ensure that they were semantically coherent (e.g. the RAINGEAR set given above). All real nouns used in the definitions for classes (a)-(c) are listed in the appendix.

Each class (a)-(d) had 10 nonce items arbitrarily assigned to it. Because they were inherently meaningless, nonce items were not randomly assigned differently across subjects. All nonce items are listed in the appendix.

Procedure: The order of trials was completely randomized, with the same random order for all subjects. Subjects were presented with items played from a tape recording, with ten-second pauses between trials. Subjects were first given five trials for practice. The main experiment consisted of 40 items. For classes (a)-(c), each item was presented in a frame like the following.

(6) danan xiang huairen, nali you xuduo danan, zheli you wu __ danan
danan is like bad guy thereare many *danan* here are five __ *danan*
'*Danan* are like bad guys; over there are a lot of *danan*; here there are five __ *danan*.'

For class (d), each item was presented in a frame like the following:

(7) nali you xuduo danan, zheli you wu __ danan
there are many *danan* here are five __ *danan*
'Over there are a lot of *danan*; here there are five __ *danan*.'

In condition A, subjects had written versions of these utterances, while in condition B they only had numbered blanks. In each case, they had to write down the most appropriate classifier for the given nonce item.

6.2 Results and discussion

The proportions (in percentages) of *ge* responses across the four classes (a)-(d) for the two conditions are given in the following table.

TABLE 3

	(a) non-*ge* items	(b) *ge*-items (human)	(c) ambiguous	(d) vague
Written cues	0	76	44	65
Auditory only	0	73	25	70

The two conditions clearly showed the same overall pattern: no *ge* responses when semantic features are present that preferentially select for other classifiers (class (a)), a majority of *ge* responses both when the HUMAN feature is present (class (b)) and when no features are present (class (d)), and a smaller but nontrivial proportion of *ge* responses when the defining exemplars compete in their classifier preference (class (c)). The two conditions did show significantly different proportions, but only in class (c) ($\chi^2(1) = 8.001$, $p < 0.05$): subjects used *ge* more often when they could actually read the competing exemplars. Why this should be is not clear, but it does show that the two conditions did involve somewhat different kinds of processing.

In both conditions, there was a dramatic difference in choice of *ge* between classes (a) and (c), consistent with the interpretation that speakers tend to fall back on the default *ge* when faced with competition among classifiers. However, this result does not provide unambiguous evidence for the use of a default rule, since for some items in class (c), at least one of the defining exemplar nouns took *ge*. Thus the difference between (a) and (c) could still be due to comparison with lexical exemplars.

Moreover, in both conditions, proportions for class (c) were significantly different from those for class (d) (condition A: $\chi^2 (1) = 8.900$, $p < 0.05$; condition B: $\chi^2 (1) = 43.736$, $p < 0.05$). This suggests that the mere fact that class (c) items presented subjects with competing classifiers was not sufficient to trigger consistent use of the default *ge* rule, since if they did, their behavior should have been identical to their behavior for class (d) items. Instead, it appears that subjects often selected one of the competing classifiers for the nonce items in class (c), e.g. choosing *ba* for something similar to *yusan*, *yuyi* and *yuxie*. This is of course no problem for the grammatical view, which accepts the reality of exemplar-based processing for non-default classifiers.

The most important result for our purposes is that in both conditions there was no significant difference in proportion of *ge* responses between class (b) and class (d) (condition A: $\chi^2 (1) = 2.900$, $p > 0.05$; condition B: $\chi^2 (1) = 0.198$, $p > 0.05$). The rate of *ge* for class (b) was not 100% since subjects occasionally used the polite human classifier *wei* instead, even though it was pragmatically odd given the exemplars (e.g. *huairen* 'bad guy'). The rate of *ge* for class (d) was not 100% since subjects occasionally used some other classifier. The choice of classifiers other than *ge* in class (d) seems to be rather haphazard; no obvious pattern carried over across conditions A and B. Collapsing across conditions, the four most common choices were *kuai* 'piece' (8 tokens), homophones of *zhi* (8 tokens), *tiao* (7 tokens), and *li* (5 tokens). The choice of the partitive classifier *kuai* seems to have been motivated by the desire to choose a classifier with no

semantic restrictions, and the others were apparently chosen for no better reason than that they are highly frequent.

