

Decoding the silence: Neural bases of zero pronoun resolution in Chinese

Shulin Zhang^{a,*}, Jixing Li^b, Yiming Yang^c, John Hale^a

^a University of Georgia, United States

^b New York University Abu Dhabi, United Arab Emirates

^c Jiangsu Normal University, China

ARTICLE INFO

Keywords:

zero pronoun
pro-drop
Chinese
MVPA
fMRI
Left Temporal Lobe
Parahippocampal Gyrus

ABSTRACT

Chinese is one of many languages that can drop subjects. We report an fMRI study of language comprehension processes in these “zero pronoun” cases. The fMRI data come from Chinese speakers who listened to an audiobook. We conducted both univariate GLM and multivariate pattern analysis (MVPA) on these data. We found increased left Temporal Lobe activity for zero pronouns compared to overt subjects, suggesting additional effort searching for an antecedent during zero pronoun resolution. MVPA further revealed that the intended referent of a zero pronoun can be decoded in the Parahippocampal Gyrus and the Precuneus shortly after its presentation. This highlights the role of memory and discourse-level processing in resolving referential expressions, including unspoken ones, in naturalistic language comprehension.

1. Introduction

Many East Asian languages, such as Mandarin Chinese, are known as radical *pro*-drop language, in which pronouns can be freely omitted in both subject and object positions given proper discourse context (Huang, 1989). These “zero pronouns” have been extensively studied in formal linguistics (e.g. Barbosa, 2011; Barbosa, 2019; Bi & Jenks, 2019; Li & Thompson, 1976; Neeleman & Szendrői, 2007; Song, 2005), yet the neural bases of zero pronoun resolution are barely discussed, especially with naturalistic stimuli that can reveal language processes at the discourse level. Here we report the results of the first fMRI (functional Magnetic Resonance Imaging) study to examine the brain regions involved in zero pronoun processing while Chinese participants listen to an audiobook.

Within linguistics, the status of zero pronouns is debated (Frascarelli & Casentini, 2019). One view has it that they are essentially the same as their overtly-expressed counterparts (e.g. Neeleman & Szendrői, 2007; Tomioka, 2003). Another view derives zero and overt pronouns from semantically-distinct noun phrases (e.g. Bi & Jenks, 2019). These alternative conceptualizations are, of course, at the level of competence grammar. However, they suggest alternative processes of omission, copying, or inference that speakers might employ during comprehension. This study examines brain responses to both zero and non-zero arguments. To the extent that the brain responses to both types are similar, it supports the view that zero pronouns are indeed analogous to

overt pronouns. To the extent that those brain responses differ, different processing mechanisms may be at work. This would raise questions about whether and how zero pronouns should be related to their overt counterparts.

1.1. Zero pronouns in Chinese

As a radical *pro*-drop language, Chinese can have a null pronoun as the subject and/or object of a clause in appropriate contexts. This is shown in example 1 below from Huang (1989).

- (1) Zero pronouns in Chinese can show up at subject and/or object positions:
“Zhangsan kanjian Lisi le ma?” (“Did Zhangsan see Lisi?”)
a. “Ta kanjian ta le.” (“He saw him.”)
b. “[] kanjian ta le.” (“[He] saw him.”)
c. “Ta kanjian [] le.” (“He saw [him].”)
d. “[] kanjian [] le.” (“[He] saw [him].”)

Compared to other *pro*-drop languages such as Spanish and Italian, Chinese lacks morphological markers for person or gender information. This information, evidently, is not the only basis on which people can recover co-reference relationships. To characterize the panoply of information that does make such resolution possible, the concept of a topic chain has been advanced (Tsao, 1977; Shi, 1989; for a review see Pu, 2019).

A topic chain is a chain of clauses sharing an identical topic that

* Corresponding author.

E-mail address: shulin.zhang@uga.edu (S. Zhang).

occurs overtly once in one of the clauses, and its boundary may cross several sentences and even paragraphs (Li, 2004). The topic chain can integrate information from multiple clauses (Sun, 2019), which makes long-distance coreference between zero pronouns and overt noun phrases possible. We can understand coreference resolution as searching for an appropriate antecedent in a topic chain. This search process likely recruits memory and discourse-related brain regions. Pu (2019, page 190) summarizes matters in this way: “The close tie between discourse anaphora and the storage, capacity and processing of working memory has been abundantly documented in the literature of discourse analysis and psycholinguistics.” If this processing demands additional effort, then we should also expect higher brain activation levels, compared to overt noun phrases.

Another a priori expectation about zero pronoun resolution is that it operates retrospectively, at the earliest point where the omission can be detected. This is illustrated in example (2). Both 2a and 2b share the same first clause. Up to “yinwei (since)” the reference of missing element [] remains undetermined. However, at the onset of subsequent Verb Phrases (i.e. at “shi (is)” and “sheng (get)”) listeners are able to determine that the zero pronoun refers in (2a) to the temporal deictic word “today” and in (2b) to the already-mentioned proper name “Zhangsan”. This kind of reasoning motivates the decision to measure zero pronoun processing at the onset of the main verb, in the fMRI study to be detailed in Section 2.5 and 2.6.

1.2. Brain regions involved in reference processing

While no prior neuroimaging study has directly investigated zero pronoun processing, there are some fMRI and MEG studies on referential processing in general (Brodbeck & Pyllkänen, 2017; Brodbeck, Gwiliams, & Pyllkänen, 2016; Hammer, Goebel, Schwarzbach, Münte, & Jansma, 2007; Li et al., 2021; Matchin, Sprouse, & Hickok, 2014; Nieuwland, Petersson, & Van Berkum, 2007; Santi & Grodzinsky, 2012). However, no consensus has been reached on the neural correlates for pronoun processing. In addition, previous studies adopted different task manipulations, making it unclear whether they tapped the same cognitive processes. For example, Nieuwland et al. (2007) compared the BOLD responses when participants read sentences containing a referentially failing pronoun (e.g., “Rose told Emily that *he* had a positive attitude towards life.”) or a coherent pronoun (e.g. “Ronald told Emily that *he* had a positive attitude towards life.”). Nieuwland et al. showed that referentially failing pronouns were associated with increased activation in the medial parietal regions and bilateral inferior parietal regions, possibly reflecting morphosyntactic processing. Hammer et al. (2007) manipulated syntactic gender matching between the antecedent and pronouns using German sentences and found that gender incongruity elicited the bilateral Inferior Frontal Gyrus (IFG), the left Medial Frontal Gyrus (MFG), and the bilateral Supramarginal/Angular Gyrus compared to congruent pronoun-antecedent pairs. Hammer et al. (2011) further investigated possible interactions between gender, and distance between the antecedents and the pronoun. The results identified a fronto-temporal network including the bilateral IFG, the Superior Temporal Gyrus (STG), and posterior Middle Temporal Gyrus (pMTG) for long-distance conditions, with the pMTG additionally driven by syntactic gender violation. These authors suggested that the temporal regions are sensitive to morphosyntactic information of the antecedents since the long distance between the antecedent and the pronoun

increased the overall syntactic complexity of the sentence. Matchin et al. (2014) also examined the effect of distance but with the backward anaphora/filler-gap dependencies contrast. Matchin and colleagues observed specific activity in the bilateral Anterior Temporal Lobes, the bilateral Angular Gyrus (AGs), and the left Precuneus activity during the processing of backward anaphora compared to *wh*-fillers. Santi and Grodzinsky (2012) compared null pronouns, a parasitic-gap and a *wh*-trace in English sentences such as “[Which paper] did the tired student submit [*wh*-trace] after reviewing [parasitic gap/it]?”. The results identified increased activity in the right Middle Frontal Gyrus (MFG), the left Ventral Precentral Sulcus, and the Left Supramarginal Gyrus for pronouns compared to parasitic gaps.

