PROSODY-LIKE ASYMMETRIES IN HANDWRITTEN CHINESE CHARACTERS*

James Myers
National Chung Cheng University, Taiwan

1. Introduction

Prosodic asymmetry, where units differ in prominence depending on their position, is observed not just in spoken languages (e.g., left-headed metrical feet as in the English word *ápple*), but in sign languages as well (Van der Kooij and Crasborn 2008). This suggests that its principles and processes do not depend on modality, but reflect something deeper about how human language is structured in the mind. Writing systems, typically considered artificial, are nevertheless highly structured, and they are also learned, perceived, and produced in tandem with natural language processing. Such considerations have led some scholars to apply the tools of grammatical analysis to writing systems as well (Fuhrhop and Peters 2023; Meletis 2020). Among the insights arising from such analyses is the idea that writing systems also have prosody (Evertz 2018; Gnanadesikan 2023). In this paper we first review arguments that prosodic asymmetry is also found in Chinese character form (section 2), and then provide new evidence from a corpus analysis of a standard Chinese font (section 3) and from a handwriting experiment (section 4). These results demonstrate that asymmetries in Chinese characters are indeed prosody-like, based on their interactions with other aspects of character structure and the effects of lexical frequency on them.

2. The grammar of Chinese characters

Building on previous analyses of Chinese character structure, from traditional (Qiu 2000) through modern formalisms (Wang 1983), Myers (2019) attempts to capture all of the major patterns in Chinese characters within a coherent grammatical framework, supported by evidence from quantitative corpus analysis, psycholinguistic experimentation, and other sources. Central to this framework is the notion of Chinese character prosody and its interactions with other aspects of character structure (section 2.1). These interactions make a number of predictions that are the focus of the present study (section 2.2).

2.1 Character prosody and its interactions

Because Chinese characters are stored in the mental lexicon rather than created on the fly, their grammar is analogous to lexical morphology and phonology rather than syntax. Myers

* This study was funded by National Science and Technology Council (Taiwan) grants 109-2410-H-194-096-MY3 and 112-2410-H-194-030-MY3. I am also grateful to Thomas Bishop, Chen Yangru, our experimental participants, and visitors to my posters at the 46th annual meeting of the German Linguistic Society (corpus analysis) and the 2024 annual conference of the Canadian Linguistic Association (experiment).

Actes du congrès annuel de l'Association canadienne de linguistique 2024.
Proceedings of the 2024 annual conference of the Canadian Linguistic Association.

(2019) argues that the orthographic analog of morphology is that aspect of a writing system that is externally interpreted. Unlike the case in spoken and signed languages, orthographic components can be interpreted not just in meaning but also in pronunciation. For example, in the characters in (1), the analog to morphology relates to the indicated components, rather than the individual strokes, because only the components are interpretable (see Watt 1975 for a similar idea applied to alphabetic orthography).

(1) a. 接 *jiē* 'connect' (扌 'hand' + 妾 *qiè*)[1]
b. 動 *dòng* 'move' (重 *zhòng* + 力 'strength')
c. 掃 *sǎo* 'sweep' (扌 'hand' + 帚 'broom')
d. 林 *lín* 'forest' (木 'tree')

Thus defined, character morphology then turns out to share another key property with its natural language namesake: it is associated with operations analogous to affixation, compounding, and reduplication. Affixation is what creates the phono-semantic characters in (1a, b), where a semantic radical, taken from a small inventory and abstractly related to the whole character's meaning, is appended (usually on the left, as in (1a)) to a component from a much larger inventory (and sometimes itself morphologically complex) that indicates the character's pronunciation (sometimes obscured by diachronic sound change), and then reduced in size. The semantic radical is thus affix-like in being closed-class, semantically abstract, relatively fixed in location, and formally reduced. Semantic compounding, as in (1c), involves components interpreted only for their meaning with few restrictions on position, and reduplication, as in (1d), involves copying of individual morphemes in fixed arrangements to indicate a fixed set of iconic meanings; both of these operations are thus similar to their namesakes in natural language as well. Affixation is by far the most productive of these morphological operations, with compounding a distant second. That is, not only do affixed characters dominate the lexicon, but virtually all new character coinages are of this type as well (see corpus analyses in Myers 2019).

Unlike character morphology, character phonology does not involve external interpretation (of course it is also silent, like sign phonology). Individual strokes are definable in terms of distinctive features (see also Wang 1983 and Peng 2017) that are distributed and alternate in regular ways (see also Wang 1983), and so Myers (2019) treats stroke-level patterns as analogous to segmental phonology (Myers 2021 goes on to argue that strokes also combine into structures analogous to syllables).

