


RESEARCH ARTICLE 

Definite and indefinite article accuracy in learner English: A multifactorial analysis

Kateryna Derkach  and Theodora Alexopoulou

University of Cambridge, Cambridge, England

Corresponding author: Kateryna Derkach; Email: kate.derkach17@gmail.com

(Received 30 December 2021; Revised 19 July 2023; Accepted 14 August 2023)

Abstract

We present a learner corpus-based study of English article use (“a”/“the”/∅) by L2 learners with four typologically distinct first languages (L1s): German and Brazilian Portuguese (both have articles), Chinese and Russian (no articles). We investigate several semantic and morphosyntactic factors—for example, specificity, prenominal modification that can affect article use. Our analysis of 660 written scripts from the Education First Cambridge Open Database confirms the lower overall accuracy of learners with no-article L1s. Our main finding is the differential effect of specificity on definite and indefinite articles: learners tend to associate specificity with “a,” which implies article omission with nonspecific indefinite singulars and overuse of “a” with specific indefinite mass nouns. Prenominal modifiers further contribute to perceived specificity, leading to article overuse with modified indefinite mass nouns. However, in definite contexts, prenominal modifiers are associated with increased article omission.

Introduction

Articles present persistent challenges for second language (L2) English learners (henceforth, learners), even at near-native level (DeKeyser, 2005; Murakami & Alexopoulou, 2016a). Articles are among the highest frequency English words, with “the” always topping frequency lists (Leech et al., 2001). However, bare nominals (i.e., those not preceded by an article, henceforth labeled ∅ contexts¹) can account for approximately 50% of all nominals in native English (Master, 1997). Therefore, learners may notice the frequently occurring articles in English but struggle to understand why they often appear to be omitted.

Article use is influenced by learners’ first language (L1) and proficiency and also the linguistic characteristics of nominals: morphosyntactic (number, countability) and

¹We remain agnostic regarding the representation of ∅ in speakers’ minds (if any). We use “∅” to mean “no article.”

semantic (definiteness, specificity) features, article discourse functions.² Many studies have demonstrated the relevance of such features (Ionin et al., 2004; Liu & Gleason, 2002; Robertson, 2000; Snape, 2008; Trenkic, 2007), but none have considered these features together, so their relative importance and interactions are unknown. We aim to fill this gap and build a more comprehensive picture of learner article use.

Background

Articles in English and other languages

Languages vary considerably with respect to articles and their use (Lyons, 1999). Many major world languages lack articles (e.g., most Slavic languages, Hindi, Japanese), and some only have the definite (Arabic, Hebrew) or the indefinite (Turkish) article (Dryer, 2013).

In English, the distribution of the three options (“the,” “a,” and \emptyset) is determined by semantic and morphosyntactic factors. Whereas the definite/indefinite distinction is semantic, the use of “a” (vs. \emptyset) is based on countability and number—namely, “a” is required for count singulars. Nevertheless, count singulars appear bare in English in such frequently used expressions as “go to school/hospital,” “travel on foot/by car/plane,” “stay in bed/at home.” Thus, extracting patterns from article distribution in target English creates challenges for learners.

The article systems of other languages ([+art], e.g., Germanic, Romance languages) may deviate from the English pattern—for example, some Romance languages do not allow bare mass nouns in argument position (“I eat *meat*”—“Je mange *de la viande*” in French).

Though definiteness is not grammaticalized in languages that have no articles, [-art], they often have notions akin to definiteness (e.g., familiarity, specificity), which can be expressed via lexical means (demonstratives “this/that”) or syntactically (word order). Additionally, definiteness may be linked to other discourse and nominal features in distinct ways—for example, in Chinese only definite and human referents can be marked as plural (Lardiere, 2009).

Given the variation in article presence and the meanings articles grammaticalize across languages, articles are expected to present challenges for [-art] and potentially for [+art] learners.

Key semantic distinctions

Definiteness

The literature on definiteness is vast and outside our scope (Frege, 1960; Hawkins, 1978; Lyons, 1999; Russell, 1905, inter alia). It suffices to state here that the definite article makes presuppositions of uniqueness and existence, whereas the indefinite article implicates nonuniqueness (Hawkins, 1991; Heim, 2019). In (1) the speaker asserts that the book exists and is uniquely identifiable by the speaker and the hearer. When the condition of uniqueness is not met, the indefinite article is used (2).

(1) I bought *the book* about definiteness.

²Although discourse functions were initially investigated, the results were inconclusive and are not discussed here, although we acknowledge they play a role in learner article use, as shown in previous research (details in the Online Supplementary Materials).

(2) I bought *a book* about definiteness.

Whether uniqueness can be presupposed depends on the knowledge shared by the speaker and the hearer, available from the discourse and/or immediate or larger, pragmatic/cultural context. In (1) there may be only one book on definiteness in existence or known within the speaker and hearer's context (presupposition based on context). Alternatively, the speaker and the hearer may have discussed a certain book on definiteness (presupposition based on discourse), although they may be aware of others. In contrast, in (2) the speaker assumes that neither common knowledge nor discourse can help the hearer to uniquely identify the book.

For mass/plural referents, uniqueness applies to the entire mass or set, respectively (Heim, 2019, p. 36). Example (3) refers to all the applicants who failed their exams, who should be known (e.g., all the applications already received). However, in (4) any applicant satisfying the description has to be rejected, but the entire set is unknown (e.g., applications are still coming in). Similarly, in (5) the speaker implies that all the expired meat was discarded, whereas in (6) there may be some that was not.

(3) We have to reject *the applicants* who failed their exams.

(4) We have to reject *applicants* who failed their exams.

(5) They discarded *the expired meat*.

(6) They discarded *expired meat*.

Specificity

A property cutting across the definiteness/indefiniteness distinction is specificity (not directly encoded in English). As with definiteness, there is extensive literature on the topic. Despite the variation in terminology and criteria for defining specificity, one common feature across the different accounts is the “referential intention,” or the speaker's communicative intention to refer to something they have in mind (von Heusinger, 2019).

In this article, we define a specific referent as one that refers to a certain entity that exists in the world and that the speaker has in mind. For example, “a book” in (2) is specific, as it is referring to a certain existing book. This definition is similar to Bickerton's “specific reference” (1981). Note the main difference from the definition adopted in some influential literature on L2 articles (most notably Ionin et al., 2004, and replications), which is that noteworthiness (i.e., whether the speaker deems the referent noteworthy for the discourse) is not an essential feature. Therefore, (2) is specific when introducing a new referent into the discourse (first-mention indefinite) even if the identity of the book is irrelevant for further discourse (not noteworthy).

Nonspecific readings become possible in semantically opaque contexts—that is, those involving opacity-creating operators, such as verbs of propositional attitude (e.g., “want,” “believe”), negation, questions, conditionals, modals, future, and intensional verbs (e.g., “look for”). In (7), the speaker may have a certain book in mind (specific reading) or may be satisfied with any book about definiteness (nonspecific reading).

(7) I need *a book* about definiteness.

The two possible readings are said to have wide scope (specific) or narrow scope (nonspecific). Such ambiguity can also occur in nonopaque contexts, as in (8) and (9) taken from Lyons (2009, p. 172).

- (8) I haven't started the class yet; I'm missing *a student*—Mary's always late.
 (9) I haven't started the class yet; I'm missing *a student*—there should be fifteen, and I only count fourteen.

In (8), “a student” refers to a certain student and is, thus, referential. In (9), “a student” has no fixed reference and is, thus, nonreferential. This referentiality distinction also applies to definites: in (10) the winner is known, unique, so the nominal is referential, whereas in (11) there is usually one winner, but their identity is unknown because the competition has not finished yet, so “the winner” has no fixed reference.

- (10) *The winner* got \$1000.
 (11) *The winner* gets \$1000. (stated before the start of the competition)

Our choice of the term “specificity” rather than “referentiality” is based on Lyons's suggestion to use “specificity” as an “informal cover term” (1999, p. 173) to include both the wide/narrow scope and the referential/nonreferential distinctions.

To summarize, our definition of specificity is looser than that in some literature on L2 articles. The only essential criteria for specificity here are that the nominal refers to an existing (in a general sense) entity and that the reference is fixed.

L2 article accuracy: Previous research

L1

Studies comparing L2 English article accuracy by [+art] and [-art] learners unequivocally suggest that the latter have significantly more difficulties. This has been observed in naturalistic spoken data (Master, 1987; Thomas, 1989), gap-fill tasks (Hawkins et al., 2006; Ionin et al., 2008; Reid et al., 2006; Snape, 2008), and in a large-scale corpus-based study with various [+/-art] learners (Murakami & Alexopoulou, 2016a).

Semantic features: Specificity

Early studies documented increased use of “the” in [+specific;-definite] contexts (following Bickerton's definition), such as first-mention indefinites (Huebner, 1985; Thomas, 1989). Butler (2002) found evidence of using “the” with [+specific;-definite] referents (following Bickerton) in L1-Japanese learners (12). Additionally, lower-level learners often used “the” in the presence of prenominal modifiers, which they thought indicated specificity (in an informal sense).

- (12) School has just begun and I have already made the terrible mistake.

Hua & Lee (2005) found L1-Chinese learners tended to accept ungrammatical bare count singulars which were nonspecific—for example, “*Computer* is an electronic device,” as opposed to specific “*Computer* stands on the top of the office desk,” in subject position, regardless of nouns' concreteness/abstractness.

Ionin et al. (2004) hypothesized that [-art] learners interpret articles as markers of both definiteness and specificity³ or of specificity alone demonstrating a fluctuation pattern. In forced-choice elicitation (gap-fill), L1-Russian and L1-Korean participants used “the” more with [+specific;-definite] (13) vs. [-specific;-definite] (14) referents and used “a” more with [-specific;+definite] (15) vs. [+specific;+definite] (16) referents.

(13) I am visiting *the* [a] friend from college—his name is Sam Brown [...].⁴

(14) I am hoping to find *a new good friend!*

(15) We are trying to find *a* [the] murderer of Mr. Peterson—but we still don’t know who he is.

(16) Tomorrow, I’m having lunch with *the creator* of this comic strip—he is an old friend of mine.

(from Ionin et al., 2004)

Ionin et al.’s (2004) replications showed similar patterns in [-art] learners (Japanese: Hawkins et al., 2006; Reid et al., 2006; Mandarin-Chinese: Snape, 2009; Trenkic, 2008) and no fluctuation in [+art] learners (Spanish: García Mayo, 2009; Ionin et al., 2008; Reid et al., 2006; Ting, 2005; Greek: Hawkins et al., 2006). However, Ionin et al. (2008) observed the effect in L1-Russians only in indefinite contexts—that is, (13) but not (15). Meanwhile, Ting (2005) did not find evidence of fluctuation in L1-Mandarin-Chinese learners. Snape et al. (2006) suggest (based on Li & Thompson, 1981) that this may be because Chinese is developing definite (demonstrative “*nei*”—“that”) and indefinite (numeral “*yi*”—“one”) articles.