In addition to morphosyntactic manipulations, Brodbeck and Pyllkänen (2017) and Brodbeck et al. (2016) used a visual world paradigm in magnetoencephalography (MEG) and found medial parietal activity in cases of successful reference resolution. More relevant to the current study is Li et al.’s (2021) study on third-person pronoun processing using the same naturalistic listening paradigm. In both fMRI and MEG, Li et al. found that the left middle temporal gyrus (LMTG) is consistently activated for third-person pronoun processing in both English and Chinese. Yet they also found additional medial parietal activity from the MEG data, consistent with Brodbeck and Pyllkänen (2017), Brodbeck et al. (2016), and Nieuwland et al. (2007).

To sum up, referential processing has been implicated in a number of regions, including the medial parietal lobe. Zero pronoun resolution, as a special case of referential processing, is expected to involve similar brain regions.

1.3. Current study

The current study examines which brain regions are responsible for the processing of the dropped pronouns in Chinese. Based on the features of zero pronoun introduced in the previous section, this study focuses on zero pronouns that show up at the subject positions in a naturalistic material. We compared Chinese listeners’ brain activation using fMRI towards zero and non-zero arguments while listening to a naturalistic material with resolution process time-locked to the main verbs. (See Section 2.2 for details on the annotation steps).

In a mass univariate analysis with a General Linear Model (GLM), we show that zero pronoun resolution prompts more activity in the Left Temporal Lobe (including LITG/LMTG, Planum Temporale), compared to overt reference resolution (see Section 3.1). This confirms earlier work (e.g. Hammer et al., 2007; Hammer et al., 2011; Li et al., 2021) and supports the view of zero pronoun resolution as an operation performed, at least in part, by the core language network.

With searchlight-based Multivariate Pattern Analysis (MVPA), we identify other brain areas that are part of an “extended” language network, one that extends beyond core areas like the Temporal Lobe mentioned above (Ferstl et al., 2008; Xiong & Newman, 2021). Ferstl’s meta-analysis (Ferstl et al., 2008) compared language comprehension with resting and perceptual baselines, and identified a bilateral fronto-temporal brain network for language comprehension, which was broader than the left temporal region. The MVPA analysis in this study identifies brain regions where above-chance decoding for zero pronouns is possible including areas that have been implicated in semantic memory such as the Parahippocampal Gyrus, the Precuneus, and bilateral Angular Gyrus (e.g. Burianova et al., 2010; Maguire & Frith, 2004;

(2) The reference of the missing subject pronoun varies depending on information from the following Verb Phrase:

a.	Zhangsan	jintian	xiuxi,	yinwei	[]	shi	xingqitian.	
	Zhangsan	today	rest	since	[today]	is	Sunday	
	“Zhangsan rests today, since [today] is Sunday.”							
b.	Zhangsan	jintian	xiuxi,	yinwei	[]	sheng	bing	le.
	Zhangsan	today	rest	since	[Zhangsan/he]	get	sick	LE
	“Zhangsan rests today, since [he] got sick.”							