Character prosody interfaces between character morphology and segmental phonology, being dependent in part on the former and crucial to shaping the latter. For example, note that all of the characters in (1) show an asymmetry in component width, with all but (1c) enlarging the component on the right. This size contrast is especially clear when we compare the sizes of the same component in different positions, as in the reduplicated form in (1d) and in the phono-semantic characters in (2): the components are "stressed" on the right and "unstressed" elsewhere. There is a similar size asymmetry along the vertical

[1] Because this study was conducted in Taiwan, all characters in this paper are traditional; see Myers (2019) for arguments that the modern Chinese system of simplified characters has virtually the same grammar.

axis, with the larger (stressed) component on the bottom, as shown in (3), but horizontally arranged characters are far more common.

(2) a. 媽 *mā* 'mother' (女 'female' + 馬 *mǎ*)
b. 騎 *qí* 'ride (horse)' (馬 'horse' + 奇 *qí*)

(3) a. 尖 *jiān* 'sharp' (小 'small' + 大 'big') [compounding]
b. 奇 *qí* 'strange' (大 'big' + 可 *kě*) [affixation]
c. 炎 *yán* 'blazing' (火 'fire') [reduplication]

Size per se is not lexically distinctive in Chinese characters, which leads Myers (2019) to consider it an aspect of character phonetics rather than character phonology, much as amplitude is a phonetic correlate of stress in spoken languages but is not itself encoded phonologically. Also like phonetics is the presumed motivation for prominence on the right and bottom, namely the same articulatory process that lengthens syllables in word- and phrase-final position in spoken and sign languages (Beckman and Edwards 1990; Sandler 1993). Stroke order in characters generally runs from left to right and top to bottom, and so generally the stroke at the lower right is the last-written one. Nevertheless, this phonetic influence is primarily historical; there is no merely motoric force that can cause an entire multi-stroke component to enlarge merely because it is written last, let alone force mechanically printed fonts to conform to this pattern.

Another reason Myers (2019) considers the size asymmetry to be a synchronic reflex of abstract prosodic asymmetry (right-headed and bottom-headed "feet") is that, as mentioned earlier, character prosody also affects segmental (i.e., stroke-level) phonology. This can be seen in examples (1a) and (1c), which contain the segmentally reduced component in (4b). The loss of the top stroke is idiosyncratic, being unique to this one component. Nevertheless, the idiosyncratically reduced allomorph is restricted to the left unstressed position, with the full form in (4a) appearing elsewhere, as in the stressed bottom position in the character in (4c). Moreover, the change of the lowest horizontal stroke in (4a) to a diagonal in (4b) is a regular alternation that affects many components, though again only in the non-head position, as illustrated in (5). Myers (2019) provides a number of other arguments for the notion of character prosody, including binary reduplicative templates that only permit doubling horizontally and/or vertically, and the restriction of stroke curving to weak prosodic positions, as modulated by component width.

(4) a. 手 *shǒu* 'hand' (as free character and full combining form)
b. 扌 'hand' (reduced bound form)
c. 拿 *ná* 'hold' (合 'shut' + 手 'hand') [compounding]

(5) a. 地 *dì* 'earth' (土 'earth' + 也 *yě*) [affixation]
b. 物 *wù* 'thing' (牛 'ox' + 勿 *wù*) [affixation]

Character prosody is distinct from character morphology, as shown by the enlargement of the lowest stroke in the monomorphemic character in (4a) (i.e., individual morphemes also have prosodic structure). There is also idiosyncratic reduction in both the affixed character in (1a) and in the compound in (1c), so both morphological operations permit similar prosodic patterns. However, prosody and morphology also interact, as shown by the small size of the affix (semantic radical) in (1b), despite the fact that it appends to the right rather than the left. Moreover, affixes tend to have fewer strokes than non-affixes (in prosodic terms, they are light rather than heavy), even if they have no other allomorph, as illustrated by the character containing a bound affix in (6). This is analogous to the tendency for affixes in spoken and signed languages to be light and intrinsically unstressable, like the English agentive suffix *-er*.

(6) 冷 *lěng* 'cold' (bound 冫 'ice' + 令 *lìng*) [affixation]

Character prosody and its interactions with character morphology and segmental (stroke-level) character phonology have been formalized in a lexical phonology framework by Myers (2024) and tested in corpus analyses and experiments reported in Myers (2016, 2019). The present study tests a number of new quantitative predictions, as explained next.

2.2 Predictions for printed and handwritten characters

The most fundamental prediction is that the character analogs of morphology, segmental phonology, prosody, and phonetics should all be intrinsically distinct from each other, even if they also interact. In quantitative terms this implies that it should be possible to tease them apart in multiple regression analyses. In particular, the effects of prosody on segmental phonology (component stroke number) should differ from its effects on phonetics (component size), and the effects of morphology (affixation vs. compounding and, for affixed characters, affix-like semantic components vs. non-affixes) should be distinguishable from the purely prosodic effect of component position (unstressed left vs. stressed right).

