

A Computational Model of the Acquisition of German Case

Daniel Freudenthal, Julian M. Pine, Fernand Gobet
Department of Psychological Sciences, University of Liverpool

Abstract

We present a computational model of the acquisition of German case that is evaluated against empirical data obtained from naturalistic speech. The model substitutes nouns into existing contexts, and proceeds through a number of stages that reflect increasing knowledge on the part of a child, both of the determiner-noun sequences that are legal in German, and of the determiner-noun sequences that are appropriate in specific sentential contexts (cases). The model provides a natural account of gender and case errors, the two most common error types produced by children, and shows the highest error rates in dative contexts and lowest error rates in nominative contexts, as is true of children learning German. However, the model's error rates in the early stages are considerably higher than those shown by children, suggesting that children possess a fairly sophisticated representation of how lexical contexts assign case from a relatively early age.

Keywords: German Case, Acquisition, Computer Modelling

Introduction

A major question in the study of language acquisition is how children learn to mark the morphological contrasts that are distinguished in the language they are acquiring. One such contrast is grammatical case. Many languages mark the grammatical role (subject, object) that nouns (and pronouns) play in the sentence. Languages differ in the complexity of their case system. English distinguishes three different cases (nominative, genitive and accusative) and only marks them on pronouns (*he, his, him*). The focus of the current paper is the German case system, a system that is considerably more complex than the English system, and poses a challenge for language-learning children.

German distinguishes 4 cases (nominative, genitive, dative and accusative), which roughly translate to subject, possessive, indirect object and direct object. German case is expressed on determiners and adjectives as well as pronouns. German case is particularly complicated because German nouns are gendered, and the correct form of the determiner depends both on the gender and case of the noun. Table 1 shows the different forms of the definite and indefinite article for the different combinations of gender and case. As can be seen in Table 1, there is a degree of syncretism in the German determiner system with only 6 forms covering the 16 cells of the paradigm for both the definite and indefinite article.

Table 1: Case marking in German

	Nom.	Gen.	Dat.	Acc.
Masc.	der/ein	des/eines	dem/einem	den/einen
Fem.	die/eine	der/einer	der/einer	die/eine
Neut.	das/ein	des/eines	dem/einem	das/ein
Plural	die/--	der/--	den/--	die/--

Early acquisition reports (Mills, 1985; Czepluch, 1996) suggest that the German case system presents a challenge for children who make errors of omission and commission, and that errors are most common on the dative, which is relatively infrequent. However, these findings are based on diary data or very small samples. More recent and detailed data (Szagun, 2004) confirm that German children's acquisition of the determiner system is protracted and error-prone. Szagun analyses spontaneous speech samples from normal-hearing children as well as children with cochlear implants, and determines rates of correct use of nominative, accusative and dative (definite and indefinite) determiners. Szagun reports data for Mean Length of Utterance (MLU) points of 2.8, 3.8 and 4.8. Because of data sparseness, she collapses data across the first two MLU points for the analyses of accusative data, and across all three points for the dative.

The pattern of errors in the normal-hearing group can be summarized as follows: Error rates on the nominative are relatively low, ~20% for the first MLU point, decreasing with increasing MLU. Early errors are largely gender errors – the child produces a determiner that is inappropriate given the noun's gender, though omissions and proto-forms also occur.

Errors on the accusative (which partly derive from a higher MLU point), are largely case errors – the child produces a determiner that is appropriate for the noun's gender, but inappropriate given its case. These errors mostly take the form of nominative for accusative errors. On the indefinite article, this error occurs at roughly 35% for the first two data points, and at 20% for the third data point. Error rates on the definite article are lower, and errors are more varied.

Errors on the dative (reported for the definite article only) occur at rates close to 50%. The most common errors are nominative for dative and accusative for dative errors (~30% combined), though omission and case+gender errors also occur. The fact that these data are collapsed over all three MLU points suggests that error rates would be considerably higher if children produced more nouns in dative contexts during the early stages. They also indicate that children continue to make errors even at a relatively late stage of development. Szagun does not report data on the genitive, which is rare in child language.