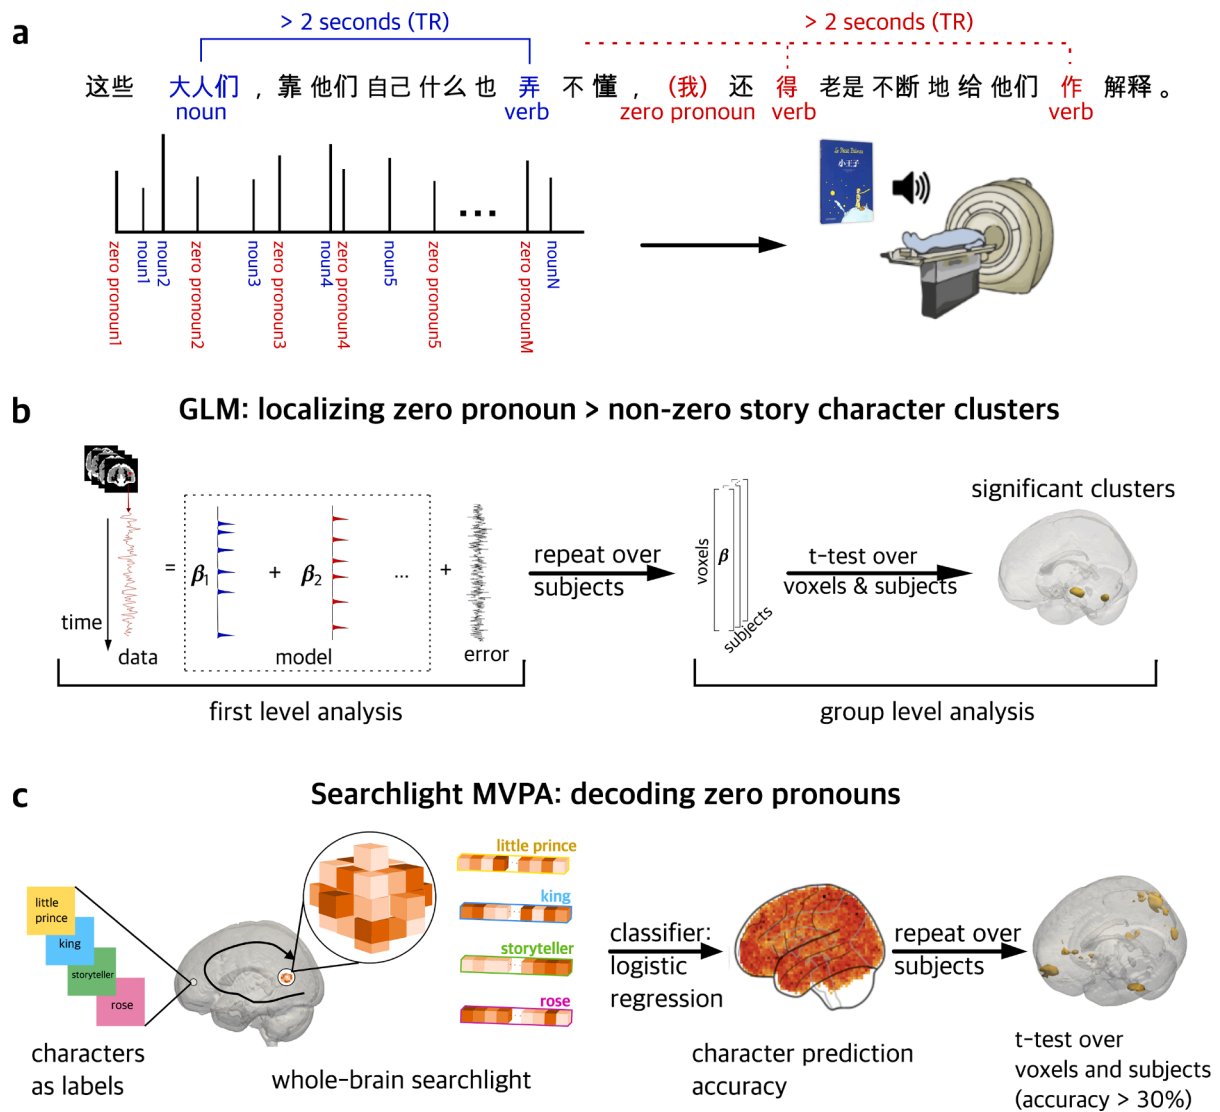


Fig. 1. Schematic illustration of the analysis procedure. **a.** Stimuli and fMRI data collection. fMRI data were collected while Chinese native speakers were listening to a naturalistic audiobook. Zero pronouns and Non-zero nouns in subject position are annotated by Chinese native speakers, and their corresponding main verbs were taken as timestamps for GLM and MVPA analyses. The distance between a zero or non-zero noun to its main verb was controlled to be longer than 2 second so that they cannot be in the same fMRI scan. **b.** Two-stage General Linear Model analyses. At the first stage, a general linear model was fitted to each participants' fMRI data, and the regressors used in the model include audio sound pressure, word frequency, zero-pronoun feature, non-zero noun feature, verb syllable count and distance between the subject and its main verb (see Section 2.5). At the second stage, a *t*-test was performed on the distribution of β values across subjects and voxels, and the significant clusters were retrieved with $p < .05$ FWE and $k > 20$. **c.** Whole-brain searchlight multivariate pattern analysis. Among all annotated zero pronouns, four story characters' (see Section 2.6) main verb scans were used in the MVPA decoding analyses. A logistic regression classifier was used to derive an average accuracy value for decodability of the four story characters, based on an *N*-voxel neighborhood. *T*-tests were performed on these accuracy values across subjects to identify clusters with above-chance accuracy. ($p < .001$ FWE and $k > 50$).

Manns et al., 2003 111 and see MVPA results in Table 1). Another MVPA, with overt nouns, converges on many of the same brain regions as the zero pronoun MVPA analysis. The overall methodology is sketched in Fig. 1 and detailed below in Section 2.

2. Materials and methods

2.1. Participants

Participants were 35 healthy, right-handed, young adults (15 female, mean age = 19.3, range = 18–25). They self-identified as native Chinese speakers and had no history of psychiatric, neurological, or other medical illnesses that could compromise cognitive functions. They were living in China at the time of scanning. The participants all started learning English in junior high school. All participants were paid for and

gave written informed consent prior to participation, in accordance with the guidelines of the Ethics Committee at Jiangsu Normal University.

2.2. Stimuli and annotations

The stimulus is a Chinese translation (xiaowangzi.org., 2021) of Saint-Exupéry's *The Little Prince*. To annotate zero pronouns and non-zero nouns, we first located all verbs (i.e., "VV"s) in the text using ZPar (Zhang & Clark, 2011). We then annotated each verb as Zero or Non-zero based on whether it has an overtly pronounced subject. For example, as shown in Fig. 1a, the verb phrase (VP) "弄不懂 (make no understanding)" is marked as non-zero as its subject "大人们 (grown-ups)" is overt; the VP "做解释 (make explanations)" is marked as "zero" since its subject "我 (I)" is omitted.

2.3. Procedure

After giving their informed consent, participants were familiarized with the MRI facility and assumed a supine position on the scanner. The presentation script was written in PsychoPy 2 (Peirce, 2007). Auditory stimuli were delivered through MRI-safe, high-fidelity headphones (Ear Bud Headset, Resonance Technology, Inc, California, USA) inside the head coil. The headphones were secured against the plastic frame of the coil using foam blocks. An experimenter increased the sound volume stepwise until the participants could hear clearly.

The Chinese audiobook lasted for about 99 min and was divided into nine sections, each lasted for about ten minutes. Participants listened passively to the nine sections and completed four quiz questions after each section (36 questions in total). These questions were used to confirm their comprehension and were viewed by the participants via a mirror attached to the head coil and they answered through a button box. The entire session lasted for around 2.5 h.

2.4. fMRI data collection and preprocessing

MRI images were acquired with a 3T MRI GE Discovery MR750 scanner with a 32-channel head coil. Anatomical scans were acquired using a T1-weighted volumetric Magnetization Prepared Rapid Gradient-Echo (MP-RAGE) pulse sequence. Functional scans were acquired using a multi-echo planar imaging (ME-EPI) sequence with on-line reconstruction (TR = 2000 ms; TEs = 12.8, 27.5, 43 ms; FA = 77°; matrix size = 72 x 72; FOV = 240.0 mm x 240.0 mm; 2 x image acceleration; 33 axial slices, voxel size = 3.75 x 3.75 x 3.8 mm). Cushions and clamps were used to minimize head movement during scanning.

All fMRI data were preprocessed using AFNI version 16 (Cox, 1996). The first 4 volumes in each run were excluded from analyses to allow for T1-equilibration effects. Multi-echo independent components analysis (ME-ICA) (Kundu et al., 2012) was used to denoise data for motion, physiology, and scanner artifacts. Images were then spatially normalized to the standard space of the Montreal Neurological Institute (MNI) atlas, yielding a volumetric time series resampled at 2 mm cubic voxels.

2.5. GLM analysis

A whole-brain GLM analysis was conducted to localize the brain regions involved in zero and non-zero reference resolution. We modeled the timecourse of each voxel's BOLD signals for each of the nine sections using binary regressors that pick out zero pronouns (510 cases) and their non-zero counterparts (1942 cases) in the audiobook. These regressors are time-locked to the onset of the following verb (complete details about these verbs are provided in Supplementary Table 2 and 3). Because the stimulus is a literary text, the co-occurrence of particular verbs in "zero" or "non-zero" configuration cannot be strictly controlled. However, it appears the two groups are approximately matched (see Supplementary Fig. 9).