Subtler predictions come from the different degrees with which morphology, segmental phonology, and phonetics are lexicalized, as diagnosed by the effects of token frequency (see Bybee 2006). Natural language morphology, like that of English, tends to be more irregular in higher-frequency words (*blow* ~ *blew*) than in lower-frequency ones (*flow* ~ *flowed*), since repeated use protects against the forgetting of exceptions, while harder-to-recall words get regularized by default. Myers (2019) observes something like this pattern in the distribution of idiosyncratic (i.e., irregular) reduction in Chinese characters, which in affixed characters is more likely the higher the character's token frequency. That is, for highly lexicalized processes like morphology or idiosyncratic lexical phonology, frequency effects are negative: the higher a lexical item's token frequency, the less likely it will be regular. At the same time, however, prosodic regularities in natural languages target more common words (*mémory* ~ *mém'ry* vs. *mámmary* ~ **mámm'ry*) through the automatization of articulatory routines. This implies a positive frequency effect: the more common a lexical item, the more likely it will undergo phonetically motivated phonology. Finally, phonetics is less influenced by rote memory than morphology and phonology (though not entirely uninfluenced, as shown by word-specific phonetics: Pierrehumbert 2002), and so tends to show little to no effect of frequency at all.

Going beyond Myers (2019), then, for Chinese characters we predict negative frequency effects for the more lexicalized patterns of character morphology and segmental phonology, perhaps particularly for compounds, which are associated with a much less productive morphological operation than affixed characters. By contrast, we should see positive frequency effects for the right-headed prosodic structure pervasive throughout the character inventory, but only in its effects on segmental phonology (including stroke counts); the prosodic correlate of component size, being analogous to phonetics, should show not show frequency effects.

These predictions could be tested using any number of different data sources. This study focuses on just two: a corpus of characters in a standard Chinese font family, and a handwriting experiment. Even mechanically printed fonts evolve naturally in a sense, as multiple human minds work to ensure some degree of commensurability with other fonts and with handwriting, and handwriting itself is likely to be influenced to some degree by the experience of reading print, given that even the most prolific writers read a lot more than they write. Corpus analyses can also include a wider variety of items than can be accommodated in a practical experiment, including many unfamiliar to the typical participant but nevertheless shaped by the same linguistic forces, at least diachronically, as experimentally testable items. Moreover, corpus-based variables can be factored out in the analysis of the experimental results, to see if handwriting patterns go beyond rote memorization of font forms. The following two sections, then, report corpus-based and experimental results, respectively.

3. Quantitative patterns in a standard Chinese font

In Chinese printing, by far the most commonly used family of fonts (or, more technically, typefaces) is called Song, after the Chinese dynasty (more formally it is called Fangsong 'imitation Song', and is sometimes also called Ming, after a later dynasty). A font from this serif family is what is used for the Chinese examples throughout this paper, as in most books, newspapers, or other printed material. As with other writing systems, fonts from sans-serif families have become more common in online text, but while fonts do differ subtly in their character grammars (see Myers 2019), the patterns discussed in this paper are consistent across all fonts.

Aside from its ubiquity in print, there is another particular advantage to choosing a Song font for a corpus analysis: it is what is used in a large database of Chinese character forms. We first describe how we used this database in section 3.1, and then report the results of our corpus analysis in section 3.2.

3.1 Methods

The corpus analysis took advantage of the Character Description Language (CDL) database (Bishop and Cook 2007) used in the Wenlin Software for Learning Chinese (https://www.wenlin.com/). This database encodes over 80,000 characters with special-purpose XML tags that, among other things, indicate each character's components, and for single-component characters, that component's stroke types. Most relevant for the present analyses, it also encodes the horizontal and vertical dimensions of each component. For example, Figure 1 shows the dimensions of the two components in a semantic compound as expressed as boxes in Wenlin's visual interface to the CDL. These dimensions are actually encoded in the XML tags with the coordinates for the upper left and lower corners of each box, which here are (0, 0) and (52, 128) for the left component and (48, 8) and (128, 128) for the right component; the dimensions of the whole character area run from 0 in the upper left to 128 in the lower right.

Figure 1. Component heights and widths in the character 好 *hǎo* 'good' (女 'woman' + 子 'child') as represented in Wenlin's visual interface to the CDL.

The CDL codes can then be parsed automatically to extract the heights and widths of each component in any character. For this study we used our own scripts written in R (R Core Team 2024), which computed component widths in terms of the proportion of the total width (e.g., a raw width of 64 would represent .5 of the full 128-point width). Our R scripts also used CDL to identify characters with a binary structure, that is, comprised of two character morphemes or of one morpheme combined with what could itself form a full character, as is typical of phono-semantic characters. For example, the character in (2b) has a fundamentally binary structure because it is composed of a unitary component to the left of a (vertically arranged two-component) component. This is not only the most common type of character, but restricting ourselves to such characters also permits our analyses to use simple binary factors, as described below. Because character decomposition in CDL is etymological, characters like that in (7) could not be included in our analysis, because the historical decomposition has become opaque in the modern form.