From the perspective of the grammatical view of classifier systems, however, what is crucial here is that it did not matter to subjects if nonce items were presented as humans or as meaningless: *ge* was still the classifier of choice, and it was chosen at the same rate. Thus subjects were not forced into giving more *ge* responses when the feature HUMAN was present. Instead, the basic *ge* response rate was maintained for human nouns as well, suggesting that subjects were selecting *ge* through the same mechanism. Since for class (d) items this mechanism must be some sort of default response, we conclude that default processing was also being used for the human nouns of class (b).

This basic *ge* selection rate, approximately 70%, deserves some attention. It appears to be stable across conditions A and B, even though, as we saw, these conditions involved separate groups of subjects and must have used somewhat different kinds of processing. We don't believe that grammatical rules necessarily come with inherent rates of application (as has been suggested, for example, in the literature on language variation; see e.g. Labov (1994)). In the case of *ge*, the basic rate is presumably a result of the fact that the production of classifiers involves both a grammatical rule and lexical access, as well as general pragmatic factors, factors specific to the demands of our task, and individual variation.

Without a complete theory of how such factors interact, we are obviously unable to explain where the 70% rate comes from, but presumably it is related in some way to the rate of *ge* use that our subjects' have been exposed to during their lifetimes. Our corpus study, which is based primarily on written language, found that *ge* use forms no more than about 35% of all individual classifier use. For fluent spoken language, Erbaugh (1986; personal communication) found in an adult story-telling study that *ge* was used 689 times, compared to only 40 times for specific classifiers, a proportion of almost 95%. The present experiment's rate of 70% lies almost perfectly at the midpoint between these two extremes. Intriguingly (but also confusingly), in experiments with 3- to 6-year-old children, Hu (1993) also elicited *ge* responses in proportions of 64-82% (where a proportion of only 8% was prescriptively "grammatical" given her materials).

The crucial point is not that a basic rate for *ge* use may or may not exist, but that the *ge* rates for meaningless and HUMAN nouns in the present experiment were the same, a result that also was stable across conditions. This provides further support for our claim that *ge* is a grammatical default uninfluenced by lexical semantics.

7. Conclusions

In this paper we have used psycholinguistic experimentation to supplement the arguments of Myers et al. (1999) that *ge* is a default selected by a grammatical rule. Although this paper focused on Mandarin, our view clearly has implications for cross-linguistic research. By demonstrating that Mandarin *ge* is truly processed as a default, we merely mean to show that such defaults are not impossible, and that therefore grammatical factors cannot be ignored in the analysis of classifier systems. However, we do not predict that all classifier systems must have such a processing default. In any case where it is in principle possible for speakers to unambiguously choose a classifier purely on the basis of lexical information, then the innate ability of the human mind to generate defaults will never be exercised.

Certainly most of the examples of classifier systems given by Zubin and Shimojo (1993) as "clear" examples of this or that function of a general classifier actually end up serving more than one function. For instance, Japanese has been claimed to have two "general" classifiers, *ko* (a cognate of *ge*) and the native *tsu*, but as Zubin and Shimojo themselves demonstrate, the true

default here is *tsu* (*ko* is semantically restricted to objects that can be held in the hand; see above for discussion of such semantically restricted "defaults"). Moreover, the classifier *tsu* is listed in their paper as an example of the unspecified referent function, and yet they show that it also serves the other two functions as well. Finally, they observe (p. 491) that in many languages, deciding on the "core semantics" for the general classifiers is "problematic," a confession that implies that most classifier languages are like Mandarin in having a meaningless grammatical default.