The GLM analysis included five control variables: the root mean square intensity (*RMS intensity*) for every 10 ms of each audio section, the binary regressor time-locked to the offset of each word in the audio (*word rate*), the unigram frequency of each word (*frequency*), estimated using Google ngrams (Version 20120701) and the SUBTLEX corpora for Chinese (Cai & Brysbaert, 2010), the distance between the verb and its syntactic subject (*distance*), and the number of syllables for each verb (*verb syllable*). These regressors were convolved with SPM12's (Penny et al., 2011) canonical HRF function and matched the scan numbers of each section. (see Supplementary Fig. 6 for the correlation matrix of the regressors, and Supplementary Fig. 5 for a visualization of the regressors.) At the group level, the contrast images for zero and non-zero pronouns were examined by a factorial design matrix. An 8 mm full-width at half-maximum (FWHM) Gaussian smoothing kernel was applied on the contrast images from the first-level analysis to counteract inter-subject anatomical variation. Significant clusters were thresholded

at $p < .05$ FWE with cluster size of $k > 20$. The GLM analysis was performed with the python package nilearn (0.7.0) (Abraham et al., 2014). All the FWE corrections used in this study were voxel-level. p -value thresholds used Bonferroni correction. Cluster size thresholds were realized with the nilearn software package.

2.6. MVPA analyses

Searchlight MVPA identifies voxels where the pattern of activation in its local neighborhood can discriminate between conditions (i.e. reference to story characters). For each participant, a spherical ROI (radius = 8 mm) centered in turn on each voxel in the brain scans time-locked to 5 seconds after the zero pronouns' presentation. A 5-second delay serves to capture BOLD signals at approximately the peak of their hemodynamic response to the zero pronoun resolution process. At the request of a reviewer, an MVPA analysis using a 6-second delay was also carried out (see the result in Supplementary Fig. 8). Each vector contains all the voxels in each sphere without feature selection. A logistic regression classifier was trained to differentiate the vectors of all four story characters. A 3-fold cross-validation process was adopted in the training process, which means 2/3 of the original labeled data were used as a training dataset, and the rest as a testing set. Prediction accuracy was averaged over the three testing results.

This whole process was repeated for the sphere for each participant. The resulting maps contain each voxel's decoding accuracy for each participant. Higher accuracy indicates better performance on decoding the reference of the zero pronouns. At the group level, a t -test was conducted for all voxels across all subjects. For the zero pronoun MVPA, voxels with an accuracy higher than 30% (higher than the chance baseline 25%) were highlighted. The empirical distribution of references to story characters in zero pronoun contexts is unbalanced as expected in naturalistic texts (see Section 2.6 for details). To help interpret accuracy levels, weighted logistic regression was applied in MVPA such that examples were weighted according to the prevalence of each class in the training data. This was realized by the scikit-learn (Pedregosa et al., 2011) `class_weight = "balanced"` option. With this parameter, the model automatically assigns class weights that are inversely proportional to those classes' respective frequencies. A greater weight applies to the minority class, resulting in a larger error calculation, and in turn, more updates to the model coefficients (King & Zeng, 2001). The average accuracy from guessing randomly according to the empirical distribution in this weighted problem is 25%. As shown in Fig. 7 a3, the MVPA threshold was set the first 3/4 quartile performance level across all subjects and voxels, which resulted in the 30% threshold for the zero pronoun condition, and 33% for the non-zero condition. Family-wise error correction was applied with an alpha level of $<.001$ and an adequate cluster size of $k > 50$. The MVPA analysis was performed using the python packages nilearn (0.7.0) (Abraham et al., 2014), and scikit-learn (Pedregosa et al., 2011).

2.6.1. MVPA for zero pronoun resolution

A whole-brain searchlight MVPA was performed to discriminate patterns of activation pertaining to the omitted story characters. The fMRI scans which contain both a zero pronoun and its previous overtly pronounced antecedent are excluded from the MVPA. We selected the four most frequent story characters for the classification, and there were 188 zero-pronoun instances used in MVPA, including: "小王子 (the little prince)", 84 instances; "我 (I/the storyteller)", 67 instances; "国王 (the king)", 25 instances; "花 (the rose)", 12 instances.

2.6.2. MVPA for non-zero subject resolution

Although the brain mechanism for non-zero subject resolution was not the main focus of this study, this supplementary analysis for non-zero subject MVPA could provide the angle of comparison between zero and non-zero resolution. We selected the four story characters that had been used in the zero pronoun MVPA, and there were 188 zero-

pronoun instances in total for the non-zero subject MVPA, including: “小王子 (the little prince)”, 578 instances; “我 (I/the storyteller)”, 331 instances; “国王 (the king)”, 69 instances; “花 (the rose)”, 59 instances. The decoding accuracy was examined at the threshold level of 33%, FWE $p < 0.001$, cluster size > 50 .

2.6.3. Control MVPA: Scrambled MVPA for zero pronoun resolution

To characterize chance performance, we carried out another MVPA analysis on a scrambled dataset where story character labels were assigned randomly. In this control analysis, a story character label randomly selected out of the four story characters was assigned to each zero pronoun. The same MVPA analysis steps introduced above were conducted to test whether there are any brain regions afford to decode of these randomly assigned labels. The story character cases ended up as: “小王子 (the little prince)”, 59 instances; “我 (I/the storyteller)”, 47 instances; “国王 (the king)”, 43 instances; “花 (the rose)”, 39 instances. (see results in Section 3.4) The decoding accuracy was examined at the threshold level of 30%, FWE $p < 0.05$, cluster size > 20 .

2.6.4. Control MVPA: Frequent character MVPA for zero pronoun resolution

Since this study is based on naturalistic materials, the occurrence of verbs with particular story characters is not strictly controlled. To assess the possibility that decoding might be possible on the basis of idiosyncrasies in the joint distribution of verbs and story characters, a further MVPA analysis was carried out, and a new regressor “frequent character” was used instead of the original story character. The steps of creating this “frequent character” regressor were: Firstly, for each zero subject’s main verb, its corresponding zero subjects in the whole story were saved in a list; Secondly, the story characters’ occurrences in this list was calculated; Last, the most frequent story character corresponding to this verb was assigned as the “frequent character”. For example, the verb “说道”(means *speak*) has 12 zero pronoun cases in the whole story: “the little prince” 8 times, “the storyteller” 2 times, “the king” 1 time, and “the rose” 1 time, and the most frequent character “the Little Prince” was assigned to all the 12 “说道” as its “frequent character” for this control analysis. Therefore, instead of decoding the correct zero pronoun reference, this analysis decoded the most frequent story character used for each main verb. The story character cases that ended up for this analysis were: “小王子 (the little prince)”, 59 instances; “我 (I/the storyteller)”, 47 instances; “国王 (the king)”, 43 instances; “花 (the rose)”, 39 instances. (see results in Section 3.5) The decoding accuracy was examined at the threshold level of 30%, FWE $p < 0.05$, cluster size > 20 .