To establish specificity—speaker intent to refer to a noteworthy item—Ionin et al. included explicit statements of speaker knowledge (ESK)—for example, “He is meeting with the director of his company. *I don’t know who that person is*” (Trenkic, 2008, p. 13, emphasis added). By adding another type of test item (specific but denying speaker familiarity), Trenkic’s replication showed that learners tended to overuse “a” in definite [-ESK] contexts and overuse “the” in indefinite [+ESK] contexts regardless of specificity.

Trenkic (2007) alternatively suggests that [-art] learners misanalyze articles as adjectives (therefore, optional elements) taking “the” and “a” to mean “definite/identifiable” and “indefinite/unidentifiable,” respectively. In production data from L1-Serbian, she observed increased article omission with adjectives. Trenkic concluded that articles in [-art] learners are not syntactically motivated but produced as lexical items motivated by the pragmatic need to express the meaning learners assign to them. Based on Trenkic’s suggestions, we hypothesize that learners may erroneously assign the meaning “specific”/“nonspecific” to “the”/“a” in addition to or instead of the meaning “definite”/“indefinite”. Trenkic argues that pronominal modification

³Ionin et al. defined specificity differently, as speaker’s intent to refer to something noteworthy. For example, in “He got lots of gifts—books, toys. And best of all—he got *a puppy!*,” “a puppy” is nonspecific because its identity is “irrelevant for the discourse” (2004, p. 23). Previously discussed authors (Butler, 2002; Huebner, 1985; Thomas, 1989, and others) would consider this example specific, as it refers to a certain existing puppy.

⁴In all examples the original spelling and grammar are kept. Nominals of interest are italicized. Article corrections are given in square brackets. Erroneous articles are struck through.

correlates with article omission because producing lexical items drains cognitive resources leading to omission of less (communicatively) important items.

To summarize, naturalistic and elicited data from participants with various L1s have shown that [-art] learners might be uncertain about whether English articles signify specificity or definiteness, although the effect is not always observable (Hua & Lee, 2005; Ting, 2005; White, 2003).

Nominal features

Number and countability play an important role in learners' article use. Lardiere (2004) suggests, based on a feature-assembly approach, that the inclusion of number and countability features is what makes the indefinite article more complex and difficult to acquire (also Hawkins et al., 2006). This is confirmed by Zhao and MacWhinney (2018) working within the competition model. Their L1-Chinese intermediate-advanced learners were more accurate on a cloze test with postmodified nominals, which are a cue for "the" (e.g., "the book which you recommended"), than with nominals without postmodifiers, which may be a cue for "a" (count singulars) or \emptyset (mass/plurals).

Trenkic's (2002) study considered multiple nominal features, including number, countability, and abstractness. Lower-intermediate to advanced L1-Serbian learners were more accurate in obligatory "the" vs. obligatory "a" contexts (text-translation task) and showed evidence of "the-flooding" (Huebner, 1983), specifically using "the" instead of "a" with count-singular indefinites.

Additionally, Trenkic (2002) revealed that learners were more accurate in supplying "a" with abstract (e.g., "environment") than with concrete (e.g., "book") nouns. She explains the initially surprising pattern by hypothesizing that learners use "a" to individuate concepts with no clear boundaries ("fuzzy" concepts), such as abstract referents. This accounts for the higher "a" omission with concrete nouns, which do not require such individuation. Trenkic also suggests that "the" may be perceived as a marker of "definite" or concrete form, explaining its higher incidence with indefinite concrete count singulars. This is despite Trenkic's initial expectation that learners would struggle more with "fuzzy" abstract referents. In comparison, Hua & Lee's (2005) L1-Chinese participants would often misdetect countability of abstract referents: in a grammaticality judgment task, they more readily accepted "much sentence" (abstract) than "much computer" (concrete). This is, however, indirect evidence, as learners accepting "much sentence" would not necessarily omit "a" before "sentence." In Butler (2002), up to 20% of L1-Japanese learners' errors were due to misdetection of countability, especially of abstract referents, although the paper does not detail error types, so it is not easily comparable with Trenkic (2002).

Snape (2008) argues that learners with L1s without a count/mass distinction (e.g., Japanese) may associate definiteness with number features—that is, use articles only with count singulars, as bare mass/plural nouns are allowed in English. In a forced-choice elicitation task (definites only), advanced L1-Japanese participants tended to omit "the" with plural/mass nouns—for example, "I've just finished our new patio. [...] Mixing [the] cement was difficult" (2008, p. 70), whereas advanced L1-Spanish learners omitted "the" to a lesser extent and only with plurals in larger situation/cultural contexts—for example, "Have you seen [the] bridesmaids?" at a wedding (2008, p. 66).

In summary, number and countability affect article accuracy of [-art] learners, who may generalize the bare mass/plural indefinite pattern in English to all mass/plural nouns. The count/mass distinction, especially in abstract nouns, further influences article accuracy if learners' L1 lacks grammatical number. Learners may also use "a" to

individuate concepts without clear boundaries (abstract referents, mass nouns), whereas “the” may mark clearly bounded entities (concrete referents, count singulars).

Research objectives

Based on previous research and cross-linguistic differences, multiple factors may influence learners’ article interpretation and use: [+/-art] L1 and proficiency level, semantic features not encoded in English (specificity, familiarity, abstractness), and morphosyntactic and syntactic features (number, countability, syntactic position, premodification).

We seek to investigate how these factors interact to predict article accuracy and error types (omission, substitution, overuse) in learners. The benefit of including multiple factors is the ability to analyze their interactions, which may reveal differential effects. By differential effects we mean that an interaction between specificity and definiteness, for example, might mean that learners tend to supply “the” with specific definites but omit “a” with specific indefinites. This is a step forward from previous research, where findings were either limited to a certain category of nominals (Ionin et al., 2004, and replications often include only concrete count singulars) or aggregated across different categories (Murakami & Alexopoulou, 2016a, a large-scale study aggregated results across number, countability, abstractness, etc.).

Including multiple factors can also help avoid overestimating the effect of any one predictor, as multiple regression modeling allows for the estimation of the effect of each variable while keeping the rest constant.

Methodology

Learner data

Corpus

We used a large open-access learner corpus, EFCAMDAT, the EF Education First Cambridge Open Language Database (Alexopoulou et al., 2015, 2017; Geertzen et al., 2013; Michel et al., 2019). EFCAMDAT contains 1,180,310 writings (scripts) responding to communicative tasks (holiday postcards, film reviews, describing personal experiences), submitted by registered learners of a large number of nationalities (anonymized) to EF’s online language school.⁵ EFCAMDAT contains 16 proficiency levels aligned with the Common European Framework of Reference (CEFR) and is pseudolongitudinal (most learners do not complete all the levels). New learners are placed in levels 1/4/7/10/13 based on placement tests. Each level comprises eight modules, each ending with a writing task. Because EFCAMDAT does not contain direct L1 data, National Language (NL), crossing nationality with country of access to the online school, is used as a proxy for L1. This has been shown to be quite reliable (Alexopoulou et al., 2015; Murakami, 2013) despite the inevitable noise in the data (e.g., multilingualism is not captured).

⁵Because the writings were completed offline (as homework) with access to resources, learners’ accuracy tends to be higher on EFCAMDAT when compared with an exam learner corpus, such as Cambridge Learner Corpus (CLC). This means learners may use external help; however, these tasks are very low stakes, and learners would not benefit in any way from submitting perfect answers. Indeed, the accuracy patterns for different morphemes are similar in EFCAMDAT and CLC across proficiency levels (see comparison in the Online Supplementary Materials).

Topic (prompt): Describe a crime

Level: 4 (CEFR A2)

L1 Brazilian Portuguese (writing ID 817284)

Hi, ~~the~~ last friday in my city two thieves stole ~~the~~ motorcycle from ~~the~~ man, but the police arrested one ~~the~~ thief. he had a camera in ~~to~~ ^{the} helmet and filmed it all, he published the film into ^{the} web, the police fired at the thief but not killed.

L1 Russian (writing ID 566460)

A Few days ago in my city. Somebody broke in ^{an} administrative building and stole ^{an} ATM. ^{The} Thieves broke down a window, then they loadet ^{the} ATM in ^a car. ^{The} ATM had approximately \$50000 cash. Two days later ^{the} police arrested ^{the} thieves. It was former security guard and his mate.

Figure 1. Example scripts with article errors marked.

Subcorpus

We sampled writings from two [+art] subgroups (German and Brazilian Portuguese, henceforth Brazilian), and two [-art] subgroups (Chinese, Russian). We included two languages from each language type to tease apart the typological effect of article presence/absence from potentially L1-independent effects (Murakami & Alexopoulou, 2016a).

We sampled 660 scripts (165 per NL), which according to our power analysis based on a simulated dataset in R (R Core Team, 2021) would provide 80% statistical power (details in the [Online Supplementary Materials](#)).

The scripts were randomly selected and equally distributed across A2–B2 CEFR levels (inclusive), corresponding to EF levels 4–12,⁶ and equally distributed across topics—that is, the specific writing prompts at the end of each module (e.g., “Write a short autobiography”)—and no more than one script was contributed by the same learner (examples in [Figure 1](#)).

Coding

We manually retrieved all nominals from the scripts. We treat as a nominal any phrase consisting of a noun, an optional article, and a prenominal modifier—for example, “book,” “an interesting book,” “the books”—but excluding demonstratives and quantifier items, such as “this book” or “many interesting books” (for full list of exclusions, see the [Online Supplementary Materials](#)).

We excluded formulaic sequences (e.g., “for example,” “all over/around the world,” “in the morning,” “twice a week,” “make a long story short,” etc.), which are expected to have higher accuracy rates (Myles, 2012; details in the [Online Supplementary Materials](#)).

Additionally, we excluded sequences provided in writing prompts and model answers, which learners often copied in their writing. All the coded

⁶We excluded lower levels, where the writings are mostly formulaic, and higher levels, where there is generally less data in the corpus.