Earlier modelling attempts

MacWhinney et al. (1989) develop a connectionist model that learns German gender, number and case. The model was shown a number of phonological, morphological and semantic cues to a noun's gender, as well as sentential cues to its case. The task of the (back-prop) model was to associate these cues on the input layer with the correct form of the (definite) article that was shown on the output layer. The

main finding was that the model was capable of learning this association, and could generalize this knowledge to previously unseen nouns – even when the (38) gender cues were replaced with a phonological representation of the stem, thus providing evidence that German noun gender is learnable from phonological cues alone. MacWhinney et al. also report that, halfway through training, the model showed relatively good mastery of the nominative, showed poor mastery of the genitive, made some over-generalisation errors, and made many errors of omission.

However, while the model produced errors that have been reported in the child literature, it does not appear to provide a good fit to the child data. Omissions (which reflect the fact that none of the output nodes was sufficiently activated to be selected), were by far the most common error type (68% of all errors), and dative errors were actually less frequent than nominative and accusative errors. This pattern is clearly different from that reported by Szagun, who finds the highest error rates on the dative, and reports omission levels well below 10%.¹ Moreover, it is not clear from MacWhinney et al.'s data which forms were substituted for which targets, and hence whether the pattern of overgeneralizations and other kinds of errors matches that reported by Szagun. Finally, while it is clearly encouraging that the model was able to learn both gender and case and generalize them to unseen items, it is unclear why the model produces the errors that it does. The current paper aims to improve on this situation by developing a model of the acquisition of German case that (a) provides a closer fit to the (developmental) data reported by Szagun, and (b) provides a clearer account of the types of errors that it produces.

The current model

The data reported by Szagun suggest that children go through an initial stage where they are unsure of a noun's gender and combine nouns with determiners that are inappropriate given their gender (*gender errors*). This stage is relatively short-lived, and children quickly become more sensitive to the determiners that are paired with specific nouns. However, children may still produce determiners that are inappropriate given the context (case) in which a noun occurs. Case errors tend to take the form of nominative for accusative, and nominative/accusative for dative. Since nominative contexts are most frequent, and dative least frequent, this finding suggests that children's knowledge of case-appropriate determiners develops more quickly for contexts that occur more frequently.

Here, we explore if the pattern of errors reported by Szagun can be understood in terms of a mechanism that substitutes nouns into existing contexts, and becomes increasingly sensitive both to the determiner-noun sequences that are legal in the language as a whole, and the determiner-noun sequences that are appropriate in specific sentential contexts. Our model proceeds through three stages. In a first stage, it substitutes nouns into existing contexts without considering

their gender. Since this stage treats nouns of different genders as equivalent, it is likely to result in gender errors. However, the rate of gender errors reported by Szagun is not particularly high, so the model's fit to these data may provide some insight into children's knowledge of noun gender.

In a second stage, we use input analyses to constrain the determiners that are paired with novel nouns. We analyse German Child-Directed Speech and note the most common determiner that precedes specific nouns and substitute the noun with its most common determiner into determiner+noun contexts. Since the most common determiner is very likely to be gender-appropriate, gender errors are not expected in this stage. However, since the most common determiner for a noun may be inappropriate for the noun's case, case errors are expected. A process of 'defaulting' to the most frequent form has been reported in several domains and languages (Laalo, 2003; Aguado-Orea, 2004; Räsänen et al., 2014; Freudenthal et al., 2015). The main question of interest here is whether a process of defaulting to the most common determiner for a given noun results in the basic pattern of (over-generalization) errors reported by Szagun.