However, not all classifier systems show clear evidence for a single default. The German gender system is an extremely impoverished classifier system with only three categories (masculine, feminine and neuter) and with famously arbitrary categorization (e.g. the German sun is feminine and the moon is masculine, reverse of the pattern in French). In this system, neuter gender appears to serve the Unspecified Referent Function, since in appropriate pragmatic contexts the neuter *das* 'that (one)' can be used for anything, regardless of gender, and even more interestingly, neuter gender tends to be used for category names (e.g. the word for 'fruit' is neuter, whereas 'apple' is masculine and 'grape' is feminine; Zubin and Köpcke (1986)). However, it does not unambiguously serve the other two functions. For example, a foreign borrowing (i.e. inherently "genderless") will tend to get its gender through a variety of lexically determined ways rather than simply taking the default. Thus *Tofu* 'tofu' is masculine, not neuter as one might expect via the Complement Function, presumably because it is similar to *Quark* 'cream cheese'. Neuter gender also fails to demonstrate the Default Function, primarily because unlike Mandarin classifiers, German gender is not very variable (though more so than one would guess from grammar books; see e.g. Zubin and Köpcke (1984)). In our view, the fact that neuter gender doesn't clearly serve all three functions in German means that it isn't really processed as a rule-governed default. German speakers simply remember which gender goes with which noun; semantic patterns such as those noted above arise solely through analogical processing, e.g. in a connectinoist network of the sort described in MacWhinney, Leinbach, Taraban, and McDonald (1989).

Nevertheless, there may be occasions where the innate tendency towards grammatical rules in language seems to force the appearance of defaults, even when they aren't logically necessary. One example is plural allomorphy in Arabic and German. In German, for instance, plurals can be formed in a bewildering variety of ways, but as Marcus, Brinkmann, Clahsen, Wiese, and Pinker (1995) have experimentally demonstrated, only one of them (not even a particularly common one) serves as a default. Of course, due to the nature of inflection, this conclusion rests primarily on the regular plural serving the Complement Function (or occasionally the Default Function); the Unspecified Referent Function doesn't make sense in the case of inflection. In any case, it may be that the processing of inflection inherently requires a default, perhaps because inflection is tied into syntax, the prototypical non-lexical module of language (though see Bates and Goodman (1997)). By contrast, some classifier systems may be tied so tightly to lexical semantics that some of them are processed entirely without lexicon-independent rules. Still, as we hope to have shown by our study of Mandarin, this is not a logical necessity, and in fact it may not be the usual case.

NOTES

[*] The research in this paper was supported by National Science Council (Taiwan ROC) grant NSC 87-2418-H-194-003 and a National Chung Cheng University seed grant to the first author.

[1]Names of semantic features in this paper are written in all capitals, in accordance with standard conventions.

[2]These two morphemes are traditionally written with different characters, though even in Taiwan, which maintains the traditional orthography, people are occasionally confused about which character should be used for which morpheme.

REFERENCES

Ahrens, Kathleen (1994) "Classifier production in normals and aphasics," *Journal of Chinese Linguistics* 22, 203-247.

Ahrens, Kathleen and Chu-Ren Huang (1996) "Classifiers and semantic type coercion: motivating a new classification of classifiers," *Language, Information and Computation* 11, 1-10.

Aikhenvald, A. Y. (1999) *Classifiers: A Typology of Noun Categorization Devices.*

Allan, Keith (1977) "Classifiers," *Language* 53, 285-311.

Bates, Elizabeth, and Judith C. Goodman (1997) "On the inseparability of grammar and the lexicon: evidence from acquisition, aphasia, and real-time processing," *Language and Cognitive Processes* 12, 507-584.

Chen, Keh-Jiann, Chu-Ren Huang, Li-Ping Chang, and Hui-Li Hsu (1996) "SINICA CORPUS: design methodology for balanced corpora," *Language, Information and Computation* 11, 167-176.

Cheng, Lisa L.-S. and Rint Sybesma (1998) "*Yi-wan Tang, Yi-ge Tang*: Classifiers and Massifiers," *Tsing Hua Journal of Chinese Studies*, New Series 28(3), 385-412.

Church, Kenneth W. and Patrick Hanks (1990) "Word Association Norms, Mutual Information, and Lexicography," *Computational Linguistics* 16, 22-29.

Erbaugh, M. (1986) "Taking Stock: the Development of Chinese Noun Classifiers Historically and in Young Children," in C. Craig (ed.), *Noun Classes and Categorization*, John Benjamins Publishing Company, Amsterdam, pp. 399-436.