Table 1. Variable coding

Variable	Levels	Note/Examples
NL Level	de—German, br—Brazilian, cn—Chinese, ru—Russian 25–96 (72 total)	Each number on the scale represents one module. Eight modules comprise one EF level, e.g., modules 25–32 comprise EF level 4. Three EF levels cover one CEFR level, e.g., EF levels 4–6 cover CEFR level A2.
Topic ID (prompt)	1–123	Each module ends in a writing prompt, referred to as “topic” in the corpus. We have more topic IDs (123) than the total number of modules (72) because some modules had two prompt options (see Shatz, 2020).
Definiteness Target article Response	definite/indefinite “a”/“the”/no article (∅) correct/incorrect	Correct if “a”/“the”/∅ is used appropriately
Error type	omission of: “a,” “the”; overuse of: “a,” “the”; substitution of “a” instead of “the,” substitution of “the” instead of “a”	Omission: failure to supply “a”/“the” in obligatory “a”/“the” contexts; Overuse: use of “a”/“the” in target ∅ contexts; Substitution: use of “a” in obligatory “the” contexts or use of “the” in obligatory “a” contexts.
Noun type (Ntype)	count singular count plural mass	Considered context: “cake” count in “they ate a cake,” mass in “they ate cake”
Abstractness	abstract/concrete	Coded as binary, but we acknowledge it is a gradient characteristic (Scott et al., 2019). Considered context: “things” concrete in “I saw interesting things,” abstract in “I learned interesting things.”
Syntactic position ⁷	(a) subject (b) object (c) predicate (d) existential	including objects of verbs and of prepositions nominals following copula “be” (and “become”). Included comparative constructions with “as” and “like” (e.g., “as tall as a giraffe,” “work like a robot”). nominals following “there” + “be” constructions, e.g., “there is dust on the table”
Specificity Prenominal modifier	specific/nonspecific present/absent	Definition in “Background” e.g., “a science book”/“a book”

variables are listed in Table 1 (see coded example in the [Online Supplementary Materials](#)).

⁷Other syntactic positions were excluded because they constituted less than 10% of the data (temporal modifiers, e.g., “last week”; appositives, e.g., “Tom, *the leading man*”; genitives, e.g., “*people’s* attitudes”).

Table 2. Distribution of nominals retrieved

EF levels	NL				Total
	de	br	cn	ru	
4	65	85	95	86	331
5	107	119	71	96	393
6	95	128	148	125	496
7	110	153	110	127	500
8	177	121	128	158	584
9	161	188	131	194	674
10	237	254	191	206	888
11	246	218	244	247	955
12	263	287	138	263	951
Total	1,461	1,553	1,256	1,502	5,772

To determine coding reliability, 100 randomly selected items were coded by another doctoral student of applied linguistics (English native speaker). The level of agreement was strong for all variables, $\kappa > 0.85$ (McHugh, 2012).

The resulting subcorpus contained 5,772 nominals (Table 2).

Statistical modeling

Binomial and multinomial mixed-effects logistic regression

We investigated the effect of our independent variables on article accuracy using a generalized linear mixed-effects logistic regression model (henceforth, accuracy model), where the dependent variable is binary (correct/incorrect article [non]use), using the `lme4` package in R (Bates, Mächler, et al., 2015).

For further analysis of error types, we used multinomial logistic regression models, which allow for more than two levels in the categorical outcome, using the `mclorfit` R package (Elff, 2021). This type of model estimates predictor variable effects on the change in the odds of the different outcomes (omission, substitution, overuse) compared with a chosen baseline outcome (no error). Thus, we could explore, for example, whether a mass noun as opposed to a count singular increases the odds of article omission versus no error.

For error type analysis, our data was split into two subsets by target article,⁸ each with different possible outcomes.⁹

- Indefinite count singular (obligatory “a,” 1,679 observations): no error, omission of “a,” substitution (“the” instead of “a”)
- Indefinite plural and mass (obligatory \emptyset , 2,060 observations): no error, overuse of “the,” overuse of “a”

⁸We intended to analyze three subsets, but the model on definites did not converge, most probably due to the uneven error distribution across noun types (only five substitution errors in mass, no substitution errors in plural contexts) and the fact that most errors were made by L1-Russians and L1-Chinese. The model converged without L1-Germans; however, adding any interactions led to more convergence issues, whereas the pseudo- R^2 for the no-interaction model was only .017. Thus, we are not reporting this model here.

⁹The split created difficulties with random effects structures: we could only include random intercepts by `wr_ID`, as adding any random slopes resulted in singular fits.

Model selection

When choosing the fixed-effects structure for each model, we started with a comprehensive model including all the independent variables of interest (without interactions). We then attempted interactions that were either theoretically motivated or that appeared likely from the visual examination of the data. We compared models using log-likelihood ratio tests (LRT) and excluded any interactions that did not significantly improve model fit. We did not exclude any independent variables unless they caused convergence issues or led to inadequate coefficient or standard error estimates.

Random effects were included in all models because each writing (*wr_ID*) contained multiple observations. We also included the topic (prompt) ID as a random effect to capture potential prompt effects where possible. In choosing the random effects structure, we adhered to the parsimonious approach by Bates, Kliegl, et al. (2015). They argue for removing any random-effects components that explain little (close to 0%) variation (as estimated using their random-effects principal components analysis function from the *RePsychLing* package in R) and that do not contribute significantly to improving model fit (as estimated by the LRT). We also excluded any random-effects components that caused convergence issues. The model selection process for each model is described in the [Online Supplementary Materials](#).

Finally, we used the *emmeans* package in R (Lenth, 2022) for post hoc pairwise comparisons¹⁰ and for generating predicted values of the outcome variable.

Analysis and results

Data distribution and observed accuracy rates

Our data distribution across the main variables is presented in mosaic plots in [Figure 2](#).

The plots reveal the following patterns:

- Top right: the three target contexts are approximately equally represented (this shows the number of contexts in which each target is expected, or would be correct, not actual article use by learners).
- Top left: definites tend to be specific—that is, refer to a certain existing (in a general sense) individual or entity, and indefinites nonspecific (especially mass/plural).
- Middle left: definites are less often premodified than indefinites, and mass nouns are not as frequently premodified as count singulars irrespective of definiteness.
- Bottom left: mass indefinites tend to be abstract, whereas in other categories abstract and concrete nouns are equally distributed.
- Bottom right: definites are more often used in subject position than indefinites.

[Figure 2](#) demonstrates that although there are tendencies and potential correlations between semantic features and individual articles, there is no one-to-one mapping, thus confirming the learning challenge.

[Figure 3](#) shows accuracy rate measured as the number of correct uses (including correct \emptyset) divided by the total number of obligatory contexts (including target \emptyset).¹¹ Note that L1-Chinese and L1-Russians (both [-art]) do not pattern together (see Murakami & Alexopoulou, 2016b, for similar findings). L1-Chinese learners' scores

¹⁰*emmeans* automatically applies the Tukey adjustment method when comparing families of > 2 estimates.

¹¹This is unlike target language use (TLU) often used in the literature, which does not include correct \emptyset contexts. See comparison in the [Online Supplementary Materials](#).

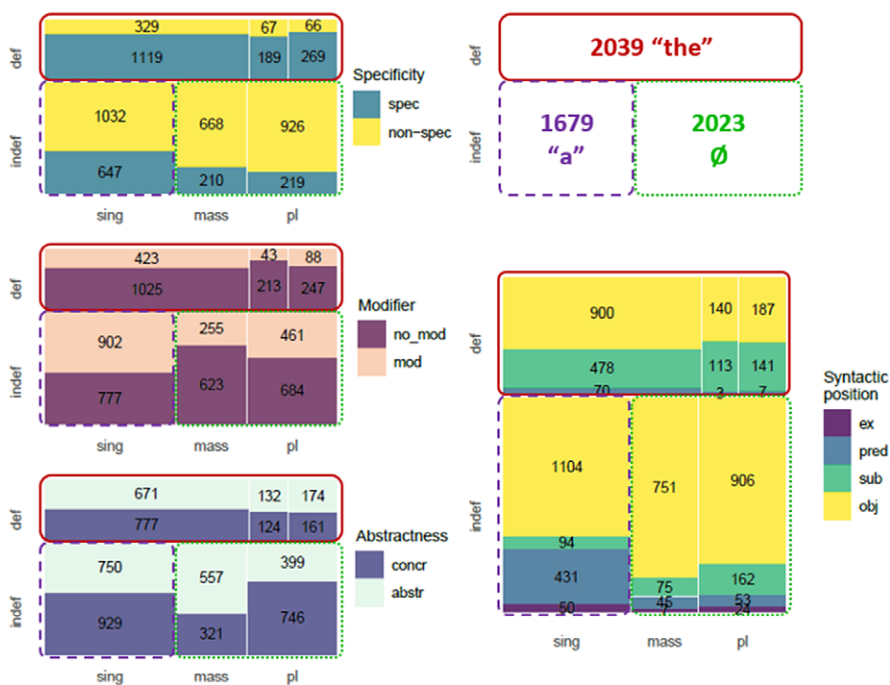


Figure 2. Observed distribution of target contexts across target, specificity, modifier, abstractness, syntactic position.




Figure 3. Development across EF levels.

may appear rather high. Nevertheless, they are consistently less accurate than L1-Germans. Because L1-Chinese behave very similarly to L1-Brazilians, we further analyze each NL separately without combining them into [+/-art] types.¹²

¹²A reviewer suggested it might still be worth comparing [+art] and [-art] groups and including NL as a random effect in the models. However, the minimum number of levels required to obtain a reasonable

Table 3. Error-type distribution

Error type	Number of errors	% of total errors
Omission	581	61%
Overuse	240	25%
Substitution	133	14%
Total	954*	100%

*16% of all observations.

The distribution of error types is detailed in Table 3 and Figure 4.

In obligatory article contexts, 81% of errors are omissions (see Figure 1 for examples). In target \emptyset contexts, overuse of “a” and “the” is mostly equal for mass nouns, whereas “the” overuse is more common in plurals. Generally, the patterns are similar across NLs, with two exceptions: (a) the use of “a” (a_sing), where Germans have a higher proportion of substitutions than others, and (b) the use of \emptyset with mass nouns (zero_mass), where Germans overuse “the” more often than “a,” whereas other NLs overuse both equally, but the overall error count for L1-Germans is small.