In stage 2 it is assumed that children always select the most common determiner for a given noun. In stage 3 we investigate how increasing knowledge of the lexical contexts that assign a specific case affects children's ability to select a case-specific determiner (and hence reduce the number of defaulting errors from stage 2). This is done by coding nouns in the input for their grammatical role (case), on the basis of their position relative to an (increasingly large) set of (main) verbs, pronouns and prepositions. We then assess which form of the determiner is most common for a given noun both overall, and in specific case contexts. The choice of determiner for a noun in a given case context is determined by the number of times that noun has been seen in that context. For nouns that have not been seen frequently (i.e. less than a threshold value) in the target context, the mechanism selects the most common determiner (across all contexts). For nouns that have been encountered in the target context more frequently the most common determiner for the given context is inserted instead. We investigate the model's performance using three sets of marker words of increasing size that reflect children's increasing knowledge of case contexts. The main questions of interest here are how this changing knowledge affects the type and rate of case errors (and their fit to the child data), and how appropriately the model is able to mark case across the system. We first provide a brief description of German case, before describing our input analysis, and the simulations.

German Case

German distinguishes four cases: nominative, genitive, dative and accusative. Nominative case is assigned to the subject of the sentence. It is also assigned to nouns in isolation (citation forms) and to subject complements expressed using linking verbs such as the copula. Thus, the noun *Hund* in the

¹ Szagun only counts omission errors from MLU 2.5 onwards.

utterance *Das ist ein Hund-NOM* (*That is a dog-NOM*) has nominative case. Genitive case which expresses possession (e.g. the book of the Man-Gen), is relatively rare, and is not considered here. Dative case is assigned to indirect objects and follows certain dative verbs and prepositions (e.g. *schicken* (send), or *mit* (with)). Finally, accusative case is assigned to *direct objects*, but is also assigned by certain prepositions (e.g. *für* (for)). A complicating factor is that some prepositions assign either dative or accusative case depending (roughly) on whether movement is involved. Thus, the preposition *in* assigns accusative case in a phrase like ‘he jumps in(to) the water’, but dative case in a phrase like ‘he swims in the water’. Overall, however, nouns that follow a preposition are more likely to be dative than accusative.

Input Analysis

The input analysis was carried out on the Child-Directed Speech addressed to 4 children (Corinne, Cosima, Pauline and Sebastian) from the Rigol corpus, available from the CHILDES data base (MacWhinney, 2000). The corpus consists of a range of short recordings of the children between the ages of 1 and 4 years and contains approximately 150,000 adult utterances.

The main aim of the input analysis was to determine what words precede specific nouns both overall, and in contexts that assign nominative, accusative and dative case. The rationale behind our modelling approach is that children are likely to provide the correct determiner for a noun in a given case context, if they have frequently encountered that noun in that case context. Errors are more likely for nouns that have been encountered in a given context less frequently. Moreover, these errors are likely to take the form of the insertion of a determiner that has frequently been encountered with the noun across all contexts. That is, children are likely to default to the most common determiner for a given noun if the child has limited experience with the noun in its target context, and they are more likely to produce the correct determiner if they have encountered the noun in its target context more often.

Needless to say, the number of times a noun has been encountered in a given case context is not only a function of the input that a child hears, but also depends on the child’s ability to correctly identify a noun’s case. For this reason, we coded the thematic role of nouns in three ways that are meant to reflect children’s increasing understanding of which lexical contexts assign which case. The schemes follow the same rules and assign thematic roles relative to a set of marker words consisting of verbs, pronouns and prepositions. The *small marker set* included the 50 most frequent verbs, the 10 most frequent pronouns, and the 10 most frequent prepositions. The *intermediate set* increased the number of verbs to 200 and included all pronouns and prepositions. The

large set further increased the number of verbs to 2000. The three different coding schemes thus reflect an increasing ability on the part of the child to correctly identify case contexts, and hence to determine the appropriate determiner for a noun in a given context.