3. Results

3.1. GLM: Localizing brain regions for zero pronoun processing

The contrast between zero pronouns and overt references to story

characters revealed significantly higher activity in the anterior and posterior left Temporal Lobe (see Fig. 2a,b). Based on the Harvard-Oxford Cortical Structural Atlas, the anterior Temporal Lobe region belongs to Inferior Temporal Gyrus (ITG)/ MTG, whereas the posterior Temporal Lobe region belongs to the Planum Temporale. The MNI coordinates for the peak of each cluster are shown in Fig. 2c. No significant cluster was found for the opposite contrast, i.e. Non-zero $>$ Zero.

3.2. MVPA: Decoding references of zero pronouns

Searchlight MVPA results are shown in Fig. 3. Brain regions with a decoding accuracy greater than 30% for the zero pronouns include the Precuneus, the LMFG, the right Inferior Temporal Gyrus (RITG), the LMTG, the left Frontal Pole, bilateral Angular Gyri, and the right Parahippocampal Gyrus.

Both the GLM and MVPA results implicate the LMTG. Not only did the LMTG show higher activity for zero pronoun resolution compared to non-zero reference resolution, but also showed high story character decoding accuracy for story characters. MVPA further revealed a network for decoding zero pronouns, including the Precuneus, the LAG, the Frontal Pole, the LMFG, and the Parahippocampal Gyrus.

3.3. MVPA: Decoding references of non-zero subjects

Searchlight MVPA results for non-zero reference decoding are shown in Fig. 4. Brain regions with a decoding accuracy greater than 33% for the zero pronouns include bilateral ITG, bilateral Frontal Pole, the right Superior Temporal Lobule, the left Frontal Orbital Cortex, bilateral Superior Frontal Gyrus (SFG), and the Cerebellum.

3.4. Scrambled Zero Pronoun MVPA: Decoding zero pronouns with randomly assigned story character labels

When the correct story character labels are replaced by randomly assigned story character labels, no significant brain region was detected by MVPA ($p < .05$ FWE, cluster size > 20) when the decoding accuracy threshold set to 30%. This null result with randomly assigned labels suggests that the main MVPA analysis in Section 3.2 is indeed identifying brain regions where story character information is represented.

3.5. Frequent Character Zero Pronoun MVPA: Decoding the story character used most frequently together with the verb

As the correct story character labels are replaced by the most frequent story character bonded with the main verbs, no significant brain region was detected from the MVPA results ($p < .05$ FWE, cluster size > 20) with decoding accuracy threshold set as 30%. The null result of decoding frequently used story character label indicates that the MVPA decoding process did not rely on the “verb-story character” joint distribution in the whole story, and further supports the idea that the




Fig. 2. GLM results for the contrast between zero and non-zero reference resolution. **a** Whole-brain view on a 3D brain. **b** Coronal slices of significant clusters. Significant clusters were thresholded at $p < .05$ FWE and $k > 20$ (see Table 1 for their MNI coordinates, cluster sizes and peak level statistics).




Fig. 3. Zero pronoun MVPA results for brain regions with decoding accuracy significantly higher than 30%. **a** Whole-brain view on a 3D brain. **b** Coronal slices of significant clusters. Significant clusters were thresholded at $p < .001$ FWE and $k > 50$ (see Table 1 for their MNI coordinates, cluster sizes and peak level statistics).




Fig. 4. Non-zero subject MVPA results for brain regions with decoding accuracy significantly higher than 33%. **a** Whole-brain view on a 3D brain. **b** Coronal slices of significant clusters. Significant clusters are thresholded at $p < .001$ FWE and $k > 50$ (see Table 1 for their MNI coordinates, cluster sizes and peak level statistics).

Table 1

MNI coordinates, cluster sizes and peak level statistics for analyses: GLM (see Fig. 2), zero pronoun MVPA (see Fig. 3), and non-zero pronoun MVPA (see Fig. 4).

Analysis Threshold	Cluster	MNI Coordinates	Peak t-value	Cluster Size (mm3)
GLM: zero > non-zero pronoun $p < 0.05, k > 20$	posterior Temporal Lobe (Including Planum Temporale)	-50,-58,-14	6.31	1424
	anterior Temporal Lobe (Including LITG/ LMTG)	-60,-20,6	6.13	1552
Zero pronoun MVPA $p < 0.001, k > 50$	RTPrecuneus	5,-74,47	12.9	3027
	LTPrecuneus	-11,-61,60	11.78	5676
	LMFG	-24,43,37	11.76	1608
	RITG	33,-45,-23	11.15	2176
	LAG	-43,-80,34	10.78	1955
	LMTG	-65,-30,-7	10.55	2522
	Frontal Pole	-27,65,9	10.29	2680
	RAG	52,-74,22	10.01	1671
	Parahippocampal Gyrus	27,2,-29	9.27	2712
	LITG	-52,-42,-19	13.56	2176
	Frontal pole	24,68,3	12.67	4446
	RITG	43,5,-48	12.39	1955
	Superior Parietal Lobule	30,-55,63	11.64	2239
Non-zero pronoun MVPA $p < 0.001, k > 50$	Frontal Orbital Cortex	-21,15,-19	11.55	1576
	Frontal Pole	-5,59,-23	11.13	2018
	RSFG	24,15,66	10.95	2018
	LSFG	-5,46,41	10.38	1986
	Cerebellum	14,-23,-35	10.13	1608

above-chance decoding performance in the main MVPA results in Section 3.2 does identify brain regions where story character information is represented.

4. Discussion

This study examines the neural bases of zero pronoun resolution in Chinese. Chinese is especially suitable for studying zero pronoun resolution because it does not have verbal inflections that might drive pronoun resolution at the morphosyntactic level. The GLM results show increased left temporal activity during zero pronoun resolution compared to non-zero reference processing. MVPA results further reveal a network of activity including the Parahippocampal Gyrus and the Precuneus in addition to the “core” language network. Zero pronoun resolution seems to require additional effort, compared to overt noun phrases. Both “core” and “extended” nodes of the language network (in the sense of Ferstl et al., 2008) appear to contribute to resolving the reference of zero pronouns. Table 2 summarizes the findings.

4.1. Core language network

Both anterior and posterior regions within the Left Temporal Lobe showed significantly higher activity for zero pronouns compared to overt references to story characters, including the LITG/ LMTG and the Planum Temporale.

Table 2

Significant brain regions found in the GLM and MVPA analyses.