(7) 施 *shī* 'exert' (etymologically 㫃 'flag' + 也 *yě*) [historically phono-semantic]

In addition to component dimensions, the analyses also depended on character type (i.e., morphological operation), whole-character token frequency, the number of strokes in each component, and (for affixed characters) affix position. Character type information was taken primarily from Wiktionary (https://www.wiktionary.org/), which distinguishes between phono-semantic characters (created through affixation) and semantic compounds; reduplicated characters were also identified, but they were so rare that they are not included in the analyses reported here. Character frequency was taken from Tsai (2006), which is based on a traditional Chinese corpus with 171,882,493 character tokens and 13,060 distinct characters. To compute component stroke counts, we made use of the traditional indexing radicals used for dictionary look-up, which are traditionally grouped by stroke number; for phono-semantic characters, the indexing radical is virtually always the affix (semantic radical), and for other character types it is arbitrary but conventionalized. The number of strokes in the character's other component was then simply the number of strokes in the indexing radical subtracted from the total stroke count. also available in Tsai (2006). Finally, affix position had already been coded manually for Myers (2019); this information also allowed our R scripts to automatically identify the arrangement of components in characters as vertical or horizontal.

From the initial set of 5,528 characters (affixed and compounded) for which we had all of the necessary information, we then extracted just the 3,511 with a horizontal arrangement. This is not only the most common type, but Myers (2019) also found that prosodic asymmetries along this axis were the most robust in previous corpus analyses and experiments; thus for the CDL-derived component dimensions, we only analyzed component widths, not heights. All this filtering gave us 3,343 affixed characters, of which 3,059 are left-affixed and 284 are right-affixed, as well as 168 semantic compounds, which ranged in frequency from 0.2 to 380,384 tokens per million.

3.2 Results and discussion

Component stroke counts and widths in affixed (phono-semantic) characters and semantic compounds were analyzed in separate linear regressions. The dependent variables were the ratio of the right component (in stroke counts or width) to the left component minus one. This variable was chosen so that components with the same number of strokes or widths have a score of zero (the null hypothesis) and higher values indicate that the right component has more strokes or is wider. All analyses had lognormed character frequency as a predictor. For affixed characters another predictor was the position of the affix (semantic radical), as well as its interaction with frequency, and for the analysis of component width, the number of component strokes was an additional predictor.

The results for the affixed characters are illustrated in Figure 2. For component stroke counts, shown in the left panel, there was a significant and positive intercept ($B = 0.37$, $t(3339) = 4.38$, $p < .0001$), indicating the dominance in the corpus of the canonical right-headed prosodic structure. There was also a significant positive effect of affix position, with left-affixed characters showing a stronger right-headed tendency ($B = 0.81$, $t(3339) = 9.73$, $p < .0001$); indeed, as can be seen by the fact that the darker line in the left panel is below zero, right-affixed characters tended to have more strokes in the left component. There was no significant effect of character frequency or its interaction with affix position ($p > .3$). For component widths, shown in the right panel, the intercept was again significant and positive ($B = 0.34$, $t(3338) = 7.91$, $p < .0001$), even with the positive effect of the number of component strokes factored out ($B = 0.42$, $t(3338) = 47.44$, $p < .0001$). Affix position also had a significant effect ($B = 0.47$, $t(3338) = 10.86$, $p < .0001$), since affixes were narrower than non-affixes. Again there was no significant effect of character frequency or its interaction with affix position ($p > .1$), though Figure 2 shows that both stroke counts and widths showed nonsignificant trends towards positive frequency effects.