Coding for thematic roles was done by first dividing the Child-Directed Speech in the Rigol corpus into phrases by breaking utterances up at commas. We then searched the resulting phrases for nouns², and compiled a list of words that preceded each noun (i.e. on a noun-by-noun basis). Coding for thematic role was done by distinguishing between nouns that were likely to fill the thematic role of subject, object or that were post-prepositional³:

- Nouns were marked as post-prepositional if they followed a preposition-determiner sequence.
- Nouns in declaratives were marked as accusative if they followed a verb (excluding linking verbs such as the copula) and/or preverbal nouns in an utterance that contained a pronoun⁴. Post-verbal nouns in questions were marked as nominative if the utterance contained no pronoun, to account for verb inversion in German question formation
- All other nouns were marked as nominative. This included nouns in isolation/citation form, as well as nouns in copula constructions, but also (for the small and intermediate marker set) nouns following verbs or prepositions that were not included in the set of high frequency words.

For each noun, this procedure resulted in four lists of words that precede it: one each for the three separate cases as well as the combined list that consists of the words preceding a noun regardless of case. Note that the coding for case context (especially nominative case) is necessarily noisy. Thus, even when coding using the large set, nouns may occur in contexts that appear nominative but are paired with non-nominative determiners. Examples of this include isolated determiner noun sequences that are elliptical answers to questions. Likewise, object nouns occur preverbally in German modal/infinitive constructions (I want a cookie-ACC eat), which may have their pronoun and modal verb elliptically omitted. However, since the main question of interest here is what is the most common form of the determiner for a given noun-case combination, such instances are unlikely to give rise to high error rates. Results of the input analysis (collapsed over all nouns) are shown in Table 2.

Table 2 shows that, as expected, nominative case is most frequent, followed by accusative and then ‘prepositional’ case. It is also apparent that the large marker set classifies fewer nouns as nominative than the small marker set. This is not surprising given the nature of the coding scheme, but it is consistent with nominative case becoming less of a default as children’s knowledge of case contexts increases.

dative, and that the distinction between accusative and dative prepositions is a subtle one that is likely to be acquired late.

⁴ To account for object topicalization

² Only phrases shorter than 7 words were analysed.

³ As was argued earlier, post-prepositional nouns can be either dative or accusative. Here we assign them to a single class on the grounds that, overall, post-prepositional nouns are likely to be

Table 2: Number of nominative, dative and accusative contexts for the small, intermediate and large marker sets.

	Nom.	Acc.	Dat.
Small	27,407	6,235	5,620
Interm.	25,713	6,979	6,570
Large	25,149	7,543	6,570

The simulations

We evaluated our model by investigating the errors it produced in 3 different stages. In a first stage, we slotted nouns into existing sentential contexts and noted the model’s propensity to substitute nouns of differing genders (and hence to produce gender errors). In a second stage, we paired nouns with the most common preceding word, and determined the pattern of defaulting errors. In a third stage, we used the results from our input analysis to pair nouns with the most common word given a specific case. This reflects a stage in which children start to use their developing knowledge of case contexts to select a case-appropriate determiner.

Stage 1: Case errors in noun substitutions

The first stage was implemented by selecting from the input 1,000 utterances that were between 2 and 5 words long and contained a determiner followed by a singular noun. Next, we took the 100 most common singular nouns from the input analysis and substituted these for the nouns in the target utterances. Since nouns in German are gendered, this procedure has the potential to substitute nouns for nouns of a different gender, and hence result in determiner-noun sequence that are illegal, or inappropriate given the gender of the substituted noun. As can be seen in Table 1, this will be the case for the majority of substitutions involving nouns of a different gender. Since the three genders have no overlap in nominative and accusative determiners, the cross-gender substitution of a noun in a context marked for accusative or nominative case is very likely to result in a determiner-noun sequence that is illegal (i.e. a gender error). The only exception is the substitution of a Feminine noun for a nominative Masculine noun. The nominative Masculine determiner (*der*) is appropriate as a dative feminine article. However, since the context into which the feminine noun is substituted is a nominative one, the appropriate determiner is *die*. Likewise, substitutions in dative contexts give rise to errors if a feminine noun is involved, but since masculine and neuter nouns share the determiner *dem*, these can be substituted without error.

