Concordia Working Papers in Applied Linguistics

Proceedings of the International Symposium on the Acquisition of Second Language Speech
Concordia Working Papers in Applied Linguistics, 5, 2014

The Acquisition of the Tap/Trill Contrast Within and Across Words in Spanish

Daniel Scarpace
University of Illinois

Abstract

This study describes the acquisition of the two rhotic phonemes in Spanish, the tap [ɾ] and the trill [r], within and across words, by English speakers learning Spanish. These two segments are especially difficult for learners to acquire as the contrast does not exist in English, and the two sounds are difficult to produce. Listeners heard fragments of phrases and had to decide whether or not the word they saw on the screen could correspond to the portion of the phrase they heard. This tests the learners' lexical encoding of such words, not only where the phoneme to grapheme correspondence is transparent (word-internally), but also where it is not (word-finally and initially). The results show much variation among speakers: not all speakers have acquired this distinction, even at very high proficiency levels.

The Spanish Rhotic System

Spanish has two rhotic phonemes: the alveolar tap, /ɾ/, which "is produced with a single rapid contact of the tip of the tongue against the alveolar region" (Hualde, 2005, p. 181) and the trill, /r/, which is described as having two or more rapid occlusions at the same place of articulation. The two sounds are in contrast in the word-internal position between vowels, leading to minimal pairs such as /ˈpeɾo/ 'but' - /ˈpero/ 'dog' and /ˈkaɾo/ 'expensive' - /ˈkaro/ 'car'. However, in all other positions within a word, there is no possible contrast, but there are several distributional tendencies (Hualde 2005): /r/ is found word-initially ([ˈroka] 'rock') or after a heterosyllabic consonant ([al.redeˈdor] 'around'), whereas /ɾ/ is found after a tautosyllabic consonant ([ˈpɾimo) 'cousin'). As a coda consonant, the rhotic may be either a tap or a trill: [paɾte] ~ [parte] 'part'. Word-finally, [ɾ] is found before a vowel (*se*[ɾ] *amigos* 'to be friends'), whereas either may be found before a consonant or a pause (*se*[r~ɾ] *poeta* 'to be a poet'; *vamos a ve*[r~ɾ] 'let's see'). The word-final trill pronunciation is especially available in emphatic speech. Thus, at the word boundary and in between vowels, there is a possible contrast between the tap and trill based on the word affiliation of the rhotic in question. If the rhotic is word-final, it is a tap: *ve*[ɾ] *ocas* 'to see geese'; if it is word-initial, it is a trill: *ve* [r]*ocas* 's/he sees rocks'. The change in segment changes the meanings of both of the words; thus this contrast is an important part of the phonological system of Spanish. Importantly, in Spanish orthography, the tap/trill contrast is only represented word-internally between vowels: /r/ is written <rr> as in *perro*, and /ɾ/ is written <r> as in *pero*. Elsewhere, only <r> is used, even in the *ve rocas*/*ver ocas* pair that does contrast the tap and trill.

Latin had no trill /r/; all written rhotics represented the tap. The appearance of /r/ came about in two ways. First, Latin had a singleton/geminate contrast for word-internal consonants in between vowels. In the development of Latin into Spanish, geminate consonants simplified to singleton consonants (CUPPA > *copa* 'wine glass'), with the exception of the alveolar sonorants /nn/, /ll/, /ɾɾ/. /nn/ and /ll/ both palatalized into /ɲ/ and /ʎ/ respectively, while /ɾɾ/ became /r/. Thus, the singleton/geminate contrast in Latin is preserved as the tap/trill contrast in Spanish. Secondly, through fortition, initial /ɾ/ became [r], matching the usual outcome of initial consonants being identical to their word-internal geminate counterparts (e.g. for /t/, TERRA > *tierra*, GUTTA > *gota*). All other positions retained the tap. Many word-final taps arose from the loss of a word-final vowel: MARE > *mare*. In similar situations where the loss of final /-e/ would have resulted in a word-final geminate, the final vowel was not lost (TURREM > *torre*, although *torre* may have resulted from paradigm leveling from the plural TURRES > *torres*, replacing former *tor* (Penny 2002, p. 82). Word-final /ɾ/ in Latin became word-internal through metathesis and retained as [ɾ] (SEMPER > *siempre*). Syllable-final /ɾ/ was maintained as [ɾ] word-internally as well: CORPUS > *cuerpo*. Where did the possibility of syllable-final [r] come from? The pronunciation of a trill may have been maintained when an intertonic vowel was lost following geminate -RR-, as in CARRICARE > *cargar* 'to load' (related to CARRU > *carro* 'cart'), leading to alternate pronunciations of <r> in syllable-final position. However, it is also possible that the free variation between the tap and trill in these positions is due to the lack of possibility of contrast, and the salience of the word-initial position (which only admits /r/).