Analysis	Significant Brain Regions
GLM Zero > Non-zero	Planum Temporale, LITG/LMTG
GLM Non-zero > Zero	Not found
MVPA Zero Subjects	Bilateral Precuneus, Right Parahippocampal Gyrus, LMFG, RITG, Bilateral AG, LMTG, Left Frontal Pole
MVPA Non Zero Subjects	LITG, RITG, Left Frontal Pole, Right Superior Parietal Lobule, Left Frontal Orbital Cortex, Bilateral SFG, Cerebellum
Control MVPA Zero Subjects - Scrambled and Frequent	Not found

Prior studies have highlighted the essential role of left temporal regions in language comprehension (e.g. [Dronkers & Turken, 2011](#); [Matchin & Hickok, 2020](#)). These same regions have been specifically associated with gender features during the processing of pronouns. For example, [Hammer et al. \(2007\)](#) showed that German sentences with congruent biological and syntactic gender evoked higher activity in the left temporal lobe. [Miceli et al. \(2002\)](#) found increased LITG and LMTG activity when the subjects were asked whether a written noun has a masculine or feminine gender. However, [Li et al. \(2021\)](#) using the same naturalistic paradigm as the current paper, showed that the anterior LMTG is also implicated for pronoun processing in Chinese. In addition, [Li, Jin, and Tan \(2004\)](#) showed a number of brain regions including the LITG and the LMTG during a lexical-judgment task while Chinese participants saw nouns, verbs, and noun/verb-ambiguous words, supporting the LITG and the LMTG's role in lexical representation. [Ferstl et al.'s \(2008\)](#) meta-analysis, introduced above in Section 1.3, further suggests that the anterior temporal lobe contributes to the comprehension of coherence, and shows stable significant results as part of the “core” language network.

In the context of this existing evidence, increased activity in the left inferior and middle temporal gyrus for zero pronouns in the current study could simply reflect greater task difficulty, rather than a qualitatively different mode of processing for zero pronouns. However the MVPA results, to be discussed below, rule out this interpretation.

4.2. Extended language network, in cooperation with the core

Whole-brain searchlight-based MVPA for zero pronouns revealed a network of brain regions implicated in the comprehension of references to story characters, including the bilateral Precuneus, the bilateral AG, the left Frontal Pole, the LMFG, the LMTG, the RITG, and the right Parahippocampal Gyrus (see [Fig. 3](#)). Among these brain regions, the RITG and the left Frontal Pole are found in both zero and non-zero MVPA results. The Precuneus, the Parahippocampal Gyrus, bilateral AG, the LMTG, and the LMFG were found only for zero pronouns (see [Fig. 4](#)). These results are all based on a presumed hemodynamic delay of five seconds; increasing this delay to six seconds yields similar results, such that story characters are decodable in both in the Parahippocampal Gyrus and the RITG (see results in Supplementary Fig. 8).

The Parahippocampal Gyrus has been previously implicated in discourse-level language processing (e.g. [Allendorfer et al., 2012](#); [Wallentin, Østergaard, Lund, Østergaard, & Roepstorff, 2005](#)) and semantic memory processing (e.g. [Bartha et al., 2003](#); [Burianova et al., 2010](#); [Maguire & Frith, 2004](#); [Manns et al., 2003](#)). For example, [Allendorfer et al. \(2012\)](#) showed higher Parahippocampal Gyrus activity while participants were generating verbs silently for a given noun. [Wallentin et al. \(2005\)](#) observed right Parahippocampal Gyrus activity while processing real motion sentences (e.g. “the man goes through the house”) and fictive motion sentences (e.g. “the trail goes through the house”). The decodability of story characters in the Parahippocampal Gyrus further suggests that the search for antecedents involves discourse-level language processing. The Parahippocampal Gyrus has also been reported to support semantic memory retrieval and semantic verbal memory processing ([Bartha et al., 2003](#)). In [Bartha et al.'s](#) fMRI study, subjects performed a semantic decision task while they heard spoken concrete nouns designating objects and made a decision on whether these objects were available in the supermarket and their costs compared to certain amounts. [Bartha et al.](#) observed activation in the Parahippocampal Gyrus, along with the medial and inferior temporal lobe, implicating these regions in semantic verbal memory processing. These previous results support the assignment of the Parahippocampal Gyrus to the extended language network. In the present study, above-chance decoding for zero pronouns in the Parahippocampal Gyrus could be explained in terms of semantic memory for story characters. As introduced in Section 1.1, topic chains among Chinese clauses may cross sentence or even paragraph boundaries so this sort of semantic memory

is indeed discourse-level. The status of Chinese as a radical *pro*-drop language is consistent with this memory-oriented interpretation, insofar as there are relatively few linguistic cues available to drive a stimulus-based (as opposed to memory-based) coreference resolution system.

Similar to the Parahippocampal Gyrus, the Precuneus has also been found to be relevant to “extra-linguistic” processing such as discourse-level information integration and memory retrieval ([Bhattasali et al., 2019](#); [Diachek et al., 2020](#); [Foudil, Kwok, & Macaluso, 2020](#); [Mashal, Vishne, & Laor, 2014](#); [Wehbe et al., 2020](#)). [Foudil et al. \(2020\)](#), for example, showed that brain activation level in the Precuneus was modulated by storyline consistency, suggesting its role in discourse information integration. [Bhattasali et al. \(2019\)](#) using the same naturalistic listening paradigm, reported that the right Precuneus was correlated with the retrieval of stored expressions. On the other hand, the Precuneus has also been characterized as a “processing core”, one that connects with the MTG and the AG and integrates multiple brain functions, including memory retrieval ([Mar, 2011](#)).

Connectivity analyses have shown that both of these regions communicate with left temporal regions (see e.g. [Xu et al., 2015](#); [Xu et al., 2019](#)). From this standpoint, story-character decodability of the sort reported above in SubSection 3.2 could reflect a qualitatively different, discourse-oriented mode of processing for zero pronouns.

Apart from the Precuneus and the Parahippocampal Gyrus, we also identified a number of brain regions within the core language network. The left Angular Gyrus (AG) is classically seen as essential for meaning-related aspects of language comprehension. (e.g. [Bonner, Peelle, Cook, & Grossman, 2013](#); [Humphreys, Ralph, & Simons, 2021](#); [Price, Bonner, Peelle, & Grossman, 2015](#); [Ramanan, Pigué, & Irish, 2018](#); [Seghier, 2013](#)). For example, using Transcranial Magnetic Stimulation (TMS), [Branzi, Pobric, Jung, and Lambon Ralph \(2021\)](#) found that the left AG is critical for integrating context-dependent information during language processing. Moreover, [Davis and Yee \(2019\)](#) suggested that the left AG's connectivity to the hippocampal regions underpins its essential role in processing semantic roles. The Frontal Pole is part of the language network's ventral pathway ([Brauer, Anwander, Perani, & Friederici, 2013](#)). The left Medial Frontal Gyrus (LMFG), especially BA46, is critical for auditory language comprehension ([Dronkers & Turken, 2011](#)). [Wu, Ho, and Chen's \(2012\)](#) meta-analysis of fMRI studies implicates the LMFG specifically in Chinese language processing. The RITG has also been associated with noun processing by [Crepaldi et al. \(2013\)](#). So, zero pronoun processing does not localize exclusively to memory-related regions — core language network regions seem to be involved as well.