Grice, H. P. (1975) "Logic and Conversation," in P. Cole and J. L. Morgan (eds.), *Syntax and Semantics, Vol. 3. Speech Acts*, Seminar Press, pp. 41-58.

Hu, Qian (1993) *The Acquisition of Chinese Classifiers by Young Mandarin-Speaking Children*, PhD dissertation, Boston University.

Huang C-R, K-J Chen, and T-H Lai (1996) *Changyong Liangci Cidian* [Dictionary of Common Classifiers], Guomin Ribao Chubanzhongxin, Taipei.

Kiparsky, Paul (1973) "'Elsewhere' in Phonology," in S. Anderson and P. Kiparsky (eds.), *A Festschrift for Morris Halle*, Hold, Rinehart and Winston, pp. 93-106.

Kuo, Yi-jun (1998) *A Semantic and Constrastive Analysis of Mandarin and English Measure Words*, MA thesis, National Chung Cheng University.

Labov, William (1994) *Principles of Linguistic Change: Internal Factors*. Blackwell.

Li, Charles N. and Sandra A. Thompson (1981) *Mandarin Chinese: A Functional Reference Grammar*. Berkeley: University of California Press.

Li, Yen-hui Audrey (1999) "Plurality in a classifier language," *Journal of East Asian Linguistics* 8, 75-99.

Loke, Kit-Ken (1994) "Is *Ge* Merely a 'General Classifier'?" *Journal of the Chinese Language Teachers Association* 29, 35-50.

MacWhinney, Brian, Jared Leinbach, Roman Taraban, and Janet McDonald (1989) "Language Learning: Cues or Rules?" *Journal of Memory and Language* 28, 255-277.

Marcus, Gary F., Ursula Brinkmann, Harald Clahsen, Richard Wiese, and Steven Pinker (1995) "German Inflection: The Exception That Proves the Rule," *Cognitive Psychology* 29, 189-256.

Myers, J., C-C Chiang, and S-P Gong (1999) "Rules vs. Analogy in Mandarin Classifier Selection," ms., National Chung Cheng University; to appear in *Language and Linguistics* 1.

Pinker, Steven (1989) *Learnability and Cognition*, MIT Press, Cambridge, Massachusetts.

Pinker, Steven (1991) "Rules of language," *Science* 253, 530-535.

Pinker, Steven (1994) *The Language Instinct: How the Mind Creates Language.* William Morrow and Company.

Prasada, Sandeep, and Steven Pinker (1993) "Similarity-Based and Rule-Based Generalizations in

Inflectional Morphology," *Language and Cognitive Processes* 8, 1-56.
Rosch, E. (1973) "Natural categories," *Cognitive Psychology* 4, 328-350.
Shi, Yu-Zhi (1996) "Proportion of Extensional Dimensions: The Primary Cognitive Basis for Shape-Based Classifiers in Chinese," *Journal of Chinese Language Teachers Association* 31, 37-59.
Tai, James H-Y. (1992) "Variation in Classifier Systems across Chinese Dialects: Towards a Cognition-Based Semantic Approach," *Chinese Languages and Linguistics 1: Chinese Dialects (= Symposium Series of the Institute of History and Philology No. 2, Academia Sinica)*, pp. 587-608.
Tai, James H-Y. (1994) "Chinese Classifier Systems and Human Categorization." In M. Y. Chen and O. Tzeng (eds.), *In Honor of William S.-Y. Wang: Interdisciplinary Studies on Language and Language Change*, Pyramid Press, Taipei, pp. 479-494.
Tai, James H-Y. and Chao F.-Y. (1994) "A Semantic Study of the Classifier *Zhang*," *Journal of Chinese Language Teachers Association* 29, 67-78.
Tai, James H-Y. and Wang L. (1990) "A Semantic Study of the Classifier *tiao*," *Journal of Chinese Language Teachers Association* 25, 35-56.
Tyan Yih-Min (1996) *Hanyu Fenleici yu Mingci de Yuyi Queren Licheng zhi Fenxi* [The Analysis of Semantic Verification Processes between Classifiers and Nouns in Mandarin Chinese], MA thesis, National Chung Cheng University.
Wang, Shaoxin (1989) "Liangci 'Ge' zai Tangdai Qianhou de Fazhan" [The Development of the Classifier 'Ge' around the Tang Dynasty], *Yuyan Jiaoxue yu Yanjiu* 1989, 98-119.
Wiebusch, Thekla (1995) "Quantification and Qualification: Two Competing Functions of Numeral Classifiers in the Light of the Radical System of the Chinese Script," *Journal of Chinese Linguistics* 23, 1-41.
Xu, F. and S. Pinker (1995) "Weird past tense forms," *Journal of Child Language* 22, 531-556.
Zubin, David A. and Klaus-Michael Köpcke (1984) "Affect Classification in the German Gender System," *Lingua* 63, 41-96.
Zubin, David A. and Klaus-Michael Köpcke (1986) "Gender and Folk Taxonomy: The Indexical Relation between Grammatical and Lexical Categorization," in C. Craig (ed.), *Noun Classes and Categorization*, John Benjamins Publishing Company, Amsterdam, pp. 139-180.
Zubin, David A. and Mitsuaki Shimojo (1993) "How 'General' Are General Classifiers? With Special Reference to *Ko* and *Tsu* in Japanese," *Proceedings of the Berkeley Linguistics Society* 19, 490-502.