Problems for Second-Language Learners

English speakers who learn Spanish are faced with a number of problems when acquiring the complete tap/trill contrast in Spanish. Previous research has mostly focused on the production of the tap and trill, but if perception precedes production, it is just as important to focus on Spanish learners' abilities to perceive the contrast. At the most basic level, learners must be able to hear the contrast between the trill, a segment that does not exist in English, and the tap, a segment that exists in English only as an allophone of /t/ and /d/. Rose (2010) performed a comprehensive set of experiments focusing on English learners of Spanish and their ability to discriminate various sounds of Spanish, including the tap and trill. To capture what might be the initial state in acquisition, Rose asked native speakers of English with little or no experience with Spanish to identify the sound present in words such as *caro/carro* as categories in English in a cross-language identification task. The Perceptual Assimilation Model (PAM) proposed by Best (1995) suggests that the success of discriminating non-native contrasts is related to the relationship of the sounds in the second language to existing categories in the first language. The trill /r/ was identified as /ɹ/ 96.2% of the time with a 2.3 goodness rating, whereas the tap /ɾ/ was identified as /ɹ/ 57.7% of the time with a 2.3 goodness rating and 30.4% as /d/ with a 2.8 goodness rating. This uncategorized vs. categorized contrast suggests that discrimination should be very good according to Best (1995), but poor according to Guion et al. (2000), due to the closeness in phonological space of the two rhotics. Rose then tested 75 native speakers of English (15 speakers in 5 different proficiency groups) in an AXB task for the minimal pairs *caro/carro* and *quería/querría* along with several other filler items in three blocks. Discrimination for the tap/trill contrast was already high for English speakers who had never studied Spanish (80.2%) and there was no statistical difference among the four groups who had studied Spanish (mean accuracy scores from 86.7% -

94.4%). Although these data suggest that English learners of Spanish can discriminate the two sounds at a nearly native-like level, the low number of stimuli used that were repeated three times might have made the task easier.

Furthermore, sound discrimination is not enough to acquire a contrast: learners must be able to correctly associate the sound that they heard with the correct word; this reflects correct lexical encoding of the contrast. When beginning to learn Spanish, they may or may not be assisted by Spanish orthography, which presents the contrast word internally between vowels as a contrast between the graphemes <r> and <rr>. On one hand, this may help learners tell that *pero* and *perro* are distinct words, but English orthography exhibits a poor grapheme-to-phoneme mapping, and English <r> and <rr> both map to /ɹ/, e.g. *arid* and *arrow*. As mentioned above, Spanish orthographic rules do not clue learners into the pronunciation of <r> in other positions, such as that the <r> in *rana* 'frog' must be pronounced as a trill. Initially, English speakers beginning to learn Spanish may be able to keep the two sounds distinct by first encoding /ɾ/ in words such as *caro* as /d/, as the tap is an allophone of /d/ in English. Thus *caro* and *carro* would be minimal pairs (/kado/ vs. /karo/), but not in the same way Spanish speakers encode them (/kaɾo/ vs. /karo/). However, not all taps in Spanish can be interpreted as allophones of English /d/ because they can occur in stressed as well as unstressed syllables, and at some point must encode them separately from /d/. Perhaps the greatest challenge presented to second-language speakers is the acquisition of allophony based on word context: in *ser o no ser* 'to be or not to be', the first <r> is necessarily a tap, but the second <r> is likely to be a trill; thus, the same word, *ser*, could appear in the same phrase with two different pronunciations. It follows that a proficient listener must pay attention to the /r/ - /ɾ/ contrast between vowels but ignore it entirely elsewhere.

THE PRESENT STUDY

The research study described in this paper seeks to add to the current research on second-language acquisition of segmental contrasts by comparing a contrast that occurs transparently at the word-internal intervocalic position, and occurs very rarely in between words, but does not occur in other word positions. The tap/trill contrast in Spanish is not only an important phonemic contrast, but it also serves as a significant cue to word segmentation because the trill is the only rhotic to appear in word-initial position.

(1) Are second-language learners of Spanish more sensitive to the tap/trill contrast in the word-internal condition than the word-boundary condition?
(2) Do second-language learners become more native-like with greater proficiency?