Predictors of accuracy

The accuracy model (Table 4) reveals significant effects of NL, proficiency level (interacting with NL), specificity, and modifier, which vary by target article and

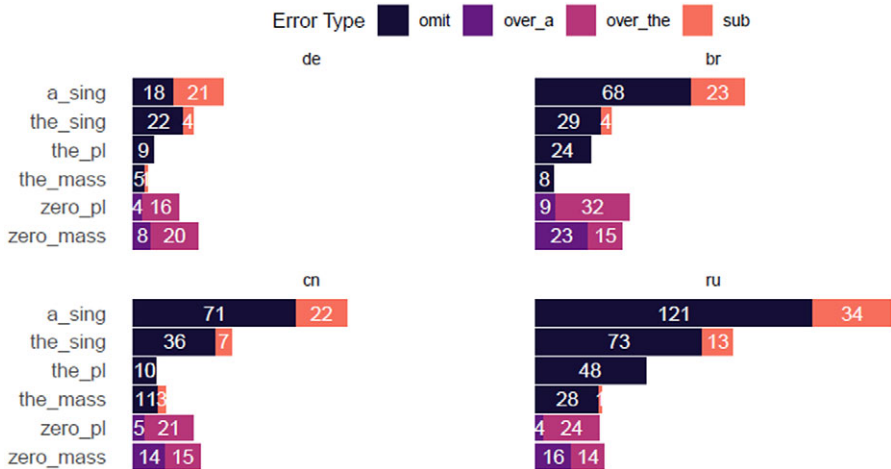


Figure 4. Error-type distribution across NL, target article, and noun type. Numbers represent instances.

random effect estimation is 5–6 (Bolker, 2022), whereas NL has only 4. Indeed, an attempted model with NL as a random effect did not converge. Had L1-Chinese performed similarly to L1-Russians, we might have grouped learners according to [+/-art] and ignored the NL subdivision. However, as L1-Chinese appeared to pattern with L1-Brazilians, grouping them with L1-Russians would make little sense.

Table 4. Accuracy model formula

Component	Formula
Fixed effects	Score ~ definiteness*Ntype*(NL*level + specificity + modifier) + abstractness + syntactic position +
Random effects	(definiteness + Ntype + specificity + abstractness wr_ID)

noun type, and also effects of syntactic position and abstractness¹³ (full results in Table 5).

NL in interaction with proficiency level

Figure 5 illustrates the effect of NL across definiteness and noun type (averaged across proficiency levels). The interaction with definiteness and noun type stems from the fact that the target is different across combinations of variable levels (“the” for all definites, “a” for count singular indefinites, Ø for mass/plural indefinites). Thus, the top three facets and the bottom-left facet of Figure 5 reflect the rate of suppliance of “the” and “a,” whereas the bottom-middle and right facets show accuracy rates in Ø contexts, where any errors would be overuse.

All but L1-German learners are significantly less accurate with “a” than with “the” in singular contexts (L1-Brazilian and L1-Chinese: $p < .001$; L1-Russian: $p = .035$).




Figure 5. The effect of NL across definiteness and noun type.¹⁴

¹³The effect of abstractness, although statistically significant ($p = 0.024$), is very small (89% predicted accuracy for concrete nouns, 95% CI [85%, 92%], and 91% predicted accuracy for abstract nouns, 95% CI [88%, 93%]) and, thus, not discussed further.

¹⁴We have labeled y-axes “Predicted accuracy rate” throughout for ease of interpretation, although the numbers in fact represent predicted probabilities of correct article (non)use in a single instance, which is conceptually the same. In all figures, error bars represent 95% confidence intervals.

Table 5. Accuracy model results

Predictors	Accuracy rate		
	Log odds	SE	95% CI
Intercept (grand mean)	2.57***	0.17	2.24 2.89
def:definite	-0.04	0.09	-0.21 0.14
Ntype:singular	-0.18	0.11	-0.39 0.03
Ntype:mass	-0.29*	0.14	-0.56 -0.02
NL:German	0.87***	0.15	0.58 1.16
NL:Brazilian	0.04	0.12	-0.19 0.27
NL:Chinese	-0.06	0.12	-0.31 0.18
Level	0.12	0.07	-0.02 0.26
spec:specific	0.02	0.08	-0.13 0.18
mod:no_modifier	0.18**	0.07	0.05 0.31
abstr:concrete	-0.12*	0.05	-0.22 -0.02
synt:existential	0.72*	0.35	0.03 1.41
synt:predicate	0.15	0.17	-0.18 0.49
synt:subject	-0.43**	0.15	-0.73 -0.12
def:definite*Ntype:singular	0.36***	0.09	0.17 0.54
def:definite*Ntype:mass	-0.03	0.13	-0.28 0.22
NL:German*Level	0.10	0.15	-0.19 0.39
NL:Brazilian*Level	-0.15	0.11	-0.37 0.07
NL:Chinese*Level	-0.07	0.12	-0.31 0.16
def:definite*NL:German	0.30*	0.15	0.02 0.59
def:definite*NL:Brazilian	0.19	0.11	-0.03 0.42
def:definite*NL:Chinese	0.21	0.12	-0.03 0.45
def:definite*Level	0.05	0.07	-0.09 0.19
def:definite*spec:specific	0.01	0.08	-0.14 0.16
def:definite*mod:no_modifier	-0.06	0.07	-0.19 0.07
Ntype:singular*NL:German	0.01	0.16	-0.31 0.33
Ntype:mass*NL:German	-0.18	0.23	-0.64 0.27
Ntype:singular*NL:Brazilian	0.10	0.13	-0.15 0.36
Ntype:mass*NL:Brazilian	0.23	0.18	-0.12 0.59
Ntype:singular*NL:Chinese	-0.12	0.14	-0.39 0.14
Ntype:mass*NL:Chinese	-0.08	0.18	-0.43 0.28
Ntype:singular*Level	-0.05	0.08	-0.21 0.11
Ntype:mass*Level	0.01	0.11	-0.21 0.22
Ntype:singular*spec:specific	0.18*	0.09	0.01 0.35
Ntype:mass*spec:specific	-0.03	0.11	-0.25 0.18
Ntype:singular*mod:no_modifier	-0.00	0.07	-0.15 0.15
Ntype:mass*mod:no_modifier	0.22*	0.11	0.01 0.43
def:definite*NL:German*Level	0.04	0.15	-0.25 0.33
def:definite*NL:Brazilian*Level	-0.17	0.11	-0.39 0.05
def:definite*NL:Chinese*Level	0.14	0.12	-0.08 0.37
Ntype:singular*NL:German*Level	-0.09	0.16	-0.41 0.23
Ntype:mass*NL:German*Level	0.35	0.23	-0.10 0.79
Ntype:singular*NL:Brazilian*Level	0.21	0.13	-0.04 0.46
Ntype:mass*NL:Brazilian*Level	-0.15	0.18	-0.49 0.20
Ntype:singular*NL:Chinese*Level	-0.18	0.13	-0.45 0.08
Ntype:mass*NL:Chinese*Level	-0.14	0.18	-0.49 0.20
def:definite*Ntype:singular*NL:German	-0.54***	0.16	-0.84 -0.23
def:definite*Ntype:mass*NL:German	0.43	0.22	-0.01 0.87
def:definite*Ntype:singular*NL:Brazilian	-0.03	0.12	-0.27 0.21
def:definite*Ntype:mass*NL:Brazilian	0.24	0.17	-0.10 0.57
def:definite*Ntype:singular*NL:Chinese	-0.07	0.13	-0.33 0.18
def:definite*Ntype:mass*NL:Chinese	-0.25	0.17	-0.59 0.08
def:definite*Ntype:singular*Level	-0.12	0.08	-0.27 0.04
def:definite*Ntype:mass*Level	0.26*	0.11	0.05 0.46
def:definite*Ntype:singular*spec:specific	-0.09	0.08	-0.25 0.07
def:definite*Ntype:mass*spec:specific	0.33**	0.11	0.12 0.54

(Continued)

Table 5. (Continued)

Predictors	Accuracy rate		
	Log odds	SE	95% CI
def:definite*Ntype:singular*mod: no_modifier	0.32***	0.07	0.17 0.46
def:definite*Ntype:mass*mod: no_modifier	-0.13	0.10	-0.33 0.08
def:definite*Ntype:singular*NL: German*Level	-0.04	0.16	-0.35 0.27
def:definite*Ntype:mass*NL: German*Level	0.41	0.22	-0.02 0.84
def:definite*Ntype:singular*NL: Brazilian*Level	0.15	0.12	-0.08 0.39
def:definite*Ntype:mass*NL: Brazilian*Level	-0.19	0.17	-0.52 0.13
def:definite*Ntype:singular*NL: Chinese*Level	0.12	0.13	-0.13 0.37
def:definite*Ntype:mass*NL: Chinese*Level	-0.22	0.17	-0.54 0.11
Random effects by wr_ID			
SD(Intercept)	0.49		
SD(def:definite)	0.55		
SD(Ntype:singular)	0.55		
SD(Ntype:mass)	0.78		
SD(abstr:concrete)	0.43		
SD(spec:specific)	0.50		
N _{wr_id}	632		
Observations	5772		
Marginal R ² /Conditional R ²	0.202/0.256		
LRT vs. null model (only random effects)	$\chi^2(63) = 379.67, p < 1 \times 10^{-15}$		
VIFs ¹⁵	Strong for some interaction terms, but low for all terms in no-interaction model		
Overdispersion ratio ¹⁶	0.616 ($\chi^2 = 3,512.700, p = 1.0$)		

Note. Reference levels: “Russian” for NL, “plural” for Ntype, “object” for syntactic position (synt).
* $p < .05$; ** $p < .01$; *** $p < .001$.

The main NL effect concerns the significantly lower accuracy of L1-Russians in obligatory “the” and “a” contexts, which drops further in definite mass (17) and plural (18) contexts, showing a sensitivity to noun type not observed in the other NLs. The difference between L1-Brazilians and L1-Germans in plural definites also reaches statistical significance at $p = .044$; however, L1-Brazilians’ accuracy is still quite high. Note the comparatively lower number of mass/plural definite observations.

(17) My dinner was horrible! [...] [The] Red wine was too sour and [the] coffee was bitter.
(L1-Russian, A2, ID:417629)¹⁷

¹⁵VIFs (variance inflation factors) are indicators of multicollinearity: <5 low, 5–10 moderate, >10 strong collinearity to be avoided (James et al., 2013).

¹⁶Overdispersion checked using performance package in R (Lüdtke et al., 2021).

¹⁷For all corpus examples, we provide learner’s L1, CEFR level, and wr_ID.

- (18) This property is ideal for investors [...]. *[The] Apartments* may be fitted with what you would like.
 (L1-Russian, B2, ID:114705)

Another NL effect is observed within singular indefinites (target “a”), where L1-Brazilian, L1-Chinese, and L1-Russian learners are all predicted to be less accurate than L1-Germans, and L1-Russians are also significantly less accurate than L1-Brazilians (19–21).

- (19) It is [...] an example for other women and ‘prove’ that is possible to get *[a] good position* [...].
 (L1-Brazilian, B2, ID:1030969)
- (20) Then I will take *[an] air plane* to Madrid.
 (L1-Chinese, B1, ID:3441)
- (21) I’m think it is helps people to get *[a] well-paid and interesting job* in the future.
 (L1-Russian, B1, ID:786665)

In target \emptyset contexts, there are two possibilities: (a) learners may tend to omit articles across the board and happen to be correct with mass or plural indefinites (coincidentally correct use) or (b) learners may be aware that mass/plural indefinites require \emptyset (genuinely correct use). Because the corpus provides performance data only, we cannot confidently distinguish between the two. We hypothesize (a) is more likely for L1-Russians, considering their overwhelming tendency to omit articles elsewhere.