To sum up, resolving zero pronouns in Chinese calls upon both the core language network and extended language network regions such as the Precuneus and Parahippocampal Gyrus. The involvement of these regions means that zero pronoun processing is discourse-related in a way that overt pronoun processing is not. This finding supports topic chains and other theories which go beyond morphosyntactic deletion.

5. Conclusion

This study examines the neural bases of zero pronoun processing in Chinese. By comparing fMRI BOLD responses for zero pronoun processing with that of non-zero reference processing during naturalistic listening, we show that zero pronoun resolution elicits increased activity in left temporal regions, nodes of the core language network. By decoding brain activity patterns for zero pronouns that refer to story characters, we identify other regions like the Precuneus and the Parahippocampal Gyrus that are outside the core language network. These regions do not seem to contain story character information when the referring expression is overt. We interpret these findings as suggesting a distinct mode of discourse-related processing for zero pronoun resolution.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant No. 1903783. The authors would like to thank Robert Mason, Leila Wehbe for valuable discussion.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.bandl.2021.105050>.

References

- Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., & Varoquaux, G. (2014). Machine learning for neuroimaging with scikit-learn. *Frontiers in neuroinformatics*, 8, 14.
- Allendorfer, J. B., Lindsell, C. J., Siegel, M., Banks, C. L., Vannest, J., Holland, S. K., & Szafarski, J. P. (2012). Females and males are highly similar in language performance and cortical activation patterns during verb generation. *Cortex*, 48(9), 1218–1233.
- Barbosa, P. P. (2011). Pro-drop and theories of pro in the minimalist program part 1: Consistent null subject languages and the pronominal-Agr hypothesis. *Language and Linguistics Compass*, 5(8), 551–570.
- Barbosa, P. P. (2019). Pro as a minimal nP: Toward a unified approach to pro-drop. *Linguistic Inquiry*, 50(3), 487–526.
- Bartha, L., Brenneis, C., Schocke, M., Trinka, E., Köylü, B., Trieb, T., & Benke, T. (2003). Medial temporal lobe activation during semantic language processing: fMRI findings in healthy left-and right-handers. *Cognitive Brain Research*, 17(2), 339–346.
- Bhattachali, S., Fabre, M., Luh, W. M., Al Saied, H., Constant, M., Pallier, C., & Hale, J. (2019). Localising memory retrieval and syntactic composition: an fMRI study of naturalistic language comprehension. *Language, Cognition and Neuroscience*, 34(4), 491–510.
- Bi, R. A., & Jenks, P. (2019). Pronouns, null arguments, and ellipsis in Mandarin Chinese. In M. Espinal, E. Castroviejo, M. Leonetti, L. McNally, & C. Real-Puigdollers (Eds.), *Proceedings of sinn und bedeutung 23* (vol. 1, pp. 127–142).: Universitat Autònoma de Barcelona.
- Bonner, M. F., Peelle, J. E., Cook, P. A., & Grossman, M. (2013). Heteromodal conceptual processing in the angular gyrus. *Neuroimage*, 71, 175–186.
- Branzi, F. M., Pobric, G., Jung, J., & Lambon Ralph, M. A. (2021). The left angular gyrus is causally involved in context-dependent integration and associative encoding during narrative reading. *Journal of Cognitive Neuroscience*, 1–14.
- Brauer, J., Anwender, A., Perani, D., & Friederici, A. D. (2013). Dorsal and ventral pathways in language development. *Brain and language*, 127(2), 289–295.
- Brodbeck, C., Gwilliams, L., & Pykkänen, L. (2016). Language in context: MEG evidence for modality-general and -specific responses to reference resolution. *eNeuro* 3, (e0145-16.2016 1-16).
- Brodbeck, C., & Pykkänen, L. (2017). Language in context: Characterizing the comprehension of referential expressions with MEG. *NeuroImage*, 147, 447–460.
- Burianova, H., McIntosh, A. R., & Grady, C. L. (2010). A common functional brain network for autobiographical, episodic, and semantic memory retrieval. *Neuroimage*, 49(1), 865–874.
- Cai, Q., & Brysbaert, M. (2010). Subtlex-ch: Chinese word and character frequencies based on film subtitles. *Plos ONE*, 5, e10729.
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research*, 29(3), 162–173.
- Crepaldi, D., Berlinger, M., Cattinelli, I., Borghese, N. A., Luzzatti, C., & Paulesu, E. (2013). Clustering the lexicon in the brain: a meta-analysis of the neurofunctional evidence on noun and verb processing. *Frontiers in human neuroscience*, 7, (303).
- Davis, C. P., & Yee, E. (2019). Features, labels, space, and time: Factors supporting taxonomic relationships in the anterior temporal lobe and thematic relationships in the angular gyrus. *Language, Cognition and Neuroscience*, 34(10), 1347–1357.
- Diachek, E., Blank, I., Siegelman, M., Affourtit, J., & Fedorenko, E. (2020). The domain-general multiple demand (md) network does not support core aspects of language comprehension: a large-scale fMRI investigation. *Journal of Neuroscience*, 40(23), 4536–4550.
- Dronkers, N. F., & Turken, A. U. (2011). The neural architecture of the language comprehension network: converging evidence from lesion and connectivity analyses. *Frontiers in systems neuroscience*, 5, 1.
- Ferstl, E. C., Neumann, J., Bogler, C., & Von Cramon, D. Y. (2008). The extended language network: a meta-analysis of neuroimaging studies on text comprehension. *Human brain mapping*, 29(5), 581–593.
- Foudil, S. A., Kwok, S. C., & Macaluso, E. (2020). Context-dependent Coding of Temporal Distance between Cinematic Events in the Human Precuneus. *Journal of Neuroscience*, 40(10), 2129–2138.
- Frascarelli, M., & Casentini, M. (2019). The interpretation of null subjects in a radical pro-drop language: Topic chains and discourse-semantic requirements in Chinese. *Studies in Chinese Linguistics*, 40(1), 1–46.
- Hammer, A., Goebel, R., Schwarzbach, J., Münte, T. F., & Jansma, B. M. (2007). When sex meets syntactic gender on a neural basis during pronoun processing. *Brain Research*, 1146, 185–198.
- Hammer, A., Jansma, B. M., Tempelmann, C., & Münte, T. F. (2011). Neural mechanisms of anaphoric reference revealed by fMRI. *Frontiers in Psychology*, 2, 1–9.
- Huang, C. T. J. (1989). Pro-drop in Chinese: A generalized control theory. In O. Jaeggli, & K. Safir (Eds.), *The null subject parameter* (pp. 185–214). Springer.
- Humphreys, G. F., Ralph, M. A. L., & Simons, J. S. (2021). A unifying account of angular gyrus contributions to episodic and semantic cognition. *Trends in Neurosciences*.
- King, G., & Zeng, L. (2001). Logistic regression in rare events data. *Political analysis*, 9(2), 137–163.
- Kundu, P., Inati, S. J., Evans, J. W., Luh, W. M., & Bandettini, P. A. (2012). Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI. *Neuroimage*, 60(3), 1759–1770.
- Li, C. N., & Thompson, S. A. (1976). Subject and topic: A new typology. In C. N. Li (Ed.), *Subject and topic* (pp. 457–489). New York, USA: Academic Press.
- Li, J., Wang, S., Luh, W. M., Pykkänen, L., Yang, Y., & Hale, J. (2021). Cortical processing of reference in language revealed by computational models. *bioRxiv* 2020.11.24.396598.
- Li, P., Jin, Z., & Tan, L. H. (2004). Neural representations of nouns and verbs in Chinese: an fMRI study. *Neuroimage*, 21(4), 1533–1541.
- Li, W. (2004). Topic chains in Chinese discourse. *Discourse Processes*, 37(1), 25–45.
- Maguire, E. A., & Frith, C. D. (2004). The brain network associated with acquiring semantic knowledge. *Neuroimage*, 22(1), 171–178.
- Manns, J. R., Hopkins, R. O., & Squire, L. R. (2003). Semantic memory and the human hippocampus. *Neuron*, 38(1), 127–133.
- Mar, R. A. (2011). The neural bases of social cognition and story comprehension. *Annual Review of Psychology*, 62, 103–134.
- Mashal, N., Vishne, T., & Laor, N. (2014). The role of the Precuneus in metaphor comprehension: evidence from an fMRI study in people with schizophrenia and healthy participants. *Frontiers in human neuroscience*, 8, 818.
- Matchin, W., & Hickok, G. (2020). The cortical organization of syntax. *Cerebral Cortex*, 30(3), 1481–1498.
- Matchin, W., Sproue, J., & Hickok, G. (2014). A structural distance effect for backward anaphora in Broca's area: An fMRI study. *Brain and Language*, 138, 1–11.
- Miceli, G., Turriziani, P., Caltagirone, C., Capasso, R., Tomaiuolo, F., & Caramazza, A. (2002). The neural correlates of grammatical gender: An fMRI investigation. *Journal of Cognitive Neuroscience*, 14, 618–628.
- Neeleman, A., & Szendrői, K. (2007). Radical pro drop and the morphology of pronouns. *Linguistic Inquiry*, 38(4), 671–714.
- Nieuwland, M. S., Petersson, K. M., & Van Berkum, J. J. (2007). On sense and reference: Examining the functional neuroanatomy of referential processing. *NeuroImage*, 37(3), 993–1004.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in python. *The Journal of machine Learning research*, 12, 2825–2830.
- Peirce, J. W. (2007). PsychoPy—Psychophysics software in python. *Journal of neuroscience methods*, 162(1–2), 8–13.
- Penny, W., Friston, K., Ashburner, J., Kiebel, S., & Nichols, T. (2011). *Statistical parametric mapping: The analysis of functional brain images*. Academic Press.
- Price, A. R., Bonner, M. F., Peelle, J. E., & Grossman, M. (2015). Converging evidence for the neuroanatomic basis of combinatorial semantics in the Angular Gyrus. *Journal of Neuroscience*, 35(7), 3276–3284.
- Pu, M. M. (2019). *Zero anaphora and topic chain in Chinese discourse* (C. Shei, Ed.). Routledge.
- Ramanan, S., Pigué, O., & Irish, M. (2018). Rethinking the role of the angular gyrus in remembering the past and imagining the future: the contextual integration model. *The Neuroscientist*, 24(4), 342–352.
- Santi, A., & Grodzinsky, Y. (2012). Broca's area and sentence comprehension: A relationship parasitic on dependency, displacement or predictability? *Neuropsychologia*, 50, 821–832.
- Seghier, M. L. (2013). The angular gyrus: multiple functions and multiple subdivisions. *The Neuroscientist*, 19(1), 43–61.
- Shi, D. (1989). Topic chain as a syntactic category in Chinese. *Journal of Chinese Linguistics*, 17, 223–261.
- Song, Z. (2005). A comparative study of subject pro-drop in old Chinese and Modern Chinese. University of Pennsylvania Working Papers. *Linguistics*, 10(2), 18.
- Sun, K. (2019). The integration functions of topic chains in Chinese discourse. *Acta Linguistica Asiatica*, 9(1), 29–57.
- Tomioka, S. (2003). The semantics of Japanese null pronouns and its cross-linguistic implications. *The interfaces: Deriving and interpreting omitted structures*, 61, 321.