To investigate these questions, a word-monitoring task was devised that could test the sensitivity of learners to the phonological system of the rhotics in Spanish. Because words were presented on the screen during the task, the question of orthography comes into play. However, because many students learn Spanish in a classroom context, it is not possible (and perhaps not advisable) to divorce the orthographic representations of the rhotics from the sounds themselves. The participants were asked to monitor for the first few syllables of a specific word represented on the screen; however, the relationship between what was presented on the screen and what was to be heard differs for each condition. For the word-internal condition, the task is more straightforward. If the word contained an intervocalic <r>, as in *barones*, the participant was to listen for an intervocalic tap [ɾ]. If the word contained an intervocalic <rr>, as in *barrotes*, the participant was to listen for an intervocalic trill [r]. The task was quite different for the word-boundary condition because the acoustic information preceding the word presented on the screen was important. If the word on the screen began with an <r>, as in *relojes*, then the participant was to listen for an intervocalic trill [r]. If the word on the screen did not begin with an <r> at all, but rather a vowel, as in *elotes*, then the participant was to listen for the first few syllables in that particular word—in each case, it would be preceded by a tap [ɾ]. This should be the most difficult condition based on orthography because the participant must reject a token such as *pedí*[r]*elo-* when monitoring for *elo(tes)* because a V[r]V sequence must be segmented as V#[r]V and never as V[r]#V. Thus, it is hypothesized that for learners of Spanish, the word-internal contrast should be the easiest, because it is both robust in the language and transparent in orthography, with no obvious differences between the conditions in which they are to monitor for a tap or a trill. The word-boundary conditions should be harder, because the orthographic representation is more opaque (in fact, the <r> represents a trill, whereas in the word-internal condition, the <r> represents a tap—within the context of the experiment, the participants must pay attention to word position carefully). Within the word-boundary conditions, it should be easier to assign <r> as a word-initial trill due to the overall salience of the trill as being associated with an <r>, and relatively difficult to reject the intervocalic trill at the word-boundary when monitoring for a vowel-initial word. All conditions should improve with improved proficiency.

Methodology

Participants. Forty total participants spent roughly thirty minutes to complete all of the experimental tasks (fifteen minutes for native Spanish speakers). A control group was formed of twenty native Spanish speakers ('NS') from various dialects. Seven of the twenty native speakers were recruited and tested at the Colegio de México in Mexico City. All native speakers had learned Spanish from birth, although three speakers were bilingual: two speakers also spoke Catalan, and one had moved to Galicia at a young age. The twenty non-native speakers ('L2') were recruited both from the graduate and undergraduate student population at the University of Illinois, Urbana-Champaign. The undergraduate participants were enrolled in upper division coursework in Spanish at the time of testing. The English native speakers were assessed in their Spanish proficiency through a written test, consisting of the vocabulary and cloze parts of the DELE (Diploma de Español como Lengua Extranjera) test, widely used in other studies of L2 learners of Spanish (Montrul, 2004; Montrul, 2005). The results are reported in Figure 1.

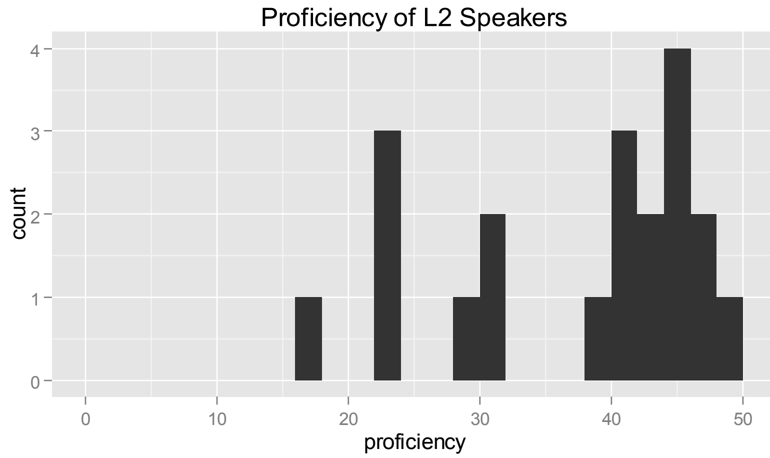

Figure 1. Histogram showing the performance of the 20 English learners of Spanish on the DELE test used to measure proficiency learners with scores of 40-50 are considered to be advanced; 30-39, intermediate; 0-29, beginner. Five out of the twenty subjects scored in the beginner range on the proficiency test; seven scored in the intermediate range and the remaining twelve scored in the advanced category.