Analysis of the noun substitutions reveals that roughly two-thirds involve nouns of different gender, and hence that the potential for gender errors is very high – considerably higher than the maximum of 15% reported by Szagun (2004). The results of this analysis thus suggest that children’s production of determiner-noun sequences is either more conservative, or considerably more sophisticated than simple substitution of

nouns into existing contexts, at least in the MLU range studied by Szagun.

Stage 2: Defaulting to the most common determiner.

We investigated the latter possibility by substituting not just the noun, but a two-word sequence consisting of the noun as well as the most common word⁵ that preceded it in the input. Since the most common word preceding a noun is very likely to be appropriate considering the gender of the noun, such an approach is unlikely to result in gender errors. However, since the most common word preceding a noun is likely to reflect the most common role of a given noun in the input, rather than the context to which it is inserted, substitution of determiner noun sequences is likely to give rise to case errors.

The analysis of case errors was restricted to the Masculine nouns in the 100 most frequent nouns. As can be seen in Table 1, Feminine and Neuter nouns take the same determiner in nominative and accusative contexts (*die* and *das*, respectively). Thus, while errors on the dative can be identified in all three genders, only the masculine allows one to identify errors on the nominative and accusative. Table 3 shows the distribution of the case assigned by the most common word preceding the different nouns. The first row only considers determiners, while the second row also considers determiners contracted onto prepositions. Contractions are common in German. Masculine dative *in dem* (in the), is commonly contracted to *im*. However, since the preposition *in* assigns dative case the contracted form can always be considered correct.

Table 3: Case of the most common determiner of masculine nouns.

	Nom	Dat	Acc
Determiners only	.68	.05	.26
Inc. contractions	.59	.18	.23

It is obvious from Table 3 that the most common determiner for most masculine nouns is the nominative. However, around a quarter of nouns occur mostly in accusative case. Dative case is least common. If children produce nouns in novel contexts with the determiner they occur with most often, they are therefore expected to make the fewest errors in nominative contexts, and the most errors in dative contexts, as reported by Szagun. Moreover, most errors would involve the insertion of a nominative determiner in a non-nominative context. A process of defaulting to the most common determiner therefore results in the basic pattern of overgeneralization shown by children acquiring German case. However, error rates on the accusative and dative would approach 75% to 80%, which is considerably higher than the rates reported by Szagun. This suggests that,

⁵ We only analysed instances where the most common word was a definite or indefinite article.

at least in the MLU range studied, children do not simply insert the most common determiner for a given noun, but instead show evidence of some knowledge of case-appropriate determiners for the nouns they produce.

Stage 3: Frequency-based insertion of case-specific determiners

The data in Table 3 reflect a situation where a child produces a noun with the word that precedes it most often, regardless of the context it occurs in. We next investigate the situation where coding for case contexts is taken into account. The coding rules described earlier were applied using a small, intermediate, and large set of marker words (verbs, pronouns and prepositions) for coding case contexts. Nouns were coded as occurring in post-prepositional (dative), nominative and accusative context. Next, nouns were probed for the most common preceding word in these contexts. Where a noun had occurred in the relevant context a minimum of 10 times, we inserted the most common determiner for that context. Where a noun had occurred fewer than 10 times, we inserted the most common determiner overall (i.e. from Table 3). The rationale behind this is that, as children become more aware of which lexical contexts assign which case, they are more likely to apply case correctly. Tables 4 and 5 give the results (response distribution in columns) for small and intermediate marker sets. Rows marked ‘Err < cutoff’ represent the proportion of (all) errors in a column where a noun has not occurred in the target context sufficiently often, and the most common determiner across all contexts is inserted.