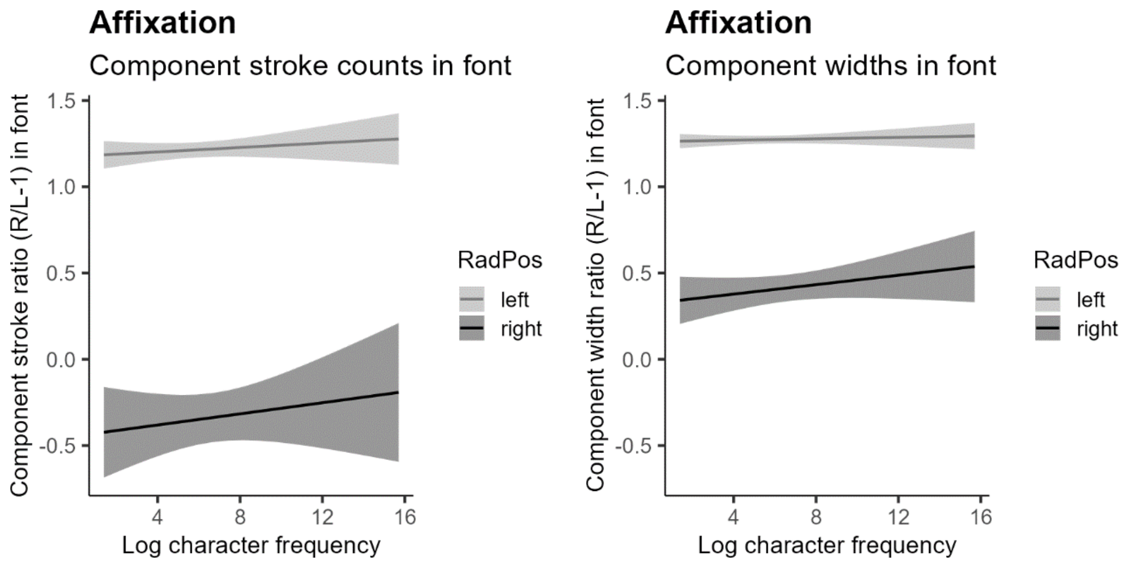

Figure 2. The effect in the standard font of the position of the affix (semantic radical) and character frequency on the relative number of component strokes (left panel) and on relative component width (right panel) in affixed (phono-semantic) characters.

Figure 3 illustrates the results for semantic compound characters. For component stroke counts, shown in the left panel, the significant positive intercept ($B = 0.97$, $t(166) = 4.72$, $p < .0001$) demonstrates that the dominance of the canonical right-headed prosodic structure applies to this character type as well. This time there was also a significant negative effect of character frequency ($B = -0.06$ $t(166) = -2.25$, $p < .05$), indicating that the less common a character, the more likely it is to have a canonical right-headed structure in terms of stroke counts. This is quite different from the nonsignificant positive trend seen in Figure 1 for affixed characters. For component widths, shown in the right panel, there was the usual significant positive intercept ($B = 0.37$, $t(166) = 3.20$, $p < .01$) and positive effect of the number of component strokes ($B = 0.58$, $t(166) = 13.95$, $p < .0001$), but no significant effect of character frequency ($p > .3$).

Note that the patterns in neither stroke counts nor width depended on idiosyncratic reduction, since the results were virtually identical when characters showing this kind of component reduction were dropped (i.e., in analyses with just 1,162 affixed characters and 84 compounds rather than the full data sets).

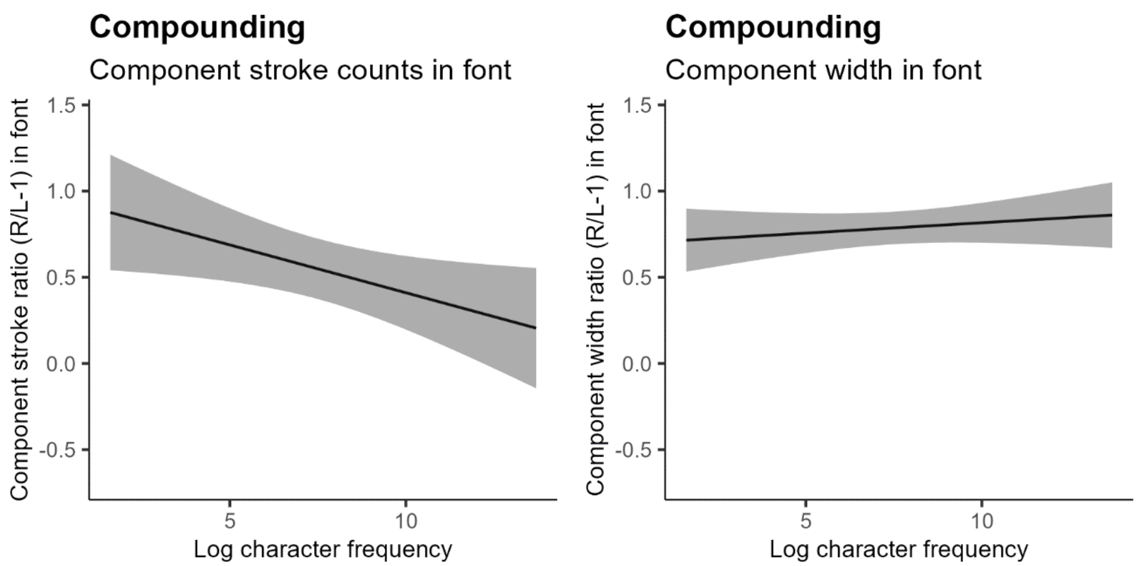

Figure 3. The effect in the standard font of character frequency on the relative number of component strokes (left panel) and on relative component width (right panel) in semantic compound characters.

Put together, these results are consistent with what we expect of character morphology, segmental phonology, prosodic phonology, and phonetics if they really behave like their namesakes in spoken and signed languages. Affixes (semantic radicals in phono-semantic characters) are prosodically weaker (smaller) both in terms of segments (stroke number) and phonetics (width). Independent of this morphological influence on prosody, both affixed and compounded characters show an overall tendency towards right-headed structures, again as reflected both in segmental phonology and phonetics. Frequency effects, which only reached statistical significance in compounds, affected segmental phonology but not phonetics. Specifically, they were negative, consistent with the idea that the more difficult a lexical item is to retrieve from memory, the more regular it is likely to become, which here meant that it is more likely to have the canonical right-headed prosodic structure. Affixed characters did not show this pattern, perhaps because they are derived via a far more productive morphological operation and so depend less on rote memory, but there was nevertheless a nonsignificant trend towards a positive frequency effect; if confirmed in further studies, this would imply a phonetically motivated regularization of prosody in affixed characters through use.