Materials. A male Native Spanish speaker from Mexico recorded the stimuli for the experiment. The stimuli were verb-noun or verb-adjective pairs where a V[r~ɾ]V sequence straddled the word boundary. The first word in these phrases were common verbs, either in the infinitive or conjugated in the imperfect (-ir verbs) or present tense (-ar/-er) because the lack of the final <r> in these verbs results in a conjugated form, e.g. *ver/ve*. The second word in each pair had the first two syllables in common, save for the target rhotic. For example, *pedi*[ɾ] *elotes* 'to ask for corn' (word final <r>: tap) vs. *pedí* [r]*elojes* 'I asked for clocks' (word-initial <r>: trill). Similar pairs to test the word-internal contrast were also constructed; for example, *veo ba*[ɾ]*ones* 'I see barons' vs. *veo ba*[r]*otes* 'I see bars'. The recordings were segmented in Praat and truncated at the disambiguation point, e.g. *pedi*[ɾ/r]*elo*-, such that the identity of the rhotic was the only cue to the identity of the second, truncated word. Sixteen pairs of each type were created and organized into four word lists, along with thirty-two filler items. The pairs were divided among each of the word lists such that any one subject did not see or hear both members of any particular pair during the experiment. For example, subject A might hear *pedir elo*- and see *elotes*, whereas subject B would hear *pedí relo*- and also see *elotes*. Neither subjects A nor B would see the target word *relojes*.

Procedures. All participants completed a language background form prior to engaging in the experimental task, which was a word-monitoring experiment. The tasks were preceded by a short training session consisting of ten items (chosen from the filler items) and gave the participant feedback after each response. The experiments were designed and ran entirely in Matlab 2011a, using Psychophysics Toolbox extensions (Kleiner, Brainard, & Pelli, 2007); all subjects used the same laptop computer when engaging in the experiment and listened to all of the stimuli at a constant volume while wearing Shure SRH440 Professional Studio Headphones. Subjects spent roughly ten minutes on the experiment. Finally, all non-native speakers then took a written proficiency test (as described above), testing their knowledge of Spanish vocabulary and grammar.

The experiment was similar to the word-fragment priming task of Soto-Faraco et al (2001) and went as follows: the target word of the critical phrase was presented visually on a computer screen in front of the participant; for example, *elotes* or *relojes*. The truncated auditory stimulus (the two-word phrase) was then presented to the participant, and the participant was asked to press the space bar if the second word in the phrase could possibly be the one visually presented on the screen. If not, the participant would do nothing, and the next trial would begin after two seconds. For example, the participant might first see the word *elotes* on the screen, and then would hear *pedi*[ɾ] *elo-*. In this "match" condition, a press of the spacebar was marked as a correct response, because in the two-word condition, the tap must be a word-final segment, and the second word must begin with a vowel. The task was the same for the word-internal condition: a subject might have heard *veo ba*[ɾ]*o-* and decide if *barones* could have been the second word. Orthography is the main difference between the two conditions: in the word-internal condition, the phoneme-to-grapheme representation is transparent: a trill is written as <rr> and a tap is written as <r>; both could appear in the visually-presented word in the experiment. In the word-boundary condition, a trill would be represented by an initial <r>, but no written <r> is presented on the screen for a tap because that segment belongs instead to the preceding word (not presented on the screen); instead, a vowel-initial word could correspond to a two-word phrase including a tap at the word boundary. Participants heard 32 target items (8 word-internal taps, 8 word-internal trills, 8 word-boundary taps, 8 word-boundary trills) as well as 32 filler items (exhibiting other segmental contrasts in Spanish such as voicing of stops) in a randomized order.

Results

Table 1. Accuracy Results for the Word-Internal Condition

see:	<r>	<r>	<rr>	<rr>
	barones	*barones*	*barrotes*	*barrotes*
hear:	[ɾ]	[r]	[r]	[ɾ]
	veo ba[ɾ]*o-*	*veo ba*[r]*o-*	*veo ba*[r]*o-*	*veo ba*[ɾ]*o-*
	MATCH	NO MATCH	MATCH	NO MATCH
NS %	95.00%	86.58%	100.00%	86.58%
L2 %	91.25%	40.96%	98.68%	37.50%

Table 2. Accuracy Results for the Word-Boundary Condition

see:	<#V>	<#V>	<r>	<r>
	elotes	*elotes*	*relojes*	*relojes*
hear:	[ɾ]	[r]	[r]	[ɾ]
	pedi[ɾ]*elo-*	*pedí*[r]*elo-*	*pedí*[r]*elo-*	*pedi*[ɾ]*elo-*
	MATCH	NO MATCH	MATCH	NO MATCH
NS %	87.50%	72.50%	94.67%	84.70%
L2 %	85.00%	51.25%	92.21%	43.37%

It is important to keep in mind what each condition represents when analyzing the data: Tables 1 and 2 provide examples of what each experimental condition was like for the participant, along with the mean accuracy data for each of the two speaker groups. For example, in Table 1, which describes the word-internal condition, the subject might have seen a word with an orthographic <rr> represented on the screen, such as *barrotes*. The participant was told to press the spacebar if the second word in the two-word phrase could have been *barrotes*. The participant then heard *veo barro-*, which included a trill. Because the target word was the same as the word in the auditory stimulus (although truncated), this condition was considered to be a "match", and a press of the spacebar was considered to be an accurate response. In such trials, native speakers had a 100% hit rate, and second-language speakers had a 98.68% hit rate. The following column describes a nearly identical condition, where the participant was told to monitor for a word with orthographic <rr> but this time heard *veo ba*[ɾ]*o-*, which included a tap instead of a trill. In this "no match" condition, a press of the spacebar was considered to be an inaccurate response. Native speakers correctly rejected these auditory stimuli 86.58% of the time, but second-language speakers did so only 37.50% of the time.

