UTRECHT UNIVERSITY

LINGUISTICS RMA PROGRAM

MASTER'S THESIS

Learning Prepositions in a Second Language

A Corpus-Based Approach to Semantic Transfer

Author
Ben Bonfil
Sn: 5920612

Supervisor
dr. Bert Le Bruyn

December, 2018

Contents

Acknowledgments

I would like to thank my supervisor Bert Le Bruyn for his patient and insightful guidance, and for the help and direction which were indispensable in the completion of this project. Thank you for always being available for long discussions and many reviews.

Many thanks to Henriëtte de Swart and Martijn van der Klis for welcoming me to the Time in Translation project and for all of their support during my internship and corpus work.

To all of the great teachers in the linguistics Master's, thank you for sharing your knowledge and inspiration. And thanks to the Master's students for a lot of fun, silliness and seriousness too. It has been a pleasure to to be your classmate.

I am very grateful to the participants in the experiment who voluntarily dedicated the time and effort required. Many thanks go to Jos Tellings as well, who assumed the role of a second reader on a short notice.

Finally, thanks to my family and friends for being who they are, and to Neta for everything.

Chapter 1

Introduction

Learners of a second language (L2) often find some parts of the language they are learning to be easier to acquire than others, which remain difficult even after extensive exposure to the language. One line of work within generative linguistics traces the varying difficulty of the L2 back to the native language (L1) and the transfer of grammatical representations from the L1 to the L2. This is a result of the assumption that learners do not approach the task of acquiring a second language with a blank slate, but rather with a fully developed L1. Such work in second language acquisition (SLA) is often done on the basis of formal analyses of variation between the L1 and the L2, motivated by linguistic theory and focused on explaining where a learner's L1 grammar might promote or inhibit the acquisition of a target-like L2 grammar.

The focus on linguistic representations within the generative literature often leaves the role of language input in SLA as an afterthought. In the current work I show that this compromise is not a necessary one, by presenting a data-driven approach to the study of (semantic) L1 transfer on the basis of corpus data. Assuming the approach presented here, generative SLA theory has the potential to evolve by adopting elements from a usage-based perspective on L2 acquisition and L2 input. These ideas are operationalized in an investigation of L1 transfer within a narrow domain of language—the domain of locative prepositions. I focus on a specific domain that exhibits cross-linguistic variation within a small set of lexical items which encode spatial relations. I then propose how predictions about learning can be based on an analysis of cross-linguistic variation, which is presented here with Hebrew and English. Finally, the predictions are tested in a production experiment with native Hebrew speakers who acquired English as an L2. While the experimental data only support a subset of the predictions, I hope to convince the reader that further refinement of the methodology is warranted.

In the next chapter, I position the current work within SLA theory and present the motivation for the chosen domain of inquiry. In Chapter 3, I present a corpus study of cross-linguistic variation in the domain of locative prepositions. The study is based on data from a parallel corpus compiled off the Hebrew and English versions of the book *Harry Potter and the Philosopher's Stone* by J. K. Rowling. An analysis of cross-linguistic variation is presented in the form of a semantic map with a binary feature decomposition, which allows us to derive predictions about L1 transfer. Chapter 4 covers an experiment designed to test the predictions of feature-based learning as defined by the corpus study. Chapter 5 provides a reflection on the main findings of the current project and suggestions for future work.

Chapter 2

Theoretical Background

The question of how a learner's L1 influences the process, progression and outcomes of L2 acquisition is motivated by both theoretical considerations and empirical observations. From a theoretical point of view, researchers approach the characterization of how L2 learners accommodate grammatical representations of a new language, given a fully developed L1 grammar. From an empirical perspective, evidence points at an effect of the learner's L1 on L2 performance, both during the initial stages of acquisition and with respect to ultimate attainment.

Within generative linguistics, approaches to language acquisition draw on linguistic theory in order to identify the mental representations that learners must possess if they are to acquire a second language (Lardiere, 2012). Earlier work in SLA drew upon the Principles and Parameters framework (Chomsky, 1981) to identify and explain constraints on second language from a Universal Grammar perspective. Yet because linguistic representations are at the cornerstone of any generative work in SLA, and given that generative linguistics later underwent a paradigm shift towards Minimalism (Chomsky, 1995) a renewed approach in SLA theory was called for. Lardiere (2009) captured some of the consequences of the minimalist shift, while discussing the limitations of the P&P framework in SLA, and proposed the Feature Reassembly hypothesis as a theoretical basis for the L2 acquisition of morpho-syntax. The detailed hypothesis, which serves to provide the main line of inquiry in the current thesis, is presented in the next section, together with several related studies. This is followed by a discussion about the role of L2 input in feature reassembly, a brief review of previous work on the semantics and variation of locatives, and the proposal put forward in the current study.

2.1 The Feature Reassembly Hypothesis

Lardiere's Feature Reassembly hypothesis (Lardiere, 2009) builds upon the Minimalist idea that the locus of cross-linguistic variation lies in the lexicon, in the configurations of features that make up lexical items. Within that theoretical framework, the L2 learning task is formulated in terms of acquiring the target-like (syntactic/semantic) feature configurations of the lexical items that make up the L2. *Lexical items* here assume the minimalist definition, and can be content words, morphemes or functional grams—anything that *"enter[s] into computations that derive hierarchically-structured representations for the pairing of meanings with forms"* (Lardiere, 2009).

According to the theory of feature reassembly, learners first map lexical items from their L1 onto L2 items based on perceived similarity in meaning or grammatical function. This process is referred to as an *initial mapping* and is further discussed below. With an initial mapping in place, learners have to gradually adjust the feature configurations carried over from their L1, based on evidence in the L2 input. In the process of reconfiguring feature bundles, learners may either *select* features from a UG-sanctioned universal set of features—in the case of features that are not productive in the L1—or *reassemble* features that do exist in the L1, but are bundled in different functional categories or across different parts of the grammar.

Consider the following example of how feature reassembly might take place, taken from Cho and Slabakova (2017): it is suggested that the English past morpheme *-ed* corresponds to a feature configuration of [+past±plural]. In Russian, past tense is marked by the morpheme *-l* which corresponds to the feature configuration [+past]. However, the Russian marker *-l* appears in combination with one of *∅, -a, -o, -i* in agreement with number/gender (masculine, feminine, neuter and plural without gender marking). Given that, an English speaker learning Russian is hypothesized to establish a mapping between the English *-ed* to one (or more) of the Russian morphemes. However, in order to arrive at a target-like configuration, a learner would have to detect the meaning contrasts between the different Russian past morphemes, which correspond to grammatical number and gender, and they would have to exclude [±plural] from the initial configuration carried over from English.

This example is one where learners have to assemble representations based on interpretable semantic features. While earlier work suggests that L2 learners fail to acquire uninterpretable features that are not present in their L1 (cf. Hawkins (2005) on deficit representations), Lardiere proposes that formal contrasts in the input are always detectable to the learner, and thus should lead the learner to construct some kind of representation based on the detected differences. It is further suggested that the necessary building blocks (features) for the target L2 representations might be available in different items in the L1. Going back to the example of an English speaker learning the Russian past marker with gender agreement, the theory suggests that while English does not have agreement of the past morpheme with gender, features for grammatical gender are still available in English, in pronouns or possessives for example. This is taken to mean that learners should be able to assemble grammatical gender in past morphemes as well.

The Feature Reassembly Hypothesis (FRH), and the theoretical consequences it bears on the process of acquisition, allow researchers to generate predictions about relative difficulty in the L2. Dominguez et al. (2011) present one approach to difficulty by considering whether feature reassembly is required or not, to arrive at target-like use of tense in L2 Spanish given an L1 English tense system as a starting point. Under such an approach, the FRH presents a possible three-way distinction between feature selection (when required features are not found in the L1 lexicon), reassembly (when features exist in the L1 but are configured differently) or overlap (when the L1 and L2 feature compositions match). Cho and Slabakova (2014) further expand on the implications that the FRH bears on difficulty by taking into account overt vs. covert expression of semantic contrasts. The authors argue that lexical items dedicated to a single grammatical meaning are easier to acquire than items with multiple grammatical meanings. For example, a demonstrative provides a definiteness function, but it is not its main function, which makes it harder to acquire than a dedicated definite article.

To summarize, the main proposal that the FRH affords with respect to difficulty in learning is based on identifying meaning contrasts and the reconfiguration of features required in order to arrive at target-like lexical items, given the grammatical (lexical) knowledge that learners carry over from their L1. On the opposite end of the difficulty scale, the feature reassembly model affords the prediction that learning should be easier in places where the L1 and L2 bundle features similarly (positive transfer), as there is less reassembly work necessary and thus less linguistic evidence required to successfully do so. Either way, the FRH sketches out a prominent role for the L1 in L2 acquisition. We find support for this kind of predictions in empirical results on L1 transfer in definiteness (Ionin and Montrul, 2010; Cho and Slabakova, 2014), tense and aspect (Dominguez et al., 2011), and specificity markers (Cho and Slabakova, 2017).

2.2 The Role of L2 Input

As presented in the previous section, the FRH allows the study of L1 transfer in terms of form-meaning contrasts and the underlying grammatical features. In addition to these factors which are motivated by theoretical linguistics and meaning representations, and come from an analysis of the native and target grammars, research in SLA is also concerned with studying the L2 input that is available to learners.

The focus on L2 input is traditionally found outside of the generative paradigm to SLA, with the study of the distributional characteristics of L2 input being central in research carried under constructionist and usage-based approaches to SLA. Work in those frameworks builds on theories of general human learning, and presents L2 learners' developmental trajectory as guided by properties such as the frequency, contingency and prototypicality of constructions in the L2 input (Wulff et al., 2009). Other work draws on extra-linguistic factors such as learners' attention, in what makes L2 input available as linguistic intake (Ellis, 2006).

The generative side of SLA has recently seen more direct calls towards incorporating the study of L2 input. This has been clearly formulated in a review of the field of generative SLA by Rothman and Slabakova (2017) who say that: *"A newer idea in generative theorizing is that L2 convergence crucially depends on how much evidence in the input there is and how clear such cues are in the input itself."* (p. 23).

The current work makes a step towards answering that call, by showing how a usage-based approach to acquisition, one centered around the distribution of form-meaning mappings in the input, can be made compatible with the generative approach of feature reassembly. As the FRH shifts the learnability burden onto lexical items, it makes sense to study the distribution of lexical items in the L2 input, in order to produce a measure of the evidence available for learning. It is easy to imagine an extreme case of a lexical item that is never seen in the input: a learner would have no evidence for it at all and would not acquire it. More realistically, for learning a lexical item that is prevalent in the input, a learner would probably require exposure to many contexts before they are able to acquire the full feature set associated with that lexical item.

The synthesis we propose here, between usage-based and feature-based learning, rests on the idea that meaning contrasts in the L2 input are identified by meaning contrasts in corresponding forms in the L1, based on the contexts in which they appear. Following Lardiere, we hypothesize that learners construct representations of meaning contrasts as they encounter them. We propose they do so by following the distribution of L1-L2 form mappings. That is to say, learners construct an L2 representation based on the usage of forms in their L1. This idea motivates our use of a translation-based corpus methodology (Chapter 3) in order to assess what learners might, or should, consider as parallel L1-L2 contexts.

In the next section, I set SLA theory aside in order to review some of the literature on the semantics of spatial reference in prepositional phrases, and to present the motivation for the choice of locative prepositions as the linguistic domain for a study of L1 transfer.

2.3 Spatial Semantics

As a crucial factor in human cognition, spatial concepts and spatial reference are prevalent in language. The relevance of spatial reference to work in semantics was nicely articulated by Feist (2000), who wrote that *"spatial relational terms provide an interesting domain for semantic* *research in part due to the fact that the relations being described could be said to exist objectively in the world, separate from linguistic conceptualization."* (p. 3). In the current project, we focus on one way in which languages relate to space, namely locative prepositional phrases and the prepositions that head them.

Following Talmy (1975), work in spatial semantics is characterized by studying possible relations between Figure and Ground: where a Figure is an object whose location is being discussed, and a Ground is an object whose location is used as a reference. In the context of the current study, we focus on locative prepositions: prepositions whose semantic content is combined with the location of a Ground, in the derivation of the Figure's location. Examples of English locatives are *in*, *on*, *at*, *above* and *near* among many others.

Approaches to the study of Figure-Ground relations may assume a theoretical perspective of semantic compositionality and geometric intention. Zwarts and Winter (2000) fit spatial considerations within a model-theoretical prepositional semantics that is based on the mathematical concept of vector-spaces. Within such a framework, the meaning of a PP has to do with the physical space occupied by figures and grounds in the world, modeled in terms of sets of points or regions.

Assuming such a geometrical approach, it is possible to identify geometry-based semantic features for prepositions. As an example that illustrates how this kind of formalization could work, one could propose that the preposition *above* contains a [+y] feature suggesting that a Figure's position lies at a higher vertical coordinate compared to a Ground (in a relevant coordinate system where y is the vertical axis).

A different approach to the semantics of spatial relations is one where geometric concerns are diminished in favor of studying how the Figure and Ground objects relate to each other in terms of their function. For example, a cup that is placed on a table is physically supported by the table—a fact that seems more relevant to our willingness to assert that the cup is *on* the table, rather than the cup being present at a higher vertical position compared to that of the table. If the cup were to float above the table, we would not say that it is on the table. We also find that spatial reference depends on properties such as the geometry and animacy of figures and grounds (Feist and Gentner, 2003). Zwarts (2017) offers a comprehensive review of the different approaches to spatial semantics and suggests that a synthesis of approaches (geometric, functional etc.) can account for a greater variety of uses of PPs, in a way that cannot be done by each approach on its own.

2.3.1 Cross-Linguistic Variation

Despite the fact that spatial meanings are grounded in properties of physical existence, Levinson and Meira (2003) found significant cross-linguistic variation in the domain of space across typologically distant languages, leading to the hypothesis that semantic categories of adpositions do not necessarily reflect universal conceptual categories. Whether language plays a role in the cognition of space is a relativistic question that will not be dealt with here. For our current purposes, it is enough to accept that different languages may categorize spatial relations in different manners.

There are many ways in which languages might differ with respect to categorization in the conceptual space of spatial relations. One kind of variation is in the grammatical categories used to refer to space, as presented by Slobin (1996) in the case of motion verbs. Variation in spatial reference can also arise from a difference in what speakers take into account in the conceptualization of a spatial relation. For example, when Korean speakers talk about containment relations, they choose different prepositions based on whether a Figure fits tightly or loosely in a Ground (Choi and Bowerman, 1991). This means that Korean brings into account a physical property of containment that is not found in the spatial semantics of other languages, and makes it explicit in the choice of prepositions.

A different kind of variation has to do with the many possible ways to divide the conceptual space of spatial relations, even when similar properties are considered in that space. In her work about L1 acquisition of spatial reference, Bowerman (1996) lists many ways in which languages vary with respect to the division of the spatial conceptual space. Bowerman goes as far as to claim that languages vary even in what relations they treat as spatial relations. The conclusion is that innate knowledge and cognition of spatial concepts is not enough for a child to acquire the spatial system of their native language, suggesting instead that children *"must work out the meanings of the forms by observing how they are distributed across contexts in fluent speech"* (p. 425). This argument cannot be applied as-is to L2 learners, as it is to be expected that adult learners possess more tools and knowledge that might allow them to perform high-level analysis and generalization in acquiring the semantic categories of the target spatial system. If, however, they utilize the knowledge from their L1, as we assume they do, the nature of variation in the spatial domain means we can expect them to run into problems.

2.3.2 Spatial Semantics in Second Language Acquisition

Within the field of SLA, spatial reference has been studied before for effects of cross-linguistic influence. Jarvis and Odlin (2000) found evidence of L1 transfer in the choice of locatives by L1 Finnish and L1 Swedish learners of English, elicited in a controlled production experiment. Their results support significant differences in the use of prepositions in the L2, between the two L1 groups; differences that the authors trace back to the ways in which each of the L1s differs from English.

More recent evidence demonstrating L1 transfer in the use of spatial prepositions is found by Alonso et al. (2016). The authors present a study comparing the performance of learners with an L1 background of either Danish or Spanish in describing spatial configurations using locative prepositions. Participants were tested in both a picture description task and a sentence completion task, in contexts where L1 English speakers use *in*, *on* and *at*. The results show that Danish learners are very close to native English speakers in their choice of prepositions, while Spanish learners vary to a greater extent. We expect to find similar effects of L1 transfer in the current study, the structure of which is explained in the next section.

2.4 The Current Study

Given the theoretical framework of the Feature Reassembly hypothesis, our aim is to study a linguistic system where a learner might be facing the task of learning a configuration of features in the L2 that is different from the one provided by their L1. We choose to look at locative prepositions in English and Hebrew.

While the FRH was introduced in the context of transfer phenomena at the syntax-semantics interface, we propose here an application of the FRH in lexical semantics. In doing so we do not intend to claim that a theory of acquisition at the interface should necessarily apply to the acquisition of lexical semantics. On the other hand, there is no a priori reason to treat lexical semantics differently under a theory that postulates contrastive learning of (interpretable) semantic features in L2 acquisition. It might also be the case that prepositions form a grammatical category that is both lexical and functional in the syntactic sense (Zwarts, 1997) and could participate in interface phenomena. However, in the current study we focus on the semantic/conceptual content of prepositions.

Based on the earlier results on L1 transfer in the prepositional domain by Jarvis and Odlin (2000) and Alonso et al. (2016), we expect to find similar effects among native speakers of Hebrew acquiring English. Both languages encode locative spatial relations by means of prepositions. Thus the learning task consists in reassembling semantic categories and not, for example, in reassembling semantic features between the category P and a grammatical case system.

The setup of the current study could potentially apply to other phenomena of cross-linguistic variation in syntax/semantics, as it is not specifically designed towards spatial reference. As detailed in Chapter 3, our methodology is data-driven and attempts to make as few assumptions as possible with respect to semantic theory. This sets the current work apart from previous studies that follow the FRH (Gil and Marsden, 2013; Cho and Slabakova, 2014, 2017), which rest on analyses of cross-linguistic variation that are theory-driven, and where the learning task is portrayed in terms of features which are independently motivated by work in theoretical linguistics.

We choose to study variation in the prepositional domain because prepositions are frequent in corpus data (and consequently in L2 input), they are simple to annotate and process using computational methods, and because an effect of L1 transfer has been attested in previous studies. Based on earlier results, we further assume that the chosen English prepositions form a class of semi-alternates, in the sense that learners who have difficulties in producing target-like prepositions would confuse or conflate members of that set.

2.4.1 Research Question

The goal of the present study is to look into the different factors that affect semantic transfer from the L1 when acquiring a second language. This goal is operationalized in the current study by looking into the extent to which native Hebrew speakers' use of prepositions in English can be modeled by variation in semantic features, as identified by cross-linguistic mapping of forms in corpora.

The structure of the current study is as follows: we begin with a contrastive analysis of locative prepositions based on data from a parallel corpus (Chapter 3). The analysis is presented in the form of a semantic map, together with a feature-based account of variation, and subsequent hypotheses about learner production. Finally, we test the predictions offered by the analysis in an empirical study using a fill-in-the-blank task (Chapter 4).

Chapter 3

Corpus Study

The first step of the current project consists in drafting a corpus-based analysis that captures the variation in the domain of spatial reference between English and Hebrew. The corpus data is compiled into a semantic map of variation, off which we infer a geometry of semantic features. This provides us with a formal proposal for feature-based learning required on behalf of learners, and allows us to trace the potential path of progression for feature (re)assembly.

Following van der Klis et al. (2017), the current project employs the method of *Translation Mining* in order to construct a semantic map of cross-linguistic variation. Translation mining relies on a combination of computational methods and manual annotation work to extract cross-linguistic form mappings out of parallel corpora. Annotation consists of three main phases: first, the desired constructions are identified in the source language version of the corpus, by selecting the relevant words or morphemes in the text. Then, translation equivalents of the annotated constructions are identified in the target language text of the corpus. Finally, each annotated construction in every language is assigned a label. These annotation steps are preferably performed manually by human annotators when the size of the corpus allows for it. The reason to prefer manual annotation is that our current computerized annotation software is not accurate enough, especially considering that in a relatively small corpus, annotation errors might lead to significant inaccuracies.

In choosing the prepositions to study, we have decided to look at a subset of what are considered as simple topological prepositions (Zwarts, 2017). As the study is concerned with L2 English, we base the selection of contexts on the subset of English topological prepositions consisting of *in*, *at* and *on*. We consider all occurrences of these prepositions in the source text where they are used in their spatial meanings, and we include all matching prepositions in the translated versions of the corpus.

Before moving on to present the analysis, it is important to make an explicit note of what is not covered in the current project. The first choice we had to make is in selecting a subset of prepositions to study, and basing the selection on English prepositions. This means that we inevitably deal with a constrained set of use cases and data. A more complete methodology to variation would include two-way mappings, going from the L1 to the L2 as well. We leave that to future research.

There is another caveat to applying translation mining with respect to contrastive analysis under the FRH framework. As discussed in the previous chapter, the FRH suggests that languages bundle features across different parts of grammar. In translation mining, we ask native speakers to identify the translation equivalents of forms from a source language in texts of their native language. This might be a problem if the annotator is not aware of the different ways in which cross-linguistic differences might surface in a given language pair. For example, suppose that one language uses a morphological suffix in a context where another language uses a combination of specific lexical items and a syntactic alternation such as clause word order. In such a case, an annotator might fail to acknowledge that both components of their native grammar are necessary to match the feature set of the source language. This is essentially the L2 learnability problem of the FRH, applied to L2 annotators in corpus work. We minimize the possibility for this kind of errors by two measures: our languages of choice use prepositions for spatial reference, in a seemingly similar way (as opposed to a case system, for example), and we recruit annotators who are trained in linguistics.

Another strong underlying assumption in our work is that the parallel texts of the corpus share the same meaning. This is not always true for literary translations. We avoid potential problems of that kind by explicitly asking annotators to verify that the translation of each context keeps the meaning of the original context, where relevant to the construction being studied.

An advantage of translation data is that compared to studying variation by means of picture description (e.g. Levinson and Meira (2003)), translation mining circumvents potential variation resulting from the different ways speakers might perceive a spatial configuration. A possible disadvantage of our approach is that spatial reference in the text of a novel has to do with an imagined reality, lacking the basic visual grounding that a picture description task allows for.

3.1 Data Collection

For the current project, we have annotated occurrences of prepositions in sentences from chapters 1, 10 and 11 of *Harry Potter and the Philosopher's Stone*. We have thus sourced 89 contexts based on the preposition *on*, 104 based on *in* and 87 contexts of *at*. The selection was further narrowed down to include only fragments where prepositions are used in their spatial meaning. Thus we have excluded temporal uses like (1a), but kept metaphorical uses that are based on spatial concepts, such as (1b) but not (1c).

(1) a. […] he left the building **at** five o'clock […]

b. "Really, Dumbledore, you think you can explain all this **in** a letter?"

c. Mr. Dursley arrived in the Grunnings car park, his mind back **on** drills.

We have also excluded contexts where a source preposition was translated by means such as nominalization or paraphrase, like (2). All of these criteria were applied to both the original contexts and the translated ones. The number of contexts left after exclusion, along with the parallel prepositions found in Hebrew is presented in Table 3.1.

(2)

dadli	*haya*	*asuk*	*belicroax*	*velehaif*	*et*	*aruxat-haboker*	*šelo*	*lekol*
Dudley	was	busy	b-to-scream	and-to-throw	ACC	breakfast	his	**l**-all

hakivunim.
directions

Dudley was now having a tantrum and throwing his cereal **at** the walls.

English		**Hebrew**[1]	
on	42	b-	60
at	42	al	35
in	34	l-	13
		el	7
		btox	3

Table 3.1: Counts of preposition contexts

If we consider the mappings between Hebrew and English we find the tuples in Table 3.2. In what follows we focus on the more frequent Hebrew forms found in the corpus: *b* and *al*. Narrowing down the set of forms that we consider leaves us with enough contexts per form to be used in the experimental part of the project (Chapter 4). It also means we lose potential sources of transfer, that might in fact be more prominent among learners given the lower frequency of contexts in the corpus. We leave those for future work using larger corpora.

	b-	*al*	*l-*	*el*	*btox*	
at	20	10	7	5		42
in	26		5		3	34
on	14	25	1	2		42
	60	35	13	7	3	

Table 3.2: English-Hebrew form mapping frequencies

[1] *b-* and *l-* are bound morphemes while *al* and *el* are free.

3.2 Semantic Mapping

One way to describe a conceptual space and how it is categorized within a certain language, or cross-linguistically, is by using semantic maps (Haspelmath, 2003; Zwarts, 2010). Underlying the idea of semantic mapping is a measure of meaning similarity transformed into geometrical distance. Semantic maps provide some indications as to how meanings within a certain semantic space compare to each other in terms of relative similarity. They show us which meanings are closer and which are farther apart.

In what follows, we produce a semantic map of the conceptual space of locatives following what Zwarts (2010) calls the 'matrix-driven' approach: a method in which a map is produced by compiling a matrix of form mappings (in our case, prepositions which occur in different contexts) and using that matrix to construct a graph of the conceptual space, in a way that every lexical item can be matched with a connected sub-graph.

For the purpose of preparing a semantic map, we set aside considerations that arise from theories of spatial semantics, and follow the distribution of forms. We begin by preparing a tuple for each context, which contains the forms used in each of the corpus languages. We then attempt to group contexts based on tuples. The underlying idea is that when any one language differentiates between two contexts by using different forms, there might exist a divide in the conceptual space that is not realized in the other languages of the corpus. Finally, tuples should be connected in the graph when they share a common form within a language.

We use a semantic map as a way to discuss the distribution and variation in meaning without making theoretical assumptions with respect to spatial semantics. The labels attached to each tuple serve only as a crude indication of the common meaning identified in the corpus examples for the category defined by the tuple. A short description of the kind of examples found in each category is presented in Table 3.3, and the resulting map is presented in Figure 3.1.

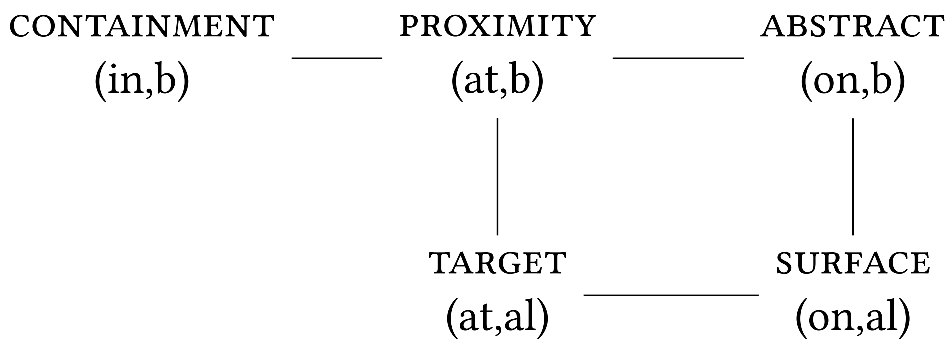

Figure 3.1: Semantic map, categories and form tuples

Label	Description
ABSTRACT	Ground is not strictly a physical object, but an abstraction.
CONTAINMENT	Ground is perceived as a container or a set that contains Figure.
SURFACE	Ground is a surface and Figure is located in a manner where contact is established.
TARGET	Directional configurations between Figure and Ground.
PROXIMITY+TARGET$_B$	Figure is in proximity to Ground; some directional configurations.

Table 3.3: Categories and labels

In (3) we list several glossed examples of the items found in the different categories: (3a - ABSTRACT, 3b - CONTAINMENT, 3c - SURFACE, 3d - TARGET, 3e - PROXIMITY)

(3) a. *mar darsli nahag tamid laševet kšegabo mufne el haxalon*
mr Dursley used always to-sit when-back-his turned to the-window
bemisrado šebakoma hatši'it.
b-office-his that-**b**-the-floor the-ninth.

Mr. Dursley always sat with his back to the window in his office **on** the ninth floor.

b. *hu haya asuk midai belexatet beglimato, mexapes mašehu*
he was busy too b-rummaging **b**-cloak-his searching something

He was busy rummaging **in** his cloak, looking for something.

c. *dambeldor hošit yad vetafax al ktefa*
Dumbledore reached hand and-patted **al** shoulder-her

Dumbledore reached out and patted her **on** the shoulder.

d. *pit'om anašim mikol xelkey hakahal hicbiu al hari*
suddenly people from-all parts-of the-audience pointed **al** Harry

Suddenly, people were pointing up **at** Harry all over the stands.

e. *[gam le-]ron, bašulxan hasamux, lo hitmazel hamazal.*
[too l-]Ron **b**-the-table the-next, NEG lucked.REFL the-luck

Ron, **at** the next table, wasn't having much more luck.

More examples of corpus items per category are found in Appendix A. Note that the category PROXIMITY+TARGET$_B$ combines contexts in which a figure is located in proximity to a ground, as in (3e) above, but also contexts exhibiting directional meanings. This is because directional meanings occur in Hebrew both with *b* (4a) and *al* (4b). For simplicity, we will label this category as PROXIMITY.

(4) a. *mar darsli micmec vebaha baxatula*
mr Dursley blinked and-stared **b**-the-cat

Mr. Dursley blinked and stared **at** the cat

b. *hari histakel al haricpa*
Harry looked **al** the-floor

Harry looked **at** the floor.

We now have an hypothesis about the layout of the conceptual space in the form of a semantic map. From here we move on to identify the underlying feature geometry that is responsible for this layout. The idea is to illustrate the geometry of variation in the conceptual space, and we avoid naming the semantic features responsible for its division. This goes back to the idea of learners having to assemble feature-based representations by means of contrastive analysis (Section 2.1). We start by highlighting the contingent sub-graph taken up by each preposition in Hebrew and English in Figure 3.2, where we illustrate the orthogonal relationship between the different divisions of the conceptual space. From here, feature decomposition should proceed in a way that would allow us to define each highlighted region using a well-formed set of features. The second constraint on decomposition is that no two categories can share the same feature set, and any two categories that are connected in the graph must differ in the value of exactly one feature. We follow these constraints and translate the map into categories based on three binary features (α, β and γ) which is the minimal number of binary features required to compose five categories. The resulting feature decomposition is presented in Figure 3.3 which follows the same layout as our map. The feature configurations that correspond to each preposition are listed in Table 3.4.

We move on to consider what the map can tell us about the possible routes of acquisition of the English prepositions by L1 Hebrew speakers, following the basic ideas of the FRH. We first discuss how feature reassembly might unfold with the identified configuration of features. This will be followed by a discussion of different ways to generate predictions based on the map, and a presentation of the predictions to be tested in our production experiment.

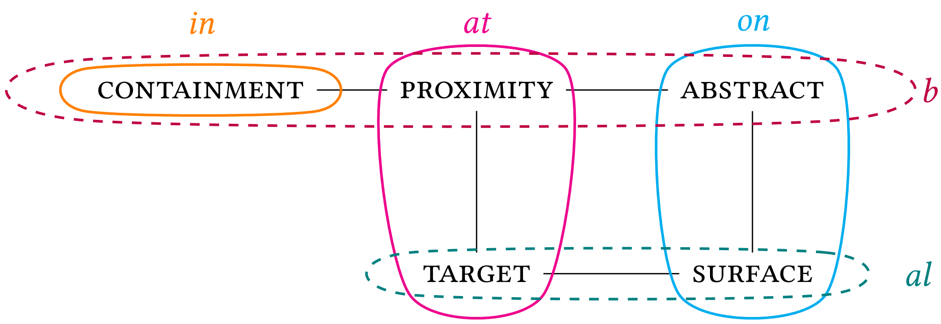

Figure 3.2: Semantic map with forms

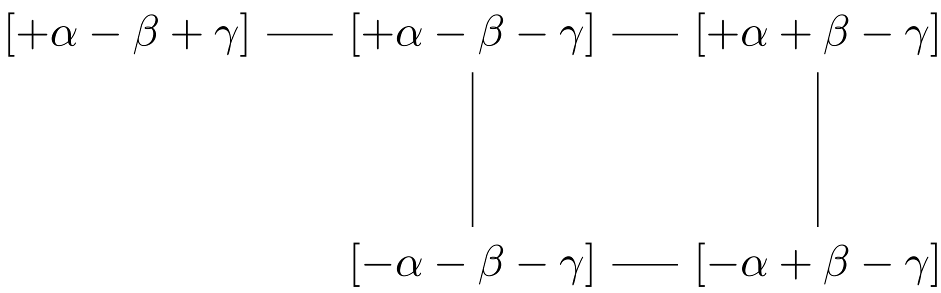

Figure 3.3: Feature decomposition of forms

b	$[+\alpha]$
al	$[-\alpha]$
on	$[+\beta]$
at	$[-\beta - \gamma]$
in	$[+\gamma]$

Table 3.4: Feature composition per form

3.3 Feature Reassembly

Since feature reassembly postulates an initial mapping between forms of the L1 and the L2, we will start by considering what such mapping might be. Lardiere (2009) proposed that learners base their initial mapping on perceived similarity in meaning, but also suggested that *"the acquisition researcher [...] will always be guessing – albeit hopefully making a professionally-informed best guess from among a smallish range of possibilities – as to which morpholexical correspondences between languages a learner is initially most likely to establish"* (p. 219). We propose several assumptions about what might affect learners in the initial mapping phase, and that allow us to make an informed guess about it.

First, we consider how frequently learners encounter an L2 form in contexts where they would use a certain L1 form in their native language. The assumption is that forms that co-occur more frequently, would be perceived as more similar. For assessing that measure, we will refer back to the frequencies of L1-L2 forms as found in our corpus (Table 3.2).

We propose that learners looking for probable candidates in English for mapping *b* would find *in* to be a good option. This is because *in* appears quite frequently in contexts where native speakers of Hebrew would use *b* (26/60, cf. Table 3.2) and because *in* never co-occurs with *al* in our data. We further assume that *al* would be mapped onto *on*, again based on frequent evidence in both directions (25/35 occurrences of *al*, 25/42 occurrences for *on*). Thus we have an informed guess as to how learners might initially map the L1 forms *b* and *al* onto the L2 forms *in* and *on*. It remains to be shown how the L2 representation of *at* is initiated. Before we look into that, we briefly go back to the theoretical discussion about initial mapping in order to present our assumptions more accurately.

In the original proposal of Lardiere, initial mapping is presented as a process in which forms from the L1 are carried over to the L2. That is to say, the direction of mapping is such that learners look for targets in the L2 onto which they map the existing representations of forms from their L1. In later work (Gil and Marsden, 2013; Cho and Slabakova, 2017), we find a slightly different presentation of the initial mapping concept, where learners are assumed to look in their L1 for lexical items that correspond to L2 forms they find in the L2 input. The difference in the direction of mapping may seem slight, but it results in two different views on initial mapping. Yet those views are not contradictory but rather complementary, and for the current discussion we assume mapping takes place in both directions.

Reconsidering the theoretical assumptions more accurately was required before discussing the mapping of *at*, because if we were to assume that mapping only goes from the L1 to the L2, we would be left with no clear view of how *at* is assembled. This is because in our domain the L1 offers less forms than the L2.

Having established the assumption that the initial mapping also goes from the L2 input back to the L1, we return to consider *at*, and suggest that based on the corpus data, *b* is a preferable candidate (20/42) compared to *al* (10/42). With the proposed initial mapping in place for all target items (*al* $\rightarrow$ *on*, *b* $\rightarrow$ {*in*, *at*}), we will now illustrate how feature reassembly proceeds, starting from the representations brought over from the L1 and arriving at target-like L2 representations.

Starting with *in*, in order to arrive at the target-like categorization which is represented by $[+\alpha{+}\gamma]$ (CONTAINMENT), a learner would have to adjust the initial representation transferred from *b* ($[+\alpha]$) by adding a single feature ($[+\gamma]$):

$$[+\alpha] \rightarrow [+\alpha + \gamma]$$

It could also be the case that instead of first acquiring the additional $[+\gamma]$, learners would pick up on $[-\beta]$. This would lead them to first categorize CONTAINMENT ($[+\alpha{+}\gamma]$) and PROXIMITY ($[+\alpha - \beta - \gamma]$) together, before adding a $[+\gamma]$ configuration to exclude PROXIMITY:[2]

$$[+\alpha] \rightarrow [+\alpha - \beta] \rightarrow [+\alpha - \beta + \gamma]$$

For the L2 form *at*, because we assume that a learner initially starts with an L1 representation of *b*, they start out with the feature composition of $[+\alpha]$. The next reassembly step to be taken is to add the feature configuration of $[-\beta]$ and arrive at $[+\alpha - \beta]$. This would lead the learner to a representation that covers the categories PROXIMITY and CONTAINMENT. Next, they would have to acquire the $[-\gamma]$ configuration, limiting the scope of the form to PROXIMITY. Finally, they would learn to remove $[+\alpha]$ and would be left with the target-like configuration $[-\beta - \gamma]$. The process can be summarized as follows:

$$[+\alpha] \rightarrow [+\alpha - \beta] \rightarrow [+\alpha - \beta - \gamma] \rightarrow [-\beta - \gamma]$$

As the process involves multiple steps, we do not attempt to predict the specific order with which they will be followed for every learner. However, we can look at the complete process and suggest that reassembly involves the addition of two features and the removal of one, a total of three steps.

Finally, we consider the reassembly process required to assemble *on*. We assumed earlier that the learner first maps *al* onto *on*, and the reassembly process that follows from that assumption is summarized as:

$$[-\alpha] \rightarrow [-\alpha + \beta] \rightarrow [+\beta]$$

The predicted reassembly of features as described above leads us to expect the following order

[2]Note that while one process assumes a final $[+\alpha + \gamma]$ configuration and the other assumes a $[+\alpha - \beta + \gamma]$ configuration, there is no difference in the corresponding conceptual categories, as $[+\alpha{+}\gamma]$ is sufficient to define CONTAINMENT.

of difficulty in acquiring target-like use of the English prepositions for L1 Hebrew speakers: *in* requires a single reassembly step and thus is the easiest to achieve, followed by *on* which requires two steps, with *at* being the most difficult, requiring three reassembly steps.

3.4 Testable Predictions

In turning the analysis into testable predictions, we propose the following approach. We examine the overall progression that is expected in order to arrive from a source L1 feature composition to the target L2 form. Then, we can suggest what is the area of the semantic space where learners are expected to face production difficulties related to specific forms, based on whether categories are "passed through" in the course of reassembly. For example, if we assume the target form *in* is assembled as portrayed above, $[+\alpha] \rightarrow [+\alpha+\gamma]$, then we derive that learners are expected to over-produce *in* in the categories ABSTRACT ($[+\alpha+\beta-\gamma]$) and PROXIMITY ($[+\alpha-\beta-\gamma]$) but they are not expected to produce *in* for contexts that fall under TARGET ($[-\alpha-\beta-\gamma]$) and SURFACE ($[-\alpha+\beta-\gamma]$), because those contain the feature $[-\alpha]$ which should at no point be included in the reassembly process of *in*. This idea is illustrated in Figure 3.4 with the preposition *in*, where the dark shaded area represents the target feature set, and the light shaded area represents the bounds of the target item earlier in the reassembly process.

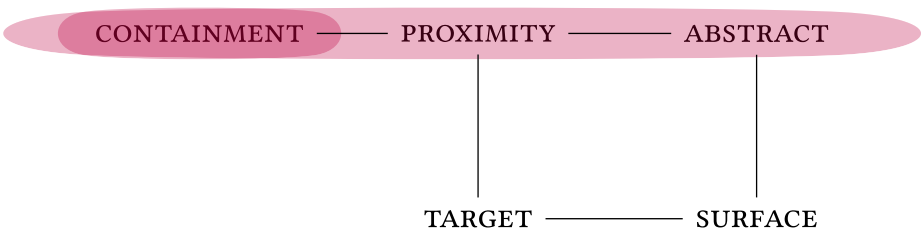

Figure 3.4: Reassembly of *in*

Leaving the feature-based terminology aside, we essentially predict over-production in contexts covered by the representations transferred from the L1, and we predict no over-production elsewhere. Using the same reasoning with *at* and *on* , we arrive at the summary of contexts where over-production is either expected or unexpected per form in Table 3.5.

3.5 Usage-Based Metrics of L2 Input

As mentioned earlier, our goal in this project is to examine how a usage-based approach to L2 acquisition is brought together with the FRH. In this section we look into usage-based variables of the L2 input, how they are gathered from corpus data, and how they might affect our predictions about L2 performance.

	at	in	on
Initial Mapping	ABSTRACT CONTAINMENT PROXIMITY	ABSTRACT CONTAINMENT PROXIMITY	SURFACE TARGET
Expected Over-production	ABSTRACT CONTAINMENT	ABSTRACT PROXIMITY	TARGET
Unexpected Over-production	SURFACE	TARGET SURFACE	COTAINMENT PROXIMITY
Grammatical	PROXIMITY TARGET	CONTAINMENT	ABSTRACT SURFACE

Table 3.5: Context type per preposition

We propose that learners might associate certain prepositions with verbs that precede them or nouns that appear in the prepositional phrase, based on frequent co-occurrences in the L2 input. To gain insight as to what associations might affect learners, we turn to extract measures of collocational distribution for our target preposition from the Corpus of Contemporary American English (Davies, 2008).[3] COCA is a general corpus and is not controlled for spatial meaning like our own corpus, and thus it includes collocations that arise because of verb-particle constructions or temporal uses, among other factors which are not studied here. However, the size of COCA allows us to have a better measure of frequency in the L2 input. As a measure of collocational frequency, the semantic context should not matter.

We used the COCA web interface to extract raw frequency counts for every verb-preposition and preposition-noun combination found in the Harry Potter corpus, using a window of -4 words between verbs and prepositions, and $+4$ words between prepositions and nouns. The frequency counts obtained from COCA were transformed into a logarithmic scale, and assigned to sentences from the HP corpus based on the verbs and nouns that appear together with the target preposition. In contexts where both a noun collocation and a verb collocation could be considered, we assign the higher frequency between the two as the frequency score of the fragment. Table 3.6 illustrates the calculation of frequency score for sample fragments from the corpus.

Several limitations of the current method should be explicitly noted: first, some co-occurring items such as the verb *to be* are too common in corpora to produce a meaningful measure. The same hold for some contexts where prepositions appear with pronouns or proper names as complements. Finally, we are only able to examine token distribution and cannot, for example, disambiguate word senses, or other sources of differences between the original fragments and the collocational corpus data.

[3]We use the fiction part of COCA based on previous positive correlation found between collocational frequencies in the fiction part of COCA and L2 learners' knowledge of collocations (Durrant, 2014).

Fragment	Preposition	Verb	Noun	V-P Frequency	P-N Frequency	Score
Harry looked at the floor.	at	look	floor	6.86	11.48	11.48
Mr Dursley always sat with his back to the window in his office on the ninth floor.	on	-	floor	-	9.66	9.66
”Well”, said Professor McGonagall, staring at the three of them	at	stare	-	10.34	-	10.34
Harry and Seamus swished and flicked, but the feather they were supposed to be sending skywards just lay on the desktop.	on	lay	desktop	8.99	4.88	8.99

Table 3.6: Fragments and frequency score

3.6 Summary

To summarize the results of the current chapter, we have seen how a process of semantic mapping of the English-Hebrew variation in prepositional use, is used to generate predictions about the progression of feature-based learning in L2 English. This has allowed us to identify specific contexts where we expect learners to over-produce L2 forms, and we identified contexts where over-production is not expected. We have also presented a measure of collocational frequency of verb-preposition and preposition-noun combinations, that we hypothesize to have an effect on learning. Based on these predictions, we ran the production experiment presented in the next chapter.

Chapter 4

Experiment

In order to test the predictions put forth in the previous chapter, we collect production data from L1 Hebrew speakers for each of the target English prepositions across different contexts, and check whether the errors that show up in the data follow the pattern that we had identified for expected or unexpected over-production. We will compare the prepositions that learners produce with those found in the sentences of the corpus. For ease of discussion, we refer to learners' production as 'correct' when they match the corpus data and as 'incorrect' or an 'error' when they do not.[1] If errors follow our predictions, this would support a process of feature-based learning that matches the contrastive L1-L2 feature geometry. For this purpose we ignore correct responses in the initial analysis and focus only on errors. The designation of context and error combinations as expected or unexpected is based on the hypothesized feature-based learning process, and is repeated here in Table 4.1.

	at	in	on
Initial Mapping	ABSTRACT CONTAINMENT PROXIMITY	ABSTRACT CONTAINMENT PROXIMITY	SURFACE TARGET
Expected Over-production	ABSTRACT CONTAINMENT	ABSTRACT PROXIMITY	TARGET
Unexpected Over-production	SURFACE	TARGET SURFACE	COTAINMENT PROXIMITY
Grammatical	PROXIMITY TARGET	CONTAINMENT	ABSTRACT SURFACE

Table 4.1: Context type per preposition

We will also look at the effect that the measure of collocational frequencies (as presented in Section 3.5) might have on learners' performance with respect to the different corpus items. This will be tested twice: first, we will look at whether frequency affects the kind of

[1]However, this does not mean that such deviations from the corpus necessarily result in sentences that are ungrammatical or infelicitous (as will be seen in Section 4.2.2).

errors that learners make, as explained above with respect to over-production. We will then also check whether collocational frequency affects a general measure of correctness. This is a coarser outcome variable that does not follow from our predictions about L1 transfer, but allows us to check for the varying level of difficulty that our items present to learners. As will be shown later in the discussion, the latter measure of correctness is found to be correlated with collocational frequencies, while the former is not.

4.1 Method

4.1.1 Procedure

Participants in the study were asked to fill in an online questionnaire consisting of two parts, built on the LimeSurvey platform. The first part was dedicated to evaluating participants' use of the target prepositions using fill-in-the-blank items derived from fragments of the original Harry Potter corpus that had the target prepositions removed. Participants were presented with one item at a time, and had to enter their response in a text box, meaning they were free to answer however they like and were not given a set of predetermined responses. Participants also had the choice to look at the Hebrew translation of any given item if they had difficulty in understanding the English text, but these translations were hidden by default.

The second part of the questionnaire was created in order to evaluate participants' proficiency in English, using the Oxford Quick Placement Test (2001) which was adapted to the questionnaire platform, so that participants could complete the entire task in a single session. The Oxford Quick Placement Test comprises of multiple-choice questions in varying forms (60 in total) and is intended for quick assessment of proficiency in CEFR compatible levels.[2] Assessment of English proficiency allowed us to select for intermediate to advanced learners, at the level of B2 or above. The assumption is that transfer effects would still be visible with advanced learners, but variation in the data would be reduced compared to testing with a wider range of learners.[3] The duration of the complete questionnaire was around 60 minutes, and participants were instructed to take a short break of several minutes between the two parts.

4.1.2 Participants

A total of 26 participants completed the questionnaire. We limit the current analysis to participants who ranked as B2 or above on the Oxford Quick Placement test (n=20, median

[2]A limited number of questions in the Oxford Quick Placement Test involve the same prepositions as our target items, in simple contexts.

[3]Because the proficiency of learners can only be assessed after the completion of the questionnaire, we have to exclude lower-ranking participants from the analysis. The selection threshold of B2 level was determined before data collection.

score 54.5 out of 60). Participants were recruited among the author's network and were not compensated for their participation.

4.1.3 Items

The prepositional use part of the questionnaire contained 10 items per category, making up for a total of 50 target items. Items were prepared by replacing target prepositions with a blank, in fragments of one or more sentences taken from the original corpus presented in the previous chapter. Fragments were chosen to contain exactly one target preposition in order to avoid any potential priming effects from prepositions that appear before or after the targets. Fragments were mostly left as-is, with some minor adjustments to control for length or added context where necessary. Filler items were created in a similar manner from corpus fragments containing the prepositions *to* (11), *out* (9) and *for* (10), making up for a total of 30 items. The underlying assumption for the choice of fillers was that participants would figure out they are being tested for preposition placement given the target items. Furthermore, an earlier pilot study confirmed that participants find the target contexts to be more difficult, and thus to somewhat stand out regardless of the filler items. Given those observations, we chose to have a slightly unbalanced ratio of target and filler items in the final study (50 vs. 30), in an effort to gather more data per participant while maintaining a feasible task duration. The items were set up for a within-subject design, meaning that every participant was tested on all items. The order of the experimental items was randomized per participant, excluding the first and last items which were always filler items.

In each category, five of the experimental items were designated as 'high frequency' and five as 'low frequency', based on their collocational score in an independent corpus (see 3.5). It is important to note that this is a relative measure per category, meaning that an item labeled as low in frequency in one category might have a frequency score in the high range of another category. This was done in order to work around the fact that our relatively small corpus contains a limited number of items per category. The mean frequency score (derived from raw frequencies on a log scale) across the low frequency items is 7.4 and the mean score for the high frequency items is 10, which makes them about 13 times as frequent, on average.

4.2 Results

We first present a brief overview of the results, followed by a detailed breakdown of the different analyses. The data provide limited support for the predictions made about the kind of errors learners would make in using the target prepositions. The largest effect is found with the preposition *in*, followed by inconclusive data for *at*, and no discernible effect for *on*. We also find no evidence that collocational frequency affects the error-based measure. However, collocational frequencies are found to be correlated with overall response correctness. We present the full analysis of over-production errors followed by the analysis for response correctness based on frequencies.

With 50 target items and 20 participants a total of 1000 responses were collected. Of these, 275 responses represent deviations from target-like preposition use, and 78 of those deviations are cases where participants provided an answer which is not one of *at*, *in*, or *on*. An overview of the responses per category (limited to responses of *at*, *in* and *on*) is presented in Table 4.2.

Category	**at**	**in**	**on**	**Total**
ABSTRACT	75	26	97	198
CONTAINMENT	31	144	14	189
PROXIMITY	155	3	15	173
SURFACE	15	6	162	183
TARGET	167	1	11	179
Total	443	180	299	922

Table 4.2: Responses by category

4.2.1 Analysis of Over-Production

The same data from Table 4.2 are presented in Table 4.3 using the division to types of contexts that we identified per preposition. As the design is not balanced, the table shows both the number of observed prepositions per context and the number of contexts per condition.

Preposition	**Expected**	**Unexpected**	**Grammatical**
at	106/400	15/200	322/400
on	11/200	29/400	259/400
in	29/400	7/400	144/200
Total	146/1000	51/1000	725/1000

Table 4.3: Responses by context type

Table 4.3 shows that the data contain more production errors in expected contexts for *at* and *in* compared to unexpected contexts. For *on*, the number of errors in expected contexts is lower than in unexpected contexts, meaning that if there is an effect, it is reversed in the case of *on*. For a statistical analysis, the unbalanced experimental design requires separate analyses per preposition. This is because the same items count as different kinds of contexts per prepositions, and thus cannot be analyzed together.

For each preposition, we look at a subset of the data that includes the expected and unexpected contexts for that preposition. Our dependent variable is a binary value, coding whether the specific preposition was given as a response. Thus we test whether participants are indeed more likely to produce the relevant preposition in contexts where we expect them to.[4] To do that, we use logistic regression and compare a null model with random intercepts for items and participants, against a respective model where a fixed effect is added for context type (expected vs. unexpected). All of the regression models reported here are generalized mixed-effects models with a logit link function, fit by maximum likelihood using the `lme4` package for R (Bates et al., 2015; R Core Team, 2018).

For *at*, a model with a fixed effect for context type improves on a null model with an AIC of 461.06 over 462.83, $\chi^2(1) = 3.764, p = 0.0524$. For *in* a similar model improves on a null model with an AIC of 247.8 over 251.74, $\chi^2(1) = 5.9355, p = 0.0149$. For *on* a model with a fixed effect for context type does not improve on a null model, with AIC increasing from 284.61 to 286.47 ($\chi^2(1) = 0.143, p = 0.7$). The model estimates are reported in tables 4.4 and 4.5. This leaves us with no support for the expected effect for *on*, and on the other hand positive results for *at* and *in* which are quite similar in terms of effect size. However, there are reasons to treat the results for *at* as inconclusive. First, the model for *at* is lacking in statistical power ($p > 0.05$). Second, as will be discussed in the following section, the analysis of over-production of *at* in some of the contexts is inaccurate.

4.2.2 The Category ABSTRACT

In Figure 4.1 we plot overall response correctness per each category. Visible in the plot is the low target-like performance of learners with items in the ABSTRACT category. We proceeded to take a closer look at the items that make up the category by consulting two native speakers of British and American English, for their judgments on the questionnaire items. Both speakers agree that 6 out of the 10 items in the category, that appear in the corpus with *on* are also acceptable with *at* instead. In fact, both speakers prefer *at* over *on* in 3 of the 6 items. As seen in Table 4.2, learners mostly answered *at* when they did not answer *on*. These data and the intuitions from our native speaker consultants make it difficult to draw conclusions based on the items in the category.

[4]We stress the idea that in both expected and unexpected contexts, producing the target preposition is treated as an error, but the difference lies in whether the error is in agreement with the projected L1 transfer.

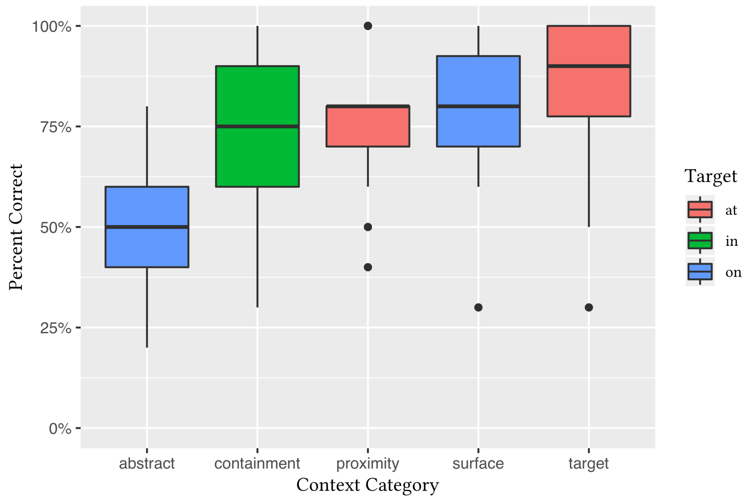

Figure 4.1: Target-like responses by context category

4.2.3 Frequency Effects

As described earlier, the experimental items were chosen per category while considering a frequency measure of verb-preposition and preposition-noun collocations in the independent COCA corpus (Section 3.5). We originally included frequency scores in the design with the intention of testing whether they have an effect on L1 transfer (as measured by over-production). However, frequency score yielded no improvements as a fixed predictor to the regression models for over-production, nor did it show an effect when used as a sole predictor. We will now discuss the effect of frequency score on the different outcome measure of production correctness.

While slightly difficult to read, Figure 4.2 shows that individual items with a higher frequency prompted less errors in general. A post-hoc analysis reveals an effect for collocational frequency with respect to production correctness. A mixed-effects model with a single fixed effect for frequency score, and random effects for participant and item, resulted in an improvement over a respective null model, with AIC of 958.32 over 960.67 ($\chi^2(1) = 4.3456, p = 0.037$). The estimates of the model are reported in Table 4.7.

The effect size for collocational frequency in terms of the odds ratio is given by $e^{0.2809}$. As the model was fit with log-transformed frequencies using the natural logarithm, interpretation is made slightly easier by transforming the estimate to a base-2 logarithm: $\exp\left(\frac{0.2809}{\log_2(e)}\right) \approx 1.215$. The way to interpret the transformed odds-ratio is that it suggest that a learner is 1.215 times likelier to use a preposition correctly in a context that appears twice as often in the corpus.

Fixed Effects	Estimate	S.E.	z	p
Unexpected	-3.452	0.748	-4.613	< 0.001
Expected	1.673	0.842	1.987	0.047
Random Effects	Variance	σ		
Item	3.423	1.85		
Participant	0.465	0.682		

Table 4.4: Model estimates for *at*

Fixed Effects	Estimate	S.E.	z	p
Unexpected	-5.943	0.92	-6.463	< 0.001
Expected	1.777	0.739	2.405	0.016
Random Effects	Variance	σ		
Item	2.295	1.515		
Participant	2.511	1.585		

Table 4.5: Model estimates for *in*

Fixed Effects	Estimate	S.E.	z	p
Intercept	-3.262	0.388	-8.415	< 0.001
Random Effects	Variance	σ		
Item	0.853	0.924		
Participant	0.697	0.835		

Table 4.6: (Null) Model estimates for *on*

Fixed Effects	Estimate	S.E.	z	p
Intercept	-1.008	1.185	-0.851	0.395
(Log-)Frequency	0.2809	0.132	2.122	0.034
Random Effects	Variance	σ		
Item	2.256	1.502		
Participant	0.601	0.775		

Table 4.7: Model estimates for response correctness

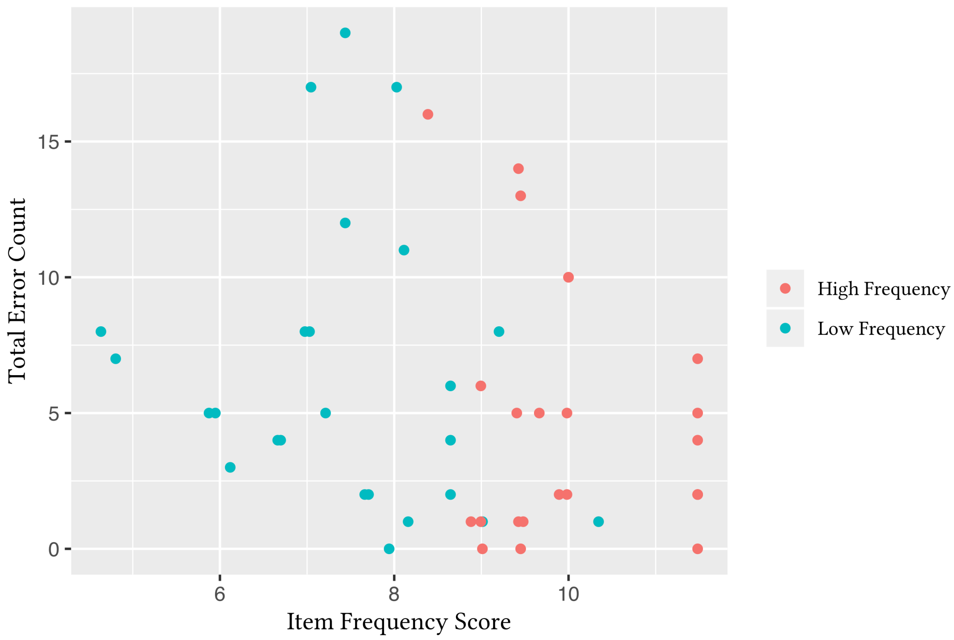

Figure 4.2: Errors by item frequency score

4.3 Discussion

Of the three English prepositions *at*, *in* and *on* we have found possible influence of the L1 only for *in* and *at*. The results for *at* are not clear enough in terms of statistics and methodological issues, which will have to be addressed in future work. A different kind of prediction that is borne out by the experimental data is that there exists a correlation between learners' knowledge of prepositions and collocational frequencies in corpus data. This is consistent with previous results (Durrant, 2014) and goes back to the basic principles of usage-based learning and corpus linguistics—in being exposed to L2 input, learners acquire the target knowledge of forms based on their use in context (in the limited interpretation where a context is a PP environment), and so they do better with frequently occurring contexts. The measure of frequency in the L2 input can in turn be extracted from corpus data, of the kind that learners are likely exposed to. While that is a positive result in favor of frequency effects in SLA, more refined