4. Patterns in handwriting

While the patterns in mechanically printed fonts presumably reflect the psychologically real processes that shaped them historically, we still need to determine if they are also mentally active in modern reader/writers. In this experiment we examined handwritten stroke counts and component widths in a component combination task. Ideally we would want to test all of the characters in the corpus analysis, but this is not practical. Not only is a handwriting task even more tiring than most psycholinguistic tasks, which tend to be limited to a few hundred items at most, but also, if our participants are meant to be representative, we cannot expect them to be familiar with characters below some frequency threshold. Both considerations mean that our sample will have to be relatively small and skewed toward sufficiently frequent characters, which, due to the great differences in productivity across the morphological operations in character formation, means that the proportion of characters that are not of the most common type (phono-semantic characters with left-affixation) will be too small for meaningful statistical analysis. In our experiment, then, we analyzed all character types together and ignored affix position as well. In other words, the experiment focused just on segmental phonology, prosody, and phonetics, not morphology. Section 4.1 describes the methods and section 4.2 the results.

4.1 Methods

As stimuli we chose 199 traditional characters with horizontal binary structure, all of which were included in the corpus analysis as well. None of the characters have idiosyncratic reduction. Because Hue (2003) estimates that the typical college student in Taiwan knows around 5,000 traditional characters, we selected our items from the top 5,000 most-frequent ones (based on Tsai 2006), so that the final set of characters ranged in frequency from 1 to 50,048 tokens per million. Of the 199 items, 182 were formed by affixation (169 left-affixed and 13 right-affixed), 13 were semantic compounds, and 4 were reduplicated; as noted above, this natural skewing made comparisons across character types impossible. Ideally we would have liked the two parts of the characters to have identical numbers of strokes and widths, based on the standard CDL font, so that any deviation from these values in handwriting would have to reflect productive processes, but this proved impossible. While the mean numbers of stroke counts were not significantly different across the left and right components, the trend was still towards the right component having more strokes (5.44 vs. 5.59, $t(396) = -0.72$, $p > .4$). The asymmetry in component widths was large and significant (.36 vs. .60, $t(396) = -30.06$, $p < .0001$). Though somewhat irritating for our experimental goals, these asymmetries at least help reconfirm the robustness of character prosody in this standard font.

The participants were 40 college students in Taiwan, two of whom were left-handed. The task was to combine visually presented components into a single character; for example, when shown the prompt in (8a), participants had to write the character in (8b) (see Figure 1 for the glosses). By prompting handwriting with visual characters, we reduced the burden on memory and the effects of non-orthographic confounds, and by displaying them in decomposed form rather than full characters, we avoided biasing participants with component asymmetries.

(8) a. 女+子
b. 好

The experiment was run using PsychoPy (Peirce et al. 2022), with the prompts displayed and the handwritten responses collected by a Wacom® One tablet and stylus.

After four practice trials, the main experiment with 199 trials began. Each trial started with a fixation symbol ‘o’ in the center of the screen (1 second), followed by the prompt (1 second), then a very brief visual mask to override any retinal traces (100 ms), and finally a 3 × 3 cm white square that remained for an unlimited amount of time, within which participants had to handwrite the full character. Participants ended each trial by tapping the stylus on a button labeled “OK” at the bottom of the screen. Even though there were only 199 experimental trials, the task was tiring enough for us to offer a break midway through. Participants required about 30 minutes to finish the experiment.

PsychoPy interpreted the signals from the tablet as if they came from a mouse, recording the coordinates, time points, and contact status of the stylus (treated as mouse button presses) approximately every 20 ms. R scripts were used to extract this information from the raw PsychoPy data files and converted them into images for visual checking; the scripts also counted strokes and their positions (see Myers 2023, for more information on how these scripts worked). The images allowed two Taiwanese lab assistants to check the handwritten forms for accuracy and to identify the first stroke of the second (right-side) component, information that was then used to compute the handwritten component widths.

4.2 Results and discussion

Six participants were dropped for having lower than 80% accuracy (one standard deviation below the mean), leaving 34 participants for analysis, including one left-hander. Trials with inaccurately written characters or handwriting that went outside the writing area were also dropped, leaving 6,146 trials for analysis. The dependent variables were the same as for the corpus analysis, namely the ratio of the number of strokes and width of the right component relative to that of the left component, minus one, so that intercepts significantly different from zero implied cross-component asymmetry. Separate linear mixed regression models were run to predict these component stroke or width ratios from lognormed character frequency and the same ratios in the standard font (i.e., font stroke ratios were used to predict handwriting stroke ratios and font widths were used to predict handwriting widths). The interactions between frequency and the font ratios were tested but proved nonsignificant, and so were dropped. All models included random intercepts for participants and characters and random slopes for participants, and p values were calculated via t tests using Satterthwaite’s method.