The accuracy results are summarized in the following graph. Native Speakers outperformed the L2 speakers in all conditions (overall, L2: 67% accurate, NS: 89%). However, when the accuracy data are separated by the match and non-match conditions (match: the spacebar press is a correct response), the two groups perform quite differently. The two groups have a higher rate of accuracy in the match condition than the non-match condition, suggesting that there is a bias towards pressing the spacebar. This difference between the two conditions is heightened in the L2 group: the L2 group is 91.7% accurate in the match condition, but 43.3% accurate in the non-match condition, whereas the NS group responded 94.2% and 82.7% respectively.

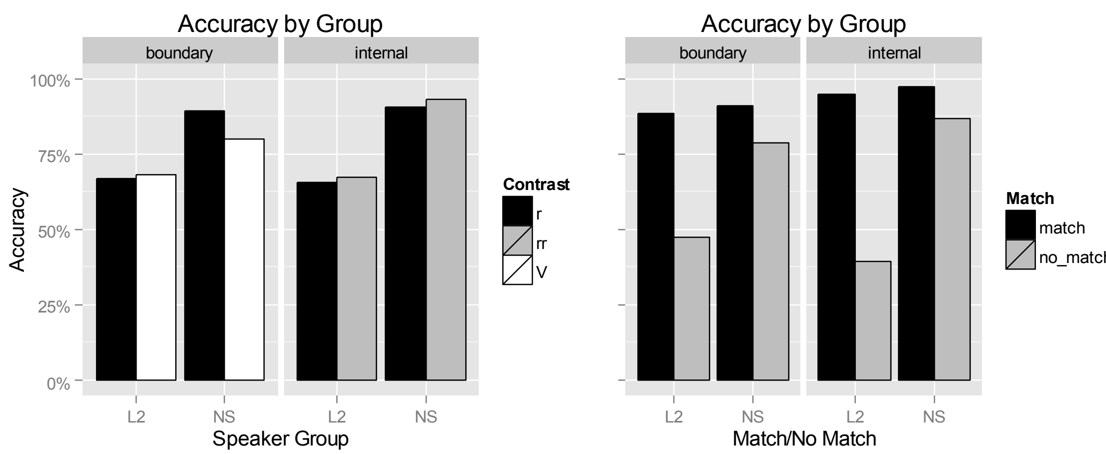

Figure 1. Bar graphs showing mean accuracy data per group in both the word-boundary and word-internal conditions when (a) the data is divided by which segment the participant monitored for; e.g. ‘r’ for word-boundary, the subject saw *relojes* on the screen and (b) the data is divided by whether or not there was a match between the presented segment and the auditory stimulus, e.g. the subject saw *relojes* on the screen and heard a word-initial trill, *pedí*[r]*elo*-.

A d-prime analysis calculating biases revealed differences between the two speaker groups, see the following figure. Within the L2 speakers, proficiency was highly correlated with d-prime ($R^2 = 0.44$).

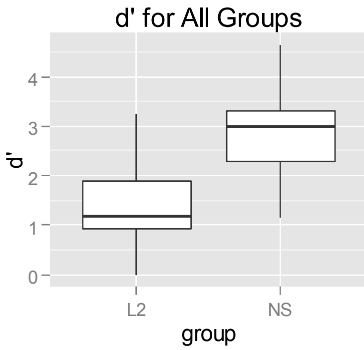

Figure 2. D-prime analysis comparing the two groups.

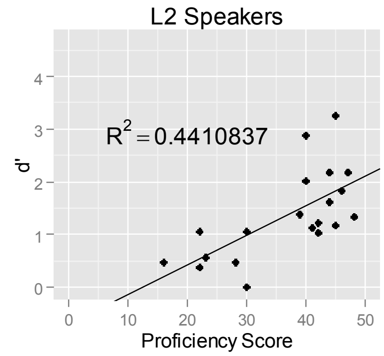

Figure 3. Correlation between proficiency and d-prime for L2 speakers.

Because only small differences were found between the groups and the conditions were found in the match condition, the rest of the data considered here will only focus on the non-match condition, thus the accuracy data represents percentage of correct rejections. The implications of choosing such an analysis can be found in the Discussion section. Figure 4 details the same information as in Figure 1(a) but only focuses on the non-match condition.