Proficiency level interacts with NL, definiteness, and noun type (Figure 6) in that it has a significant effect only in indefinite singulars (target “a”), with L1-Russians significantly growing ($p < .001$) and L1-Chinese declining ($p = .023$) in accuracy, with the two slopes being significantly different from each other ($p < .001$). Note that the estimates for mass/plural definites are rather unreliable with large confidence intervals.




Figure 6. Proficiency level by NL across definiteness and noun type.

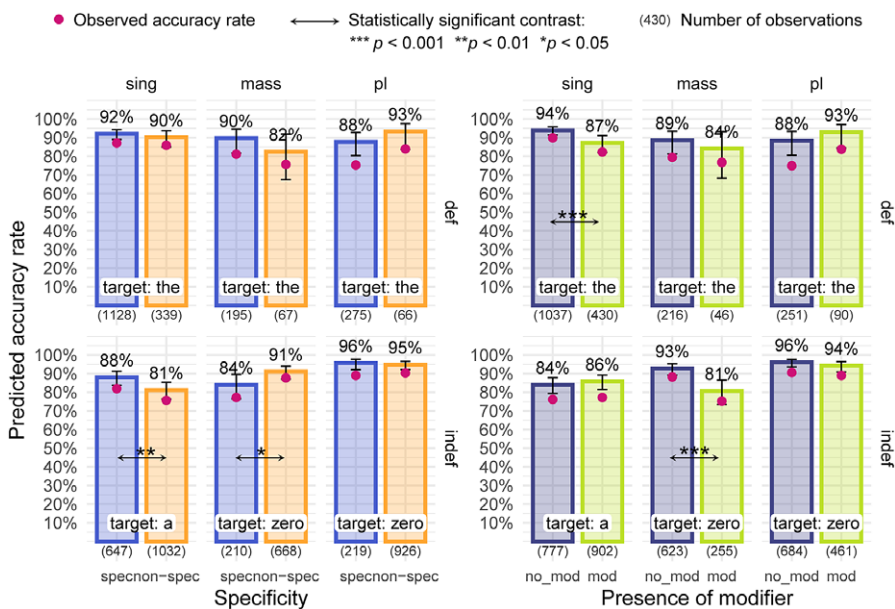


Figure 7. The effect of specificity (left) and modifier presence (right) across definiteness and noun type.

Specificity

Specificity, as defined in this study, affects only two contexts (Figure 7, left): (a) indefinite singulars (target “a”), where accuracy is significantly lower for nonspecific (22) than for specific (23) referents; (b) indefinite mass (target \emptyset), where the effect is the opposite, with significantly higher accuracy for nonspecific (24) than for specific (25) referents.

- (22) I have many dreams [...] I’d make [a] *career* in my business and have a fulfilled and balanced live.
(L1-German, B1, ID:249369)
- (23) Alexander had *an accident* last summer when he was arrived cinema.
(L1-Russian, A2, ID:807885)
- (24) You can use it to cook not only *rice* but also *congee* and *soup*.
(L1-Chinese, B1, ID:779091)
- (25) When police got the home they noticed that one servant’s face was covered with \emptyset *red paint*.
(L1-Russian, B2, ID:854608)

Modifier presence

A pronominal modifier (Figure 7, right) negatively affects accuracy in singular definites, where it increases “the” omission (26), and in mass indefinites (target \emptyset), where it increases article overuse (27).




Figure 8. The effect of syntactic position.

- (26) I first met my friend, Kolya, when I was working in advertising project five years ago. [...] Kostya and I enjoy working on [the] advertising project together.
(L1-Russian, A2, ID:887991)
- (27) Nowadays there is a [Ø] great respect and not be seen anyone smoking indoors.
(L1-Brazilian, B2, ID:664765)

Syntactic position

Errors are significantly more likely in subject (28) and object (29) positions (both at 86% predicted accuracy) than in predicate (30) position (91% predicted accuracy) at $p = .02$ and $p = .002$ (Tukey adjusted), respectively (Figure 8, note the scale starts at 60%). The 95% predicted accuracy in existential position (31) is not significantly higher than that in subject or object positions due to a larger confidence interval.

- (28) [An] Online study program give me opportunity to learn when I have free time and desire.
(L1-Russian, B1, ID:15851)
- (29) I like watching them [reality TV programmes] [...]. I can learn the [Ø] life experience from other people.
(L1-Chinese, B1, ID:135026)
- (30) [...] we were supposed to have a cosy and comfortable vessel but in fact that was just a terrible little boat.
(L1-Chinese, B1, ID:372641)
- (31) First, there was an insect in my soup!
(L1-Russian, B1, ID:157548)

Predictors of error type

Count singular indefinites

The error-type model for count singular indefinites (target “a”; Table 6) confirms the significant interaction between NL and proficiency level (Figure 9) and the significant effect of syntactic position (similar pattern to that in the accuracy model) while revealing that the differences in accuracy rates are driven by omission errors, with low substitution error rates across NLs. Additionally, we find significant interactions between specificity and NL, specificity and modifier,

Table 6. Error-type model formula for count singular indefinites

Component	Formula
Fixed effects	Error type ~ NL*level + specificity*(NL + modifier + abstractness) + syntactic position
Random effects	+ (1 wr_ID)

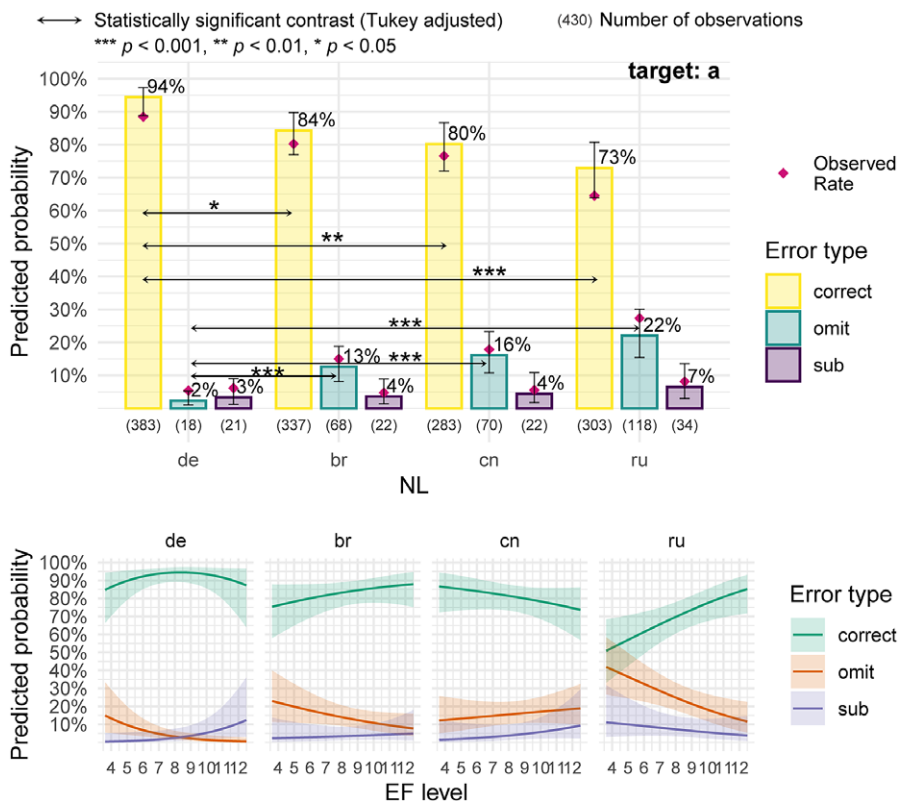


Figure 9. The effect of NL alone (top) and in interaction with level (bottom) on predicted probabilities of error types in count singular indefinites.

specificity and abstractness¹⁸ (Table 7). The model predicts 82.5% probability for correct “a” suppliance, 13.5% omission, 4% substitution (averaged across other variables).

Interaction between NL and specificity. The model reveals that only in nonspecific contexts (i.e., those not referring to certain existing entities) do all NLs show significantly lower accuracy than L1-Germans (Figure 10). Meanwhile, in specific

¹⁸The interaction between specificity and abstractness is not discussed further because of its weakness: the effect of specificity is only significant in concrete nouns, but the trend is the same for abstract nouns and is, in fact, approaching significance ($p = .06$).

Table 7. Error-type model for count singular indefinites results

Predictors	Accuracy rate		
	Log odds	SE	95% CI
omit-Intercept (grand mean)	-2.61***	0.25	-3.10 -2.12
sub-Intercept (grand mean)	-3.81***	0.33	-4.47 -3.15
omit-NL:German	-1.85***	0.40	-2.62 -1.07
sub-NL:German	-0.38	0.43	-1.22 0.46
omit-NL:Brazilian	0.19	0.25	-0.30 0.68
sub-NL:Brazilian	-0.18	0.41	-0.98 0.61
omit-NL:Chinese	0.57*	0.24	0.09 1.04
sub-NL:Chinese	-0.02	0.42	-0.84 0.80
omit~Level	-0.45**	0.14	-0.72 -0.18
sub~Level	0.32	0.24	-0.14 0.78
omit~spec:specific	-0.49***	0.14	-0.75 -0.22
sub~spec:specific	-0.20	0.18	-0.56 0.17
omit~mod:no_modifier	-0.12	0.10	-0.31 0.07
sub~mod:no_modifier	0.32*	0.15	0.03 0.61
omit~abstr:concrete	0.09	0.10	-0.10 0.28
sub~abstr:concrete	0.31*	0.16	0.01 0.62
omit~synt:existential	-0.97	0.54	-2.03 0.10
sub~synt:existential	0.07	0.64	-1.17 1.32
omit~synt:predicate	-0.09	0.26	-0.60 0.42
sub~synt:predicate	-1.75***	0.50	-2.73 -0.77
omit~synt:subject	0.61	0.33	-0.04 1.27
sub~synt:subject	1.22**	0.43	0.38 2.06
omit-NL:German*Level	-0.56	0.30	-1.14 0.03
sub-NL:German*Level	0.68	0.44	-0.19 1.55
omit-NL:Brazilian*Level	0.06	0.21	-0.36 0.47
sub-NL:Brazilian*Level	-0.15	0.39	-0.91 0.61
omit-NL:Chinese*Level	0.65**	0.21	0.23 1.07
sub-NL:Chinese*Level	0.29	0.42	-0.54 1.13
omit-NL:German*spec:specific	-0.32	0.30	-0.91 0.27
sub-NL:German*spec:specific	0.00	0.31	-0.61 0.62
omit-NL:Brazilian*spec:specific	-0.09	0.19	-0.46 0.28
sub-NL:Brazilian*spec:specific	-0.02	0.32	-0.64 0.59
omit-NL:Chinese*spec:specific	0.01	0.18	-0.35 0.37
sub-NL:Chinese*spec:specific	-0.43	0.32	-1.06 0.19
omit~spec:specific*mod:no_modifier	-0.11	0.10	-0.30 0.07
sub~spec:specific*mod:no_modifier	-0.15	0.15	-0.44 0.14
omit~spec:specific*abstr:concrete	-0.09	0.10	-0.27 0.10
sub~spec:specific*abstr:concrete	0.17	0.16	-0.14 0.47
Random effects by wr_ID			
SD omit~(Intercept)	0.51		
SD sub~(Intercept)	0.61		
N _{wr_id}	541		
Observations	1,679		
Pseudo-R ² (Nagelkerke, 1991)	0.193		
LRT vs. null model (only random effects)	$D(36) = 309.89, p < 1 \times 10^{-15}$		
VIFs			
NL	20.28 in final model 8.6 in no-interaction model		
other variables	moderate or high for some interaction terms <3 in no-interaction model		

Note. omit~ estimates for omission errors vs. correct; sub~ substitution errors vs. correct. Reference levels: "Russian" for NL, "object" for syntactic position (synt).