Appendix

Exemplar nouns used to define nonce items in Experiment 2
(for each noun, a collocating classifier is given in parentheses)

Exemplar nouns	Nonce forms
(a) Non-*ge* items	
gou 'dog' (zhi)	lu3jie2
bi 'writing implement' (zhi)	kui2nie4
shuzhi 'branch' (gen)	fen2yan2
shengzi 'rope' (tiao)	xiao1zha2
yifu 'clothing' (jian)	ka3li2
laoshi 'teacher' (wei)	man2pie3
mi 'rice grain' (li)	juan1huai2
jingzi 'mirror' (mian)	bo1'e4
yezi 'leaf' (pian)	tong2re3
zhi 'paper' (zhang)	duan1qing4
(b) *Ge* items (human)	
huairen 'bad guy'	da1'nan3
laoren 'old person'	xing2gang4
nianqingren 'young person'	shan4ke1
meimei 'little sister'	fen1zhua1
heshang 'monk'	yong1'ou1
guer 'orphan'	mi2ye2
nongfu 'farmer'	lei4wei4
shuidiangong 'plumber-electrician'	jiang1la4
qigai 'beggar'	ke4su1
liulanghan 'vagrant'	xun2zong3
(c) Ambiguous items	
feiji 'airplane' (jia), qiche 'automobile' (liang), lunchuan 'ship' (sao)	wen3hun2
mobu 'rag' (tiao), saoba 'broom' (zhi), shuitong 'bucket' (ge)	jia2pi3
yusan 'umbrella' (ba), yuyi 'raincoat' (jian), yuxie 'galoshes' (shuang)	kan1'gao1
pifeng 'cape' (jian), maozi 'hat' (ding), pidai 'belt' (tiao)	qu2qian2

shu 'book' (ben), chi 'ruler' (zhi), xiangpica 'eraser' (ge)	du4ga1
he 'river' (tiao), shan 'mountain' (zuo), yun 'cloud' (pian)	ji4qi2
zidan 'bullet' (ke), qiang 'gun' (ba), gangkui 'army helmet' (ding)	pian1dou3
shi 'poetry' (shou), ci '(poetry form)' (shou), hua 'painting' (zhang)	wo1'ai3
binggan 'cookie' (kuai), tangguo 'candy' (ke), bingbang 'popsicle' (zhi)	luo2lai4
dipian 'film' (juan), xiangji 'camera' (tai), zhaopian 'photograph' (zhang)	xiu1lu2
(d) Vague items (no exemplar nouns)	
qing1suo3	
he2chi2	
tao1yu1	
sha1'na4	
ya4'nuo4	
duo4yan2	
cu4se4	
jing1yuan3	
she2fang1	
sou3cuo4	