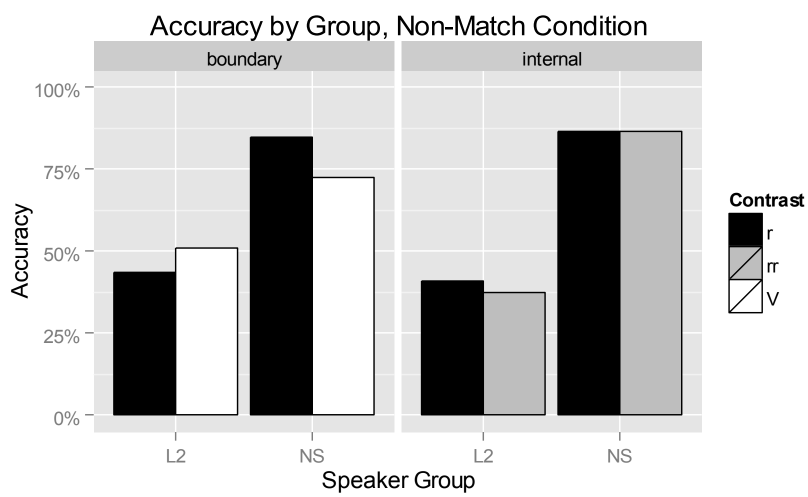

Figure 4. Percentage of Correct Rejections for each group divided by position of the contrast in the non-match condition: either at the word boundary (left) or word –internally (right). Each condition is further divided by the orthographic cue present in the visually-presented stimulus. For example, for <r> in the boundary condition, the word presented might be *relojes*. Because this is the non-match condition, the auditory stimulus heard would have involved a tap: *pedir elo-*. Thus, L2 speakers were below chance at correctly rejecting *relojes* as the second word in *pedir elo-*.

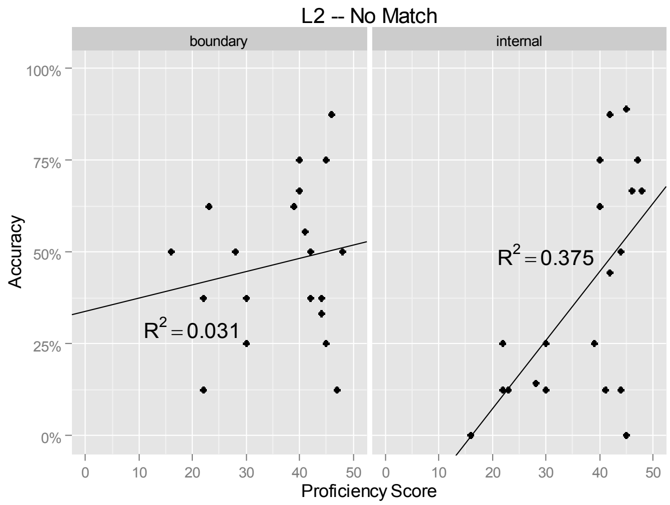

Figure 5. This scatterplot shows the correlation between the percentage of correct rejections and each individual second language speaker's proficiency score for (a) the word-boundary condition and (b) the word-internal condition.

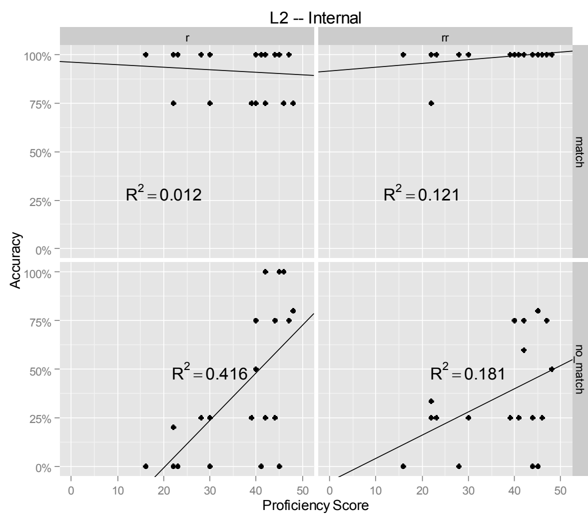

Figure 6. Scatterplots and correlations between proficiency scores and accuracy in each of the four conditions from Table 1, the word-internal condition. The "match" conditions represent hits and the "no match" conditions represent correct rejections. Each point on the plot represents a single L2 participant.