The results are shown in Figure 4. For component stroke counts, illustrated in the left panel, there was a significant positive intercept ($B = 0.40$, $t(135.32) = 4.39$, $p < .0001$), showing that there tended to be fewer handwritten strokes in the left than in the right component. This was so despite stroke counts in the standard font being factored out; these also had a positive effect on handwritten stroke counts ($B = 1.21$, $t(198.99) = 8.03$, $p < .0001$). Together these results indicate that participants dropped significantly more strokes in the left component (mean 1.08) than in the right component (mean 0.76). There was also a negative frequency effect ($B = -0.02$, $t(188.24) = -2.34$, $p < .05$), with left component strokes dropped more often in less common characters. Figure 5 shows an example of left-component stroke dropping in a relatively low-frequency character.

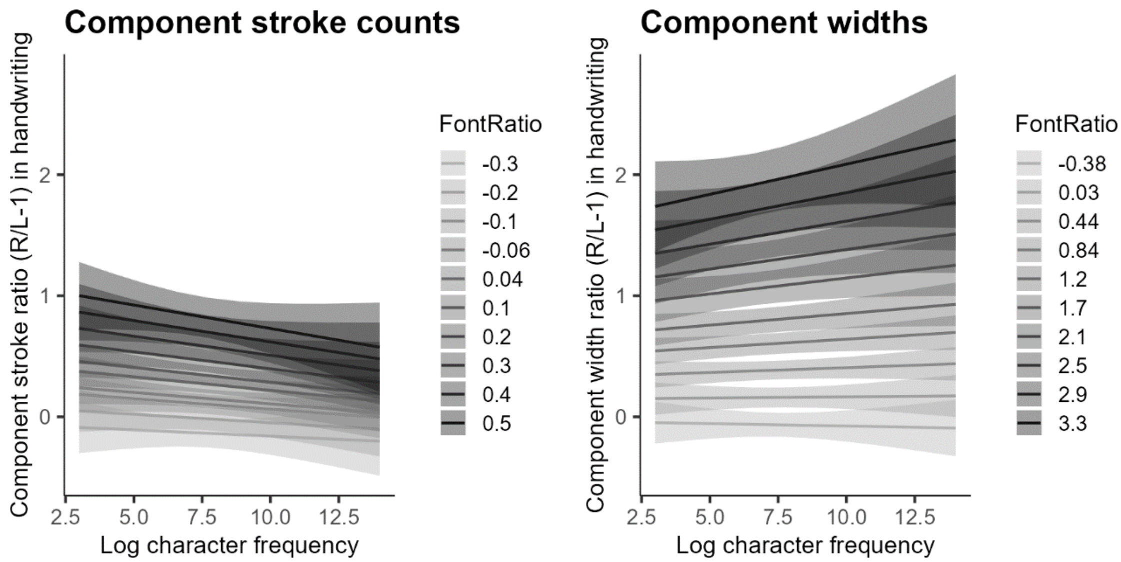

Figure 4. The effect in handwriting of character frequency on the relative number of component strokes with the font stroke count ratio factored out (left panel) and on relative component width with the font width ratio factored out (right panel).

For component widths, illustrated in the right panel of Figure 5, neither the intercept nor character frequency were significant ($p > .1$), but there was a significant positive effect of font component widths ($B = 0.55$, $t(133.10) = 12.87$, $p < .0001$). This width effect was not a mere side-effect of the number of component strokes in either the standard font or handwriting, since it remained in a new model with these factored out as well ($B = 0.54$, $t(137.84) = 12.08$, $p < .0001$). Thus even though writers did not enhance the cross-component asymmetry in widths as they did for the asymmetry in stroke counts, their handwriting still showed this width asymmetry. This either means that this gradient detail is stored in the lexicon (character-specific phonetics), or, more likely, the width asymmetry is implemented productively in the course of handwriting.

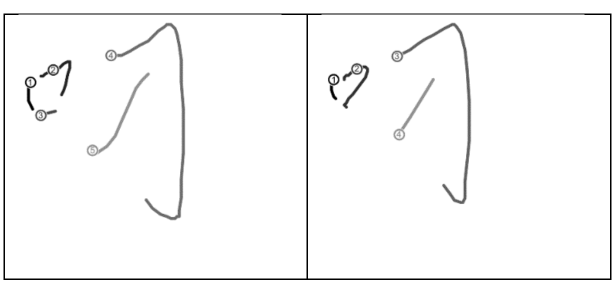

Figure 5. Variable stroke dropping in the relatively low-frequency character (8 tokens per 10 million) 叼 *diāo* ‘hold in mouth’ as produced by different writers. The component 口 is canonically written with three strokes, as in the left panel. Numbers and shading indicate stroke order (dark to light).