* $p < .05$; ** $p < .01$; *** $p < .001$.



Figure 10. The effect of specificity in interaction with NL on predicted probabilities of error types in count singular indefinites.

contexts only L1-Russians appear to be behind, although L1-Chinese also demonstrate significantly higher “a” omission than L1-Germans. In other words, in indefinite count singulars specificity has no effect on L1-Germans, who are at ceiling, or on L1-Russians, who tend to omit “a” regardless of specificity. However, L1-Brazilians ($p = .006$) and L1-Chinese ($p = .023$) omit “a” significantly more often with non-specific referents.

Interaction between specificity and modifier. There is a significant effect of specificity, as defined in this study, in non-premodified nouns, with the odds of omitting “a” dropping to 8% for specific nouns as opposed to 17% for nonspecific ones (22; Figure 11, left). The trend in premodified nouns (Figure 11, right) is the same but the difference at only 5% becomes statistically nonsignificant.

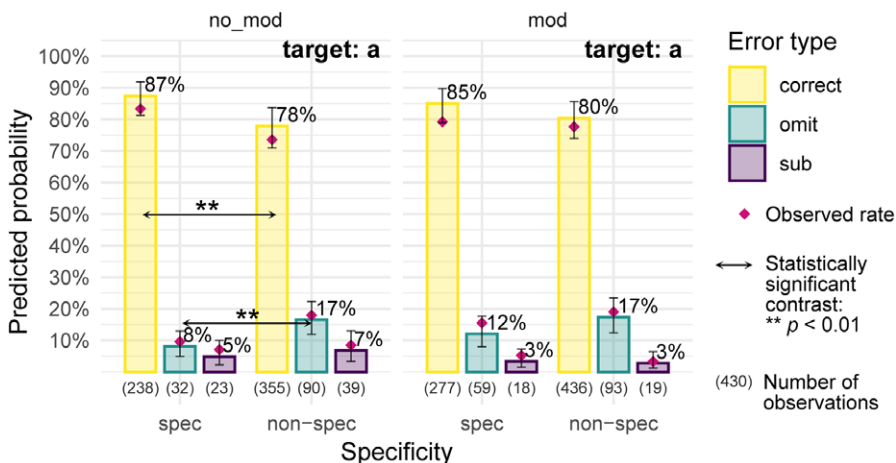


Figure 11. The effect of specificity in interaction with modifier on predicted probabilities of error types in count singular indefinites.

Table 8. Final model formula for error-type model for indefinite mass nouns

Component	Formula
Fixed effects	Error type ~ NL + level + specificity + modifier + abstractness
Random effects	+ (1 wr_ID)

Table 9. Error-type model for indefinite mass nouns results

Predictors	Accuracy rate		
	Log odds	SE	95% CI
over_a~Intercept (grand mean)	−2.57***	0.22	−3.01 −2.13
over_the~Intercept (grand mean)	−2.67***	0.23	−3.12 −2.21
over_a~NL:German	−0.46	0.41	−1.26 0.34
over_the~NL:German	0.37	0.33	−0.28 1.01
over_a~NL:Brazilian	0.42	0.33	−0.22 1.07
over_the~NL:Brazilian	−0.06	0.34	−0.72 0.60
over_a~NL:Chinese	0.15	0.36	−0.56 0.86
over_the~NL:Chinese	0.08	0.35	−0.60 0.76
over_a~mod:no_modifier	−0.77***	0.16	−1.09 −0.45
over_the~mod:no_modifier	−0.36*	0.17	−0.69 −0.04
over_a~spec:specific	0.52**	0.18	0.16 0.88
over_the~spec:specific	0.02	0.21	−0.39 0.43
over_a~abstr:concrete	−0.08	0.19	−0.46 0.29
over_the~abstr:concrete	−0.06	0.18	−0.41 0.30
over_a~level	0.07	0.22	−0.36 0.50
over_a~level	0.15	0.21	−0.26 0.57
Random effects by wr_ID			
SD over_a~(Intercept)	0.55		
SD over_the~(Intercept)	0.44		
N _{wr_id}	371		
Observations	878		
Pseudo-R ² (Nagelkerke, 1991)	0.171		
LRT vs. null model (only random effects)	$D(14)=70.273, p=1.7 \times 10^{-9}$		
VIFs	<3.5		

Note. over_a~ estimates for overuse of “a” errors vs. correct; over_the~ overuse of “the” errors vs. correct. Reference level: “Russian” for NL.

* $p < .05$; ** $p < .01$; *** $p < .001$.

Mass indefinites¹⁹

The model (Table 8) predicts 82% probability for correct \emptyset use, 11% “a” overuse, 7% “the” overuse (averaged across other variables). This is slightly lower than the predicted accuracy rate from the accuracy model, which was above 85% (full results in Table 9).

The model confirms that the only two significant variables in this context are specificity and modifier presence (Figure 12). Both effects are driven by the rate of “a” overuse: learners tend to overuse “a” more often before specific (25) and before premodified nominals (32).

(32) So I become fit, get fresh air and see \emptyset beautiful nature.
(L1-German, B2, ID:1087916)

¹⁹Fitting the model on both mass and plural indefinites produced negative pseudo-R² values (Nagelkerke, 1991). The same problem occurred when fitting a separate model on plurals. Therefore, we fitted the model on mass nouns only ($n = 878$).

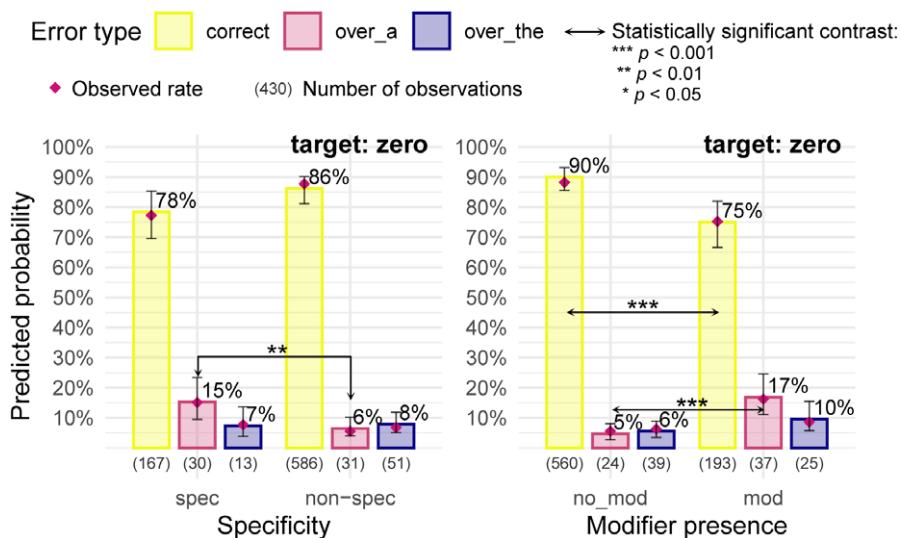


Figure 12. The effects of specificity (left) and modifier presence (right) on predicted probabilities of error types in mass indefinites.

Discussion

Summary of findings

Using manually coded learner corpus data and statistical modeling, we have revealed that the main factors affecting article accuracy are involved in several complex interactions—that is, their effects are not uniform across contexts. The most important findings are the following:

1. Clear L1 effect: article accuracy is generally higher in [+art] than in [-art] learners, although some NL-groups do not perform as expected (notably, L1-Chinese), and the effect varies across contexts (definite/indefinite, singular/mass/plural).
2. Specificity effect (defined as reference to a certain existing entity): only affects indefinite contexts (target “a”/∅), where “a” is generally more likely to appear with specific referents.
3. Prenominal modifier effect: distinct in definites versus indefinites—namely, modifiers increase “the” omission with definite singulars but increase “a” overuse with indefinite mass nominals (target ∅). There is, however, no modifier effect in indefinite singulars (target “a”).
4. Syntactic position effect: higher accuracy in existential and predicate as opposed to subject and object positions.

We will discuss each finding in more detail, combining the intricately related Findings 2 and 3 into one subsection.

NL in interaction with proficiency, definiteness, number, and countability

The higher accuracy of [+art] learners in our study largely confirms previous findings (Ionin et al., 2008; Murakami & Alexopoulou, 2016a; Snape, 2008). As expected

(Lardi re, 2004), all learners are more accurate in using “the” than “a,” except L1-Germans, who are at ceiling for both. However, a third of all nominals used by learners require \emptyset (these are mass and plural indefinites). Considering that omission is the most common error, this explains ~90% accuracy in \emptyset contexts across all NLs, including even L1-Russians (in stark contrast with their 55% accuracy in mass definites).

What requires additional explanation is the results of L1-Chinese learners. First, their accuracy seems rather high (~80%), although studies involving this population have demonstrated accuracy rates > 70% on article gap-fill tasks (Snape, 2009; Ting, 2005; Trenkic, 2008; Zhao & MacWhinney, 2018).²⁰ As for comparable production data, two studies using a picture-story task showed relatively high article suppliance rates in count singulars: overall 70% in 13 out of 15 lower to higher intermediate L1-Chinese learners (Goad & White, 2008), 98% for definites, and 89% for indefinites (only non-premodified) in 15 upper-intermediate learners (Snape, 2009). Considering that our data comprises writings produced offline as homework, the higher accuracy rates are not unexpected.