Figures 7 and 8 reveal correlations between proficiency and accuracy in each of the conditions described in Tables 1 and 2. These results are obscured by the mean data presented in the earlier bar graphs: where it is appears that there is no-to-little difference between the overall accuracy (of correct rejections) of the L2 speakers in the word-boundary and word-internal conditions, the R^2 values presented in these scatterplots tell a different story. First, Figure 7, which describes the word-internal condition: in the match conditions, all subjects perform close to ceiling. In the non-match condition, proficiency has a strong correlation with percentage of correct rejections, especially in rejecting [r] as a match for <r>. In Figure 8, which describes the word-boundary condition, many of the participants perform close to ceiling in the match conditions, but there is an exceptional amount of individual variety in the non-match conditions. Proficiency does not seem to help either: R^2 is 0 in the non-match/vowel-initial word condition, as even advanced learners perform equally poorly and native-like.

A series of linear mixed effects models with speaker and item as a random effect were created in order to determine if various fixed effects were significant. With data from both the native speakers and second-language learners combined, the best model associated accuracy with group and match condition ($p < 0.001$); there was a weak association with position of contrast ($p = 0.05$), but there was also a significant interaction between group and position of contrast ($p = 0.012$), and group and match condition ($p < 0.01$). Within the group of second-language speakers, the best model predictor of accuracy was an interaction between position of contrast and match condition ($p = 0.002$) and a three-way interaction between proficiency, position of contrast, and match condition ($p = 0.014$). There was also a weak association with proficiency alone ($p = 0.05$).

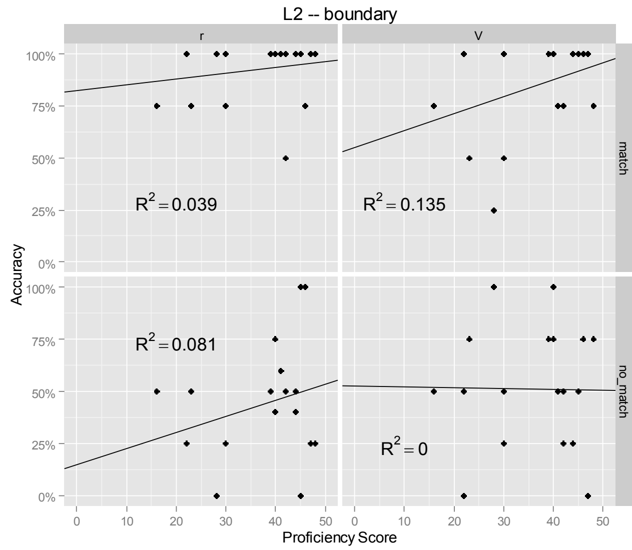

Figure 7. Scatterplots and correlations between proficiency scores and accuracy in each of the four conditions from Table 2, the word-boundary condition. The "match" conditions represent hits and the "no match" conditions represent correct rejections. Each point on the plot represents a single L2 participant.

DISCUSSION

There are a number of surprising findings in this study. First of all, the Second Language speakers perform notably worse in the non-match conditions; that is, when a correct response means rejecting a match between the visually and auditorily presented stimuli. There may be inherent bias in the experimental design to pressing the space bar, thereby over-accepting matches (the opposite response is to not press any button at all), or it may be the case that this high rate of false alarms is indicative of a lack of acquisition of the contrast, even in the basic word-internal context (e.g. *caro* vs. *carro*). This result neither confirms nor denies the results given in Rose 2010: recall that this previous study investigated learners' abilities to simply discriminate between minimal pairs, a much more explicit task. The task presented here requires participants to not only identify which sound is being heard but also to correctly assign it to the orthographic (or lexical) form represented on the screen. That is, even some of the lower proficiency speakers who have nonetheless had at least of year of Spanish study are not able to quickly reject [r] as a pronunciation of <r>. Simply put, both /r/ and /ɾ/ (as well as <r> and <rr>) belong to the same category, and the two sounds are in free variation for these lower proficiency speakers. A production task might bear this finding out; however, due to the difficulty in producing the two rhotics in

Spanish, production of /r/ and /ɾ/ in this context as [ɹ] would not necessarily reflect one underlying category, but it may in fact reinforce it. Face (2006) found that in the word-internal intervocalic position, there was a stage in which learners of Spanish were producing [ɾ] for both <r> and <rr> in a reading task, which supports this result.

Although the tap/trill contrast in Spanish is an important distinction in Spanish, functionally, there are very few minimal pairs. For a beginning learner, the pairs *pero* 'but' / *perro* 'dog' and *caro* 'expensive' / *carro* 'car' are common, but it is likely that at an early stage of lexical encoding, all of these entries include the English /ɹ/ or some non-distinct rhotic. Because of the lack of other common minimal pairs, the contrast may be very slowly acquired, and these early entries may never be repaired. The lack of contrast may be reinforced by the variation between [r] and [ɾ] in syllable final position, which also has a lack of orthographic transparency.