- Tsao, F. (1977). *A functional study of topic in Chinese: The first step towards discourse analysis (Unpublished doctoral dissertation)*. Los Angeles, California: USC.
- Wallentin, M., Østergaard, S., Lund, T. E., Østergaard, L., & Roepstorff, A. (2005). Concrete spatial language: See what I mean? *Brain and language*, *92*(3), 221–233.
- Wehbe, L., Blank, I.A., Shain, C., Futrell, R., Levy, R., von der Malsburg, T., & Fedorenko, E. (2020). Incremental language comprehension difficulty predicts activity in the language network but not the multiple demand network. *bioRxiv*.
- Wu, C. Y., Ho, M. H. R., & Chen, S. H. A. (2012). A meta-analysis of fMRI studies on Chinese orthographic, phonological, and semantic processing. *Neuroimage*, *63*(1), 381–391.
- xiaowangzi.org. (2021). 小王子网站. <http://www.xiaowangzi.org/>. (Accessed: 2021-04-03).
- Xiong, Y., & Newman, S. (2021). Both activation and deactivation of functional networks support increased sentence processing costs. *Neuroimage*, *225*, 117475.
- Xu, J., Lyu, H., Li, T., Xu, Z., Fu, X., Jia, F., & Hu, Q. (2019). Delineating functional segregations of the human middle temporal gyrus with resting-state functional connectivity and coactivation patterns. *Human brain mapping*, *40*(18), 5159–5171.
- Xu, J., Wang, J., Fan, L., Li, H., Zhang, W., Hu, Q., & Jiang, T. (2015). Tractography-based parcellation of the human middle temporal gyrus. *Scientific Reports*, *5*(1), 1–13.
- Zhang, Y., & Clark, S. (2011). Syntactic processing using the generalized perceptron and beam search. Retrieved from *Computational Linguistics*, *37*(1). https://doi.org/10.1162/coli_a_00037. <http://aclweb.org/anthology/J11-1005>.