Table 4: Response distribution for small marker set

	Targets		
	Nom	Dat	Acc
Nom.	.70	.29	.35
Dat.	.00	.52	.06
Acc.	.30	.19	.59
Err. < cutoff	.17	.80	.71

Table 5: Response distribution for intermediate marker set

	Nom	Dat	Acc
Nom.	.85	.24	.31
Dat	.00	.52	.06
Acc.	.15	.24	.62
Err. < cutoff	.32	.70	.67

Tables 4 and 5 show that error rates are lowest for nominative case, and highest for dative case, as reported by Szagun. Errors on nominative case mostly occur over the threshold value, while errors for dative and accusative case mostly occur under the threshold value. This reflects the fact that the coding scheme (correctly) identifies many contexts as nominative, and most nouns are thus seen in nominative contexts more than 10 times. The main difference between coding using the small and intermediate marker set is that accuracy for the nominative is higher in intermediate coding. This reflects the fact that the larger set of marker words in the intermediate set marks fewer contexts as nominative (see

Table 2), but marks these with higher accuracy, and hence is less likely to insert an incorrect determiner.

The large marker set (Table 6), which mainly uses a larger verb category, identifies more contexts as accusative, leading to a reduction in error in both the nominative and accusative, but not the dative. This is not surprising since the number of (frequent) prepositions is limited and the stricter coding scheme does not identify more contexts as dative/post-prepositional, and thus suggests that many nouns do not occur in dative contexts very often. We investigated this by lowering the threshold to 0, meaning we always inserted the most common determiner for a given context. This lowered error rates to .06 for the nominative, .13 for the accusative, and .20 for the dative. These remaining errors reflect noise in the coding scheme, or in the case of dative/post-prepositional errors, reflect the fact that, while most prepositions assign dative case, some actually assign accusative case.

Taken together, the results of analysis 3 show that the pattern of errors reported by Szagun is consistent with a mechanism that inserts the most common determiner for nouns that have been encountered in a given context infrequently, and a case-appropriate determiner in contexts that have been seen frequently. However, they also suggest that error rates on the small marker set are too high, and thus that children have relatively sophisticated knowledge of the contexts that assign (nominative and accusative) case.

Table 6: Response distribution for large marker set

	Nom	Dat	Acc
Nom.	.89	.24	.20
Dat	.00	.52	.07
Acc.	.11	.24	.73
Err. < cutoff	.50	.70	.50

Conclusions

Our model provides a clear account of the types of errors produced by children acquiring German case. Gender errors result from cross-gender substitutions into existing contexts, while case errors result from a process of defaulting to the most common determiner for a given noun. Analysis 2 showed that such a defaulting mechanism supports the basic pattern of errors shown by children acquiring German case. However, both mechanisms produce errors at rates that are considerably higher than those shown by children, and thus suggest that children’s knowledge is more sophisticated.

A third stage employed information regarding lexical case contexts obtained using a small, intermediate or large set of marker words, and inserted a determiner appropriate for nominative, accusative or dative contexts when a noun had been seen in this context sufficiently often. This reduced error rates, in particular for nominative and accusative case, and brought the error rates more in line with those reported by Szagun. However, it could be argued that error rates only really match those reported by Szagun for the intermediate (or large) marker set, thus suggesting that children have a fairly sophisticated grasp of which sentential contexts assign which case. It should also be noted, however, that the

(accusative and dative) data reported by Szagun are collapsed across several MLU points, and thus may obscure high error rates at lower MLUs.

The pattern of case errors is the same across all three sets of marker words. Rates of case error are lowest for the nominative. This reflects the fact that nominative case is the most common, and hence that nouns are very likely to be paired with the correct determiner. Errors on the dative and accusative are more common. For the accusative, they are largely restricted to nominative errors. For the dative they are split between nominative and accusative errors (as reported by Szagun). Error rates on the nominative and accusative are reduced for the larger sets of marker words, but remain high for the dative. Thus, the low frequency of nouns in post-prepositional contexts means that error rates remain high, even for the large set of marker words. This matches the finding that children continue to make dative errors at high MLUs. However, the finding that error rates are reduced when the threshold value is lowered does indicate that the (noisy) coding scheme employed here can result in reasonably error-free performance.