The experiment thus confirms that prosodic asymmetry also appears in handwriting, affecting both character phonology (component stroke counts) and character phonetics (component width). Moreover, character phonology but not character phonetics shows frequency effects. Specifically, the effects are negative, which suggest that writers are more likely to apply productive prosodic knowledge when lexical access fails, just as been observed in natural language.

5. Conclusions

Both font developers and handwriters place prominence in a consistent position in Chinese characters, namely on the right component in horizontally arranged characters. Consistent with previous arguments that this asymmetry is analogous to prosody, this prominence pattern interacts with the character analogs of morphology, segmental phonology, phonetics, and lexical frequency in ways similar to what has been observed in natural language. This study thus adds to the growing body of research suggesting that prosodic structure is fundamental to the knowledge and processing of complex linguistic structures regardless of modality.

References

Beckman, Mary E., and Jan Edwards. 1990. Lengthening and shortening and the nature of prosodic constituency. In *Papers in laboratory phonology I: Between the grammar and the physics of speech*, ed. John C. Kingston and Mary E. Beckman, 152–178. Cambridge: Cambridge University Press.

Bishop, Tom, and Richard Cook. 2007. A character description language for CJK. *Multilingual* 18(7): 62–68.

Bybee, Joan. 2002. Word frequency and context of use in the lexical diffusion of phonetically conditioned sound change. *Language Variation and Change* 14(3): 261–290.

Evertz, Martin. 2018. *Visual prosody: The graphematic foot in English and German*. Berlin: Walter de Gruyter.

Fuhrhop, Nanna, and Jörg Peters. 2023. *Einführung in die Phonologie und Graphematik* [Introduction to phonology and graphematics]. 2nd ed. Berlin: Springer.

Gnanadesikan, Amalia E. 2023. Segments and syllables in Hangeul and Thaana: A comparison and optimality theoretic analysis. *Written Language & Literacy* 26(2): 238–265.

Hue, Chih-Wei. 2003. Number of characters a college student knows. *Journal of Chinese Linguistics* 31(2): 300–339.

Meletis, Dimitrios. 2020. *The nature of writing: A theory of grapholinguistics*. Brest: Fluxus Editions.

Myers, James. 2016. Knowing Chinese character grammar. *Cognition* 147: 127–132.

Myers, James. 2019. *The grammar of Chinese characters: Productive knowledge of formal patterns in an orthographic system*. London: Routledge.

Myers, James. 2021. Levels of structure within Chinese character constituents. In *Grapholinguistics in the 21st Century: From graphemes to knowledge. Part II*, ed. Yannis Haralambous, 645–582. Brest: Fluxus Editions.

Myers, James. 2023. Nonstatistical things that linguists can do with R: A review with a case study on Chinese handwriting. In *Proceedings of the 34th North American Conference on Chinese Linguistics* (NACCL-34), ed. Chien-Jer Charles Lin, Alex Cherici, and Bihua Chen, 32pp. Bloomington: Department of East Asian Languages and Cultures, Indiana University Bloomington. https://sites.google.com/view/naccl-34/proceedings

Myers, James. 2024. Interactions among patterns in Chinese character form. In *Looking at language from all sides*, ed. James Myers and James H-Y. Tai, 129–152. Taipei: Crane Publishing.

Peirce, Jonathan, Rebecca Hirst, and Michael MacAskill. 2022. *Building experiments in PsychoPy*. 2nd ed. London: Sage.

Peng, Xuanwei. 2017. Stroke systems in Chinese characters: A systemic functional perspective on simplified regular script. *Semiotica* 2017(218): 1–19.

Pierrehumbert, Janet. 2002. Word-specific phonetics. In *Laboratory Phonology 7*, ed. Carlos Gussenhoven and Natasha Warner, 101–139. Berlin: Mouton de Gruyter.

Qiu, Xigui. 2000. *Chinese writing*, compiled and translated from Chinese by Gilbert Louis Mattos and Jerry Norman. Berkeley, California: The Society for the Study of Early China and The Institute of East Asian Studies.

R Core Team. 2024. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/

Sandler, Wendy. 1993. A sonority cycle in American Sign Language. *Phonology* 10(2): 243–279.

Tsai, C. H. 2006. *Frequency and stroke counts of Chinese characters*. http://technology.chtsai.org/charfreq/.

Van der Kooij, Els and Onno Crasborn. 2008. Syllables and the word-prosodic system in Sign Language of the Netherlands. *Lingua* 118(9): 1307–1327.

Wang, Jason Chia-Sheng. 1983. Toward a generative grammar of Chinese character structure and stroke order. Doctoral dissertation, University of Wisconsin-Madison.

Watt, William C. 1975. What is the proper characterization of the alphabet? I: Desiderata. *Visible Language* 9(4): 293–327.