What is surprising is that [-art] L1-Chinese learners pattern with [+art] L1-Brazilians and not with [-art] L1-Russians. There are proposals that Mandarin Chinese is developing a definite article, which may even assume a functional projection (Cheng et al., 2017; Huang, 1999) as well as an indefinite article (Chen, 2003). Cultural differences may also be at play—for example, L1-Chinese learners may be more performance driven than L1-Russians. Finally, it is unclear why the accuracy of L1-Chinese learners decreases with proficiency (unlike the other NLs), particularly with count singular indefinites (target “a”). More data from higher level learners might clarify whether this is true and significant decline or part of a fluctuating or U-shaped curve.

Specificity and modifier in interaction with definiteness, number, and countability

Definites

As noted in Findings 2 and 3, definites are not affected by specificity, but there is increased “the” omission with premodified count singulars—that is, “the” is more likely omitted in “the advertising company” than in “the company.” Trenkic (2007), based on similar findings from L1-Serbians, suggests the article is structurally an adjective for learners, making it optional. So, when a modifier has already sufficiently narrowed down the range of potential referents, an article may be redundant. As we can see in our own data, in many cases the modifier leaves only one plausible referent option—for example, “the departure lounge of Oslo airport,” “the following recipe,” “the top score.” We could adopt Trenkic’s cognitive explanation, which suggests increased omission of redundant elements when cognitive resources are limited. However, we still need to explain why in our data this redundancy effect is found in definites but not in indefinites, which we will address in the following subsection.

Two findings remain unclear. First, definite mass and plural contexts are unaffected by modifier presence: “the” omission is not increased before premodified mass/plurals—for example, “the red wine”/“the new shoes.” Second, L1-Russians have considerably lower accuracy in mass and plural contexts (predicted 55% and 66%)

²⁰The only exception is the 63% accuracy rate in specific definite contexts with explicit denial of speaker familiarity with the referent in Trenkic (2008).

than in count singulars (82%). Austin et al. (2015) also observe a higher “the” omission rate with plurals than with singulars in 20 intermediate L1-Thai learners (prompted story recall). They attribute this to L1–L2 structural competition, which predicts that cognitively more demanding contexts, such as those requiring the suppliance of multiple functional morphemes (e.g., “the” and plural “-s” in plural definites), impede the suppression of competing L1 forms (i.e., bare plural definites). However, this only explains the higher omission in definite plurals but not in mass nominals. We cannot fully explain these patterns, which might also be rather uncertain due to lower numbers in these contexts (256 mass, 335 plural) and larger standard errors.

Count singular indefinites

The first question is how and why count singulars (target “a”) are significantly affected by specificity. Essentially, “a” is more consistently supplied with specific referents but more often omitted with nonspecific ones. We claim learners may associate “a” with the function of introducing a certain existing referent (i.e., specific, by our definition) into the discourse. In contrast, in nonspecific contexts, where “a” is not introducing an existing referent (as there is none), the semantic contribution of “a” may be unclear to learners.

However, the effect is only significant for L1-Brazilians and L1-Chinese (Figure 10; L1-Germans are at ceiling, and L1-Russians’ predicted accuracy rate is at 73% regardless of specificity). We suggest L1-Brazilians may draw on the indefinite article in their L1, especially in specific contexts. This could be because the use of bare singulars in argument position (which is allowed in Brazilian Portuguese) is more restricted in specific contexts²¹ (Ferreira, 2021). We could also argue that L1-Chinese learners benefit from an emerging indefinite article in their L1 (numeral “yi” meaning “one”), which is also more common in specific than in nonspecific contexts (Chen, 2003, pp. 1159–1160). In this case, L1-Russians are the only ones with nothing to rely on in their L1.

The second question is why there is no modifier effect in count singular indefinites (Finding 3). The only slight influence of modifier presence is that the effect of specificity described above becomes nonsignificant in premodified contexts. To illustrate, consider the non-premodified example (33), where less omission is predicted because the speaker has a specific question in mind (as opposed to a context where “a question” does not refer to any specific question). When a nominal is premodified, however (34), this specificity effect becomes statistically nonsignificant. Nevertheless, the pattern is in the same direction (Figure 11, right), so the tendency is similar, albeit smaller.

(33) So I make *a question*. Why the people needs to put their lives in risk?
(L1-Brazilian, B2, ID:780256)

(34) So I decided to take *a telemarketing course*.
(L1-Brazilian, B2, ID:650393)

²¹For example, bare singulars are ungrammatical in subject position of episodic predicates, where the referent is often specific, e.g., “*Cachorro está latindo”—“A dog is barking.” Moreover, bare singulars in episodic sentences can be interpreted as number neutral rather than necessarily singular, e.g., “Maria comeu maçã”—“Maria ate (an/some) apple” (examples from Ferreira, 2021).

Returning to the question, we need to explain why the modifier does not appear to make “a” redundant in the same fashion as it can make “the” redundant. If we assume, as suggested above, that learners associate “a” with the function of introducing a specific (existing) referent, we have to admit that a modifier cannot fulfil this function. There is also no evidence that learners use “a” to signal referent identifiability, which is the function of “the,” as there are few substitution errors. Therefore, although a modifier can narrow the range of possible referents, it may still only indicate a type—for example, “We are seeking an experienced analyst” as opposed to “any analyst”—if we accept that learners do not consider “analyst” identifiable in the first place.²²

This is unlike the findings in Trenkic (2007), whose [-art] L1-Serbian participants tended to omit both “the” and “a” with premodified nominals. The discrepancy is partly explained by the different task types. Trenkic used an oral information gap task (map completion) and a written task asking participants to translate as many stories as they could within the time limit, ensuring less reliance on metalinguistic knowledge. These online tasks revealed higher omission rates than the tasks in our corpus, which were untimed and unsupervised. Nevertheless, in Trenkic’s written task, the modifier effect was overall more pronounced in definite than in indefinite contexts, which is in the same direction with our pattern of a significant (but smaller) effect in definites and no significant effect in indefinites.

Note also that our results clearly differ from those in Ionin et al. (2004) and replications, as we detect specificity effects in both [+art] and [-art] learners and we observe few substitution errors. This is partly due to the differences in defining specificity (see fn3) and partly due to the different types of data.

Mass indefinites

In mass indefinites (target \emptyset), learners overuse “a” more often both with specific referents and with premodified nominals. We argue that this is consistent with our explanation for count singular indefinites above. If learners use “a” to introduce a certain existing referent, they would not use “a” with most mass nouns, which typically denote unbounded or vaguely defined entities. However, when a mass noun is used to refer to something specific, it will often refer to a portion or an instance of the entity, and learners might be using “a” to indicate this (35–38). A prenominal modifier can additionally specify a subclass or a type of entity, which is arguably more likely to occur when a specific portion or instance is referred to.²³

- (35) When we seat the server brought us $\# [\emptyset]$ *corn soup* to start.
(L1-Brazilian, B1, ID:1082859)
- (36) But one day,ⁱ did $\# [\emptyset]$ *great help* to her.
(L1-Chinese, B1, ID:1038505)
- (37) I ordered waffles with $\# [\emptyset]$ *whipped cream*.
(L1-Russian, B1, ID:157548)

²²We could also speculate that definite contexts are cognitively more demanding than indefinite ones due to the need to keep track of the discourse, but this would need to be confirmed in an online experiment.

²³In fact, > 40% of specific mass indefinites have prenominal modifiers in our data, whereas of nonspecific mass indefinites, only ~ 25% are premodified.

- (38) Also he said I should stay at home and take α [\emptyset] *medicine*.
(L1-Russian, A2, ID:534142)

Syntactic position

The higher article accuracy in existentials and predicates as opposed to subject and object positions is broadly in line with the literature (Hua & Lee, 2005, only for nonspecific contexts). One possible explanation is that the discourse and semantic properties of existential and predicate constructions are almost fixed regardless of the noun inserted: stating existence and denoting properties, respectively. They are also explicitly taught early on and may first be learned as formulaic sequences—for example, “There is a book on the table,” “I am a student.”

In contrast, the use of nominals in subject and object positions is much more varied, so it is difficult for learners to infer any patterns of distribution, as in the case of existentials and predicates.

Conclusion

Our study has revealed previously unnoticed and complex interactions between specificity, modifier presence, definiteness, number/countability, and L1 (NL). Overall, our data points to a semantic interpretation of articles by learners (except L1-Germans), broadly in line with Trenkic (2007). We conclude that learners associate “the” with definiteness (in the sense of an identifiable discourse referent) and “a” with introducing a specific (i.e., existing and certain) referent into the discourse.

The practical implications for learning and teaching are mainly around focusing learners’ attention on the structural features of the indefinite article (number and countability) rather than semantic features (specificity). Crucially, learners’ ability to use “a” correctly may depend on their understanding of countability in English.

An important limitation of this study is that the writings in the corpus were completed offline as homework, which implies preparation and the possibility to edit responses. As a result, the observed accuracy rates are probably overestimated and could be considerably lower in spontaneous, timed, or unprepared (written or oral) production.

Further research would benefit from extending this analysis to other [-art] L1s to ensure the observed patterns are not specific to L1-Russians (as L1-Chinese behaved similarly to [+art] groups). Larger sample sizes would improve the ability to detect effects of such cumbersome multilevel variables as syntactic position, especially in contexts where error rates are already low (e.g., plural definites or mass/plural indefinites).

Supplementary material. The supplementary material for this article can be found at <http://doi.org/10.1017/S0272263123000463>.

Data Availability Statement. The experiment in this article earned Open Data badge for transparent practices. The data is available at <https://www.iris-database.org/details/42UaQ-jWDMt>.

Competing interest. The authors declare none.

References

- Alexopoulou, T., Geertzen, J., Korhonen, A., & Meurers, D. (2015). Exploring big educational learner corpora for SLA research: Perspectives on relative clauses. *International Journal of Learner Corpus Research*, 1, 96–129.