While the model provides an account of the basic type and pattern of errors it also has some shortcomings. For one, it does not distinguish between definite and indefinite articles, and hence cannot account for differences in the errors they attract. Thus, Szagun (2004) reports that nominative for accusative errors are particularly frequent for the indefinite article where the nominative (*ein*) is reduced relative to the accusative (*einen*). Szagun only reports dative data for the definite article, but does suggest that accusative for dative errors may reflect the greater phonological similarity between the dative and accusative determiners (*dem, den* – pronounced as English *dame*, and *dane*), relative to the nominative (*der*, English *deer*). Clearly, such phonological considerations are beyond the scope of the current model. However, the results reported here do suggest that a frequency-based account can explain the overall pattern of errors, including accusative for dative errors.

In a similar vein, the current model equates dative and post-prepositional contexts. While it is true that most prepositions assign dative case, some actually assign accusative case or accusative and dative case. The task faced by the child is thus considerably more complex since the child needs to learn which case is assigned by which prepositions. Again, this task is beyond the scope of the current model.

Finally, the model presented here is purely lexical, and only considers around 35 masculine nouns contained in the most frequent 100 nouns. While this partly reflects the amount of input available to our model, children are ultimately able to generalize their knowledge to nouns they have not seen before. It should be noted, however, that while the process of acquiring German case may be protracted, the system is highly regular. That is, a noun's gender (and hence its case-appropriate determiners) can be determined on the basis of a single occurrence in a nominative or accusative context. Such adult-like performance is also beyond the current model.

Taken together, the results reported here suggest that a relatively simple frequentist model that defaults to the most common form of the determiner for a given context, cannot only provide a plausible account of children's errors in the acquisition of case, but could easily be extended to result in error-free performance on a larger set of nouns. However, they also suggest that, despite poor performance in dative contexts, children actually have quite sophisticated knowledge of the contexts that assign nominative and accusative case.

Acknowledgments

Daniel Freudenthal, Julian Pine, and Fernand Gobet are members of the International Centre for Language and Communicative Development (LuCiD) at the University of Liverpool, for which the support of the Economic and Social Research Council [ES/L008955/1] is gratefully acknowledged.

References

- Aguado-Orea, J. (2004). *The acquisition of morpho-syntax in Spanish: Implications for current theories of development*. Unpublished doctoral dissertation. University of Nottingham, UK.
- Czepluch, H. (1996). Case morphology and case system in L1 acquisition: A pilot study of a German child. In R. Sackmann (ed.), *Theoretical linguistics and grammatical description*. Amsterdam: Benjamins
- Freudenthal, D., Pine, J., Jones, G., & Gobet, F. (2015). Defaulting effects contribute to the simulation of cross-linguistic differences in Optional Infinitive errors. In *Proceedings of the 37th annual meeting of the Cognitive Science Society* (pp. 746-751).
- Laalo, K. (2003). Early verb development in Finnish: A preliminary approach to mini-paradigms. In D. Bittner, W.U. Dressler & M. Kilani-Schoch (eds.), *Development of verb inflection in first language acquisition: A cross-linguistic perspective*, 323-350. Berlin: Mouton de Gruyter.
- MacWhinney, B. (2000). *The CHILDES Project: Tools for analyzing talk* (3rd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.
- MacWhinney, B., Leinbach, J., Taraban, R. & McDonald, J. (1989). Language learning: cues or rules? *Journal of Memory and Language*, 28, 255-277.
- Mills, A. E. (1985). The acquisition of German. In D. Slobin (ed.), *The cross-linguistic study of language acquisition*. Hillsdale, NJ: Erlbaum.
- Räsänen, S. H. M., Ambridge, B. & Pine, J.M. (2014). Infinitives or bare stems? Are English-speaking children defaulting to the highest frequency form? *Journal of Child Language*, 31, 1-30.
- Szagun, G. (2004). Learning by ear: On the acquisition of case and gender marking by German-speaking children with normal hearing and with cochlear implants. *Journal of Child Language*, 31, 1-30.