It is therefore of no great surprise that the word-boundary condition shows even greater variability among the higher proficiency speakers. This condition is also more difficult for the native speaker group, but even some of the highest proficiency second-language speakers have a lack of awareness that word-initial <r> is always a trill. The task is harder in this condition because of the necessity of paying attention to the preceding word context, but it might be additionally harder because /r/ presents an exception to the normal rules of resyllabification between words in Spanish. The usual rule is that /V#CV/ and /VC#V/ sequences are homophonous; that is, *la salas* 'you salt it' and *las alas* 'the wings' are identical, especially in fast speech. However, as mentioned above, *ve rocas* and *ver ocas* are never identical: the tap or trill is the cue to the placement of the word boundary. Pairs of phrases that are ambiguous are non-existent in the input of beginning learners, and the distribution of /r/ is never explicitly taught in classes. A follow-up production study would be illustrative to find out what realizations of /r/ second language speakers use in contexts other than the word-internal, intervocalic context.

The amount of individual variation calls into question the usefulness of pooling second-language speakers together as a single group. Although overall scores (such as d-prime) improved with greater proficiency, in the most difficult conditions, proficiency was not correlated with accuracy. The individuals with high proficiency were mostly graduate students in the Department of Spanish, and some may have had explicit pronunciation instruction and may be metalinguistically aware of the tap/trill contrast at the word boundary. Any further research should explicitly ask participants after the experimental tasks have been completed if they are aware of the rule in question.

CONCLUSION

This study has shown that with second language speakers are able to acquire the tap/trill contrast in Spanish with varying levels of results. In the word-internal position (*pero*/*perro*), an increase in proficiency is well correlated with an increase in performance. However, in the word-boundary condition (*ve rocas*/*ve rocas*), only some second language speakers were able to correctly assign the tap or trill to the lexical item to which it belongs. Therefore, while there is an increase in accuracy as the proficiency level of the learner increases, it seems that contrasts at the word boundary are harder to acquire. This difficulty may be a result of the weak functional load that the tap/trill contrast holds (only between vowels, in free variation elsewhere except word-intially), or it may be due to the lack of transparency in the orthography (written as <r> in all cases). The word-boundary condition also has implications for word-segmentation strategies in the second language: there is weak evidence here that second-language learners are able to use the allophonic cue of the tap/trill to correctly place a word boundary. This study also shows that the lexical encoding of these rhotics is fairly weak for beginning learners of Spanish, as they are willing to accept either the tap or the trill as being in any of the presented words. Future research should couple perception tasks with production tasks in order to further understand the state of the acquisition of the rhotic system in second-language learners' Spanish.

ACKNOWLEDGEMENTS

I would like to thank Annie Tremblay and José Ignacio Hualde for their helpful comments and support on this project. I would also like to thank the participants and members of the 4th Annual SLATE Graduate Research Symposium at the University of Illinois, Urbana-Champaign, where an early version of this project was presented.

REFERENCES

Best, C. T. (1995). A direct realist perspective on cross-language speech perception. In W. Strange (Ed.), *Speech perception and linguistic experience: Theoretical and methodological issues in cross-language speech research* (pp.167–200). Timonium, MD: York.

Face, T. L. (2006). Intervocalic rhotic pronunciation by adult learners of Spanish as a second language. In Carol A. Klee & Timothy L. Face (Eds.), *Selected proceedings of the 7th conference on the acquisition of Spanish and Portuguese as first and second languages* (pp. 47–58). Somerville, MA: Cascadilla Proceedings Project.

Guion, S.G., Flege, J. E., Akahane-Yamada, R. & Pruitt, J. C. (2000). An investigation of current models of second language speech perception: The case of Japanese adults' perception of English consonants. *Journal of the Acoustical Society of America*, *107*(5), 2711-2724.

Hualde, J. I. (2005). *The sounds of Spanish.* Cambridge: Cambridge University Press.

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perceptions, 36 (ECVP 2007 Abstract Supplement).*

Montrul, S. (2004). *The acquisition of Spanish.* Amsterdam: John Benjamins.

Montrul, S. (2005). Second language acquisition and first language loss in adult early bilinguals: Exploring some differences and similarities. *Second Language Research, 21*, 199-249.

Soto-Faraco, S., Sebastián-Gallés, N., & Cutler, A. (2001). Segmental and suprasegmental mismatch in lexical Access. *Journal of Memory and Language*, *45*, 412-432.

Rose, M. (2010). Differences in discriminating L2 consonants: A comparison of Spanish taps and Trills. *Selected Proceedings of the 2008 Second Language Research Forum*, ed. Matthew T. Prior et al., 181-196. Somerville, MA: Cascadilla Proceedings Project.

Penny, R. 2002. *A history of the Spanish language*. Cambridge: Cambridge University Press.