- Alexopoulou, T., Michel, M., Murakami, A., & Meurers, D. (2017). Task effects on linguistic complexity and accuracy: A large-scale learner corpus analysis employing natural language processing techniques. *Language Learning*, 67, 180–208.
- Austin, G., Pongpaiboj, N., & Trenkic, D. (2015). Structural competition in second language production: Towards a constraint-satisfaction model. *Language Learning*, 65, 689–722.
- Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). *Parsimonious mixed models*. Arxiv. <https://arxiv.org/pdf/1506.04967v1.pdf>
- Bates, D., Mächler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67, 1–48.
- Bickerton, D. (1981). *Roots of language*. Karoma.
- Bolker, B. (2022). *GLMM FAQ*. <http://bbolker.github.io/mixedmodels-misc/glmmFAQ.html>
- Butler, Y. G. (2002). Second language learners' theories on the use of English articles. *Studies in Second Language Acquisition*, 24, 451–480.
- Chen, P. (2003). Indefinite determiner introducing definite referent: A special use of “yi ‘one’ + classifier” in Chinese. *Lingua*, 113, 1169–1184.
- Cheng, L. L.-S., Heycock, C., & Zamparelli, R. (2017). Two levels for definiteness. In M. Y. Erlewine (Ed.), *Proceedings of GLOW in Asia XI* (Vol. 1, pp. 79–93). MIT Working Papers in Linguistics.
- DeKeyser, R. M. (2005). What makes learning second-language grammar difficult? A review of issues. *Language Learning*, 55, 1–25.
- Dryer, M. S. (2013). Definite articles. In M. S. Dryer & M. Haspelmath (Eds.), *The world atlas of language structures online* (chap. 37). Max Planck Institute for Evolutionary Anthropology. <http://wals.info/chapter/37>
- Elff, M. (2021). *mclogit: Multinomial logit models, with or without random effects or overdispersion* [Computer software] (0.8.7.3). R Foundation for Statistical Computing.
- Ferreira, M. (2021). Bare nominals in Brazilian Portuguese. In P. Cabredo Hofherr & J. Doetjes (Eds.), *The Oxford handbook of grammatical number* (pp. 497–521). Oxford University Press.
- Frege, G. (1960). On sense and reference. In P. Geach & M. Black (Eds.) & M. Black (Trans.), *Translations from the philosophical writings of Gottlob Frege* (pp. 56–78). Blackwell.
- García Mayo, M. del P. (2009). Article choice in L2 English by Spanish speakers: Evidence for full transfer. In M. del P. García Mayo & R. Hawkins (Eds.), *Second language acquisition of articles: Empirical findings and theoretical implications* (pp. 13–35). John Benjamins Publishing.
- Geertzen, J., Alexopoulou, T., & Korhonen, A. (2013). Automatic linguistic annotation of large scale L2 databases: The EF-Cambridge Open Language Database (EFCamDat). In R. T. Miller, K. I. Martin, C. M. Eddington, A. Henery, N. M. Miguel, A. M. Tseng, A. Tuninetti, & D. Walter (Eds.), *Selected proceedings from the 31st Second Language Research Forum (240–254)*. *Cascadilla Proceedings Project*.
- Goad, H., & White, L. (2008). Prosodic structure and the representation of L2 functional morphology: A nativist approach. *Lingua*, 118, 577–594.
- Hawkins, J. A. (1978). *Definiteness and indefiniteness*. Croom Helm.
- Hawkins, J. A. (1991). On (in)definite articles: Implicatures and (un)grammaticality prediction. *Journal of Linguistics*, 27, 405–442.
- Hawkins, R., Al-Eid, S., Almahboob, I., Athanasopoulos, P., Chaengchenkit, R., Hu, J., Rezaei, M., Jaensch, C., Jeon, Y., Jiang, A., Leung, Y. L., Matsunaga, K., Ortega, M., Sarko, G., Snape, N., & Velasco-Zárate, K. (2006). Accounting for English article interpretation by L2 speakers. *EUROSLA Yearbook*, 6, 7–25.
- Heim, I. (2019). Definiteness and indefiniteness. In P. Portner, K. Heusinger, & C. Maienborn (Eds.), *Semantics: Noun phrases and verb phrases* (pp. 33–69). De Gruyter Mouton.
- Hua, D., & Lee, T. H. (2005). Chinese ESL learners' understanding of the English count-mass distinction. In L. Dekydtspotter, R. A. Sprouse, & A. Liljestrang (Eds.), *Proceedings of the 7th Generative Approaches to Second Language Acquisition Conference (138–149)*. *Cascadilla Proceedings Project*.
- Huang, S. (1999). The emergence of a grammatical category definite article in spoken Chinese. *Journal of Pragmatics*, 31, 77–94.
- Huebner, T. (1983). *A longitudinal analysis of 'the' acquisition of English*. Karoma.
- Huebner, T. (1985). System and variability in interlanguage syntax. *Language Learning*, 35, 141–163.
- Ionin, T., Ko, H., & Wexler, K. (2004). Article semantics in L2 acquisition: The role of specificity. *Language Acquisition*, 12, 3–69.

- Ionin, T., Zubizarreta, M. L., & Maldonado, S. B. (2008). Sources of linguistic knowledge in the second language acquisition of English articles. *Lingua*, 118, 554–576.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning*. Springer.
- Lardiere, D. (2009). Some thoughts on the contrastive analysis of features in second language acquisition. *Second Language Research*, 25, 173–227.
- Lardiere, D. (2004). Knowledge of definiteness despite variable article omission. In A. Brugos, L. Micciulla, & C. E. Smith (Eds.), *BUCLD 28 proceedings* (pp. 328–339). Cascadia Press.
- Leech, G. N., Rayson, P., & Wilson, A. (2001). *Word frequencies in written and spoken English: Based on the British National Corpus*. Longman.
- Lenth, R. V. (2022). *emmeans: Estimated marginal means, aka least-squares means* [Computer software] (R package version 1.7.4-1). R Foundation for Statistical Computing.
- Li, C. N., & Thompson, S. A. (1981). *Mandarin Chinese: A functional reference grammar*. University of California Press.
- Liu, D., & Gleason, J. L. (2002). Acquisition of the article “the” by nonnative speakers of English: An analysis of four nongeneric uses. *Studies in Second Language Acquisition*, 24, 1–26.
- Lüdtke, D., Ben-Shachar, M. S., Patil, I., Waggoner, P., & Makowski, D. (2021). performance: An R package for assessment, comparison and testing of statistical models. *Journal of Open Source Software*, 6, 31–39.
- Lyons, C. (1999). *Definiteness*. Cambridge University Press.
- Lyons, C. (2009). Defining definiteness. In C. Lyons, *Definiteness* (253–281). Cambridge University Press.
- Master, P. (1987). *A cross-linguistic interlanguage analysis of the acquisition of the English article system*. UCLA.
- Master, P. (1997). The English article system: Acquisition, function, and pedagogy. *System*, 25, 215–232.
- McHugh, M. L. (2012). Interrater reliability: The kappa statistic. *Biochemia Medica*, 22, 276–282.
- Michel, M., Murakami, A., Alexopoulou, T., & Meurers, D. (2019). Effects of task type on morphosyntactic complexity across proficiency: Evidence from a large learner corpus of A1 to C2 writings. *Instructed Second Language Acquisition*, 3, 124–152.
- Murakami, A. (2013). *Individual variation and the role of L1 in the L2 development of English grammatical morphemes: Insights from learner corpora*. University of Cambridge.
- Murakami, A., & Alexopoulou, T. (2016a). L1 influence on the acquisition order of English grammatical morphemes: A learner corpus study. *Studies in Second Language Acquisition*, 38, 365–401.
- Murakami, A., & Alexopoulou, T. (2016b). Longitudinal L2 development of the English article in individual learners. In S. Denison, M. Mack, Y. Xu, & B. C. Armstrong, (Eds.), *Proceedings of the 38th Annual Meeting of the Cognitive Science Society (1050–1055)*. Cognitive Science Society.
- Myles, F. (2012). Complexity, accuracy and fluency: The role played by formulaic sequences in early interlanguage development. In A. Housen, F. Kuiken, & I. Vedder (Eds.), *Dimensions of L2 performance and proficiency: complexity, accuracy and fluency in SLA* (pp. 71–94). John Benjamins Publishing.
- Nagelkerke, N. J. (1991). A note on a general definition of the coefficient of determination. *Biometrika*, 78, 691–692.
- R Core Team. (2021). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Reid, J., Battaglia, P., Schuldt, M., Narita, E., Mochizuki, M., & Snape, N. (2006). *The article choice of learners of English as a second language*. University of Essex.
- Robertson, D. (2000). Variability in the use of the English article system by Chinese learners of English. *Second Language Research*, 16, 135–172.
- Russell, B. (1905). On denoting. *Mind*, 14, 479–493.
- Scott, G. G., Keitel, A., Becirspahic, M., Yao, B., & Sereno, S. C. (2019). The Glasgow Norms: Ratings of 5,500 words on nine scales. *Behavior Research Methods*, 51, 1258–1270.
- Shatz, I. (2020). Lessons from creating a new corpus from an existing large-scale English learner language database. *International Journal of Learner Corpus Research*, 6, 220–236.
- Snape, N. (2008). Resetting the nominal mapping parameter in L2 English: Definite article use and the count-mass distinction. *Bilingualism: Language and Cognition*, 11, 63–79.
- Snape, N. (2009). Exploring Mandarin Chinese speakers’ L2 article use. In N. Snape, Y. I. Leung, & M. Sharwood Smith (Eds.), *Representational deficits in SLA: Studies in honor of Roger Hawkins* (pp. 27–52). John Benjamins Publishing.

- Snape, N., Leung, Y. I., & Ting, H.-C. (2006). Comparing Chinese, Japanese and Spanish speakers in L2 English article acquisition: Evidence against the fluctuation hypothesis? In M. G. O'Brien, C. Shea, and J. Archibald, (Eds.), *Proceedings of the 8th Generative Approaches to Second Language Acquisition Conference (GASLA 2006)* (132–139). Cascadilla Proceedings Project.
- Thomas, M. (1989). The acquisition of English articles by first- and second-language learners. *Applied Psycholinguistics*, 10, 335–355.
- Ting, H.-C. (2005). *The acquisition of articles in L2 English by L1 Chinese and L1 Spanish speakers*. University of Essex.
- Trenkic, D. (2002). Form–meaning connections in the acquisition of English articles. *EUROSLA Yearbook*, 2, 115–133.
- Trenkic, D. (2007). Variability in second language article production: Beyond the representational deficit vs. processing constraints debate. *Second Language Research*, 23, 289–327.
- Trenkic, D. (2008). The representation of English articles in second language grammars: Determiners or adjectives? *Bilingualism: Language and Cognition*, 11, 1–18. <https://doi.org/10.1017/S1366728907003185>
- von Heusinger, K. (2019). Specificity. In P. Portner, K. Heusinger, and C. Maienborn (Eds.), *Semantics: Noun phrases and verb phrases* (pp. 70–111). De Gruyter Mouton.
- White, L. (2003). Fossilization in steady state L2 grammars: Persistent problems with inflectional morphology. *Bilingualism: Language and Cognition*, 6, 129–141.
- Zhao, H., & MacWhinney, B. (2018). The instructed learning of form–function mappings in the English article system. *Modern Language Journal*, 102, 99–119.

Cite this article: Derkach, K., & Alexopoulou, T. (2024). Definite and indefinite article accuracy in learner English: A multifactorial analysis. *Studies in Second Language Acquisition*, 46: 710–740. <https://doi.org/10.1017/S0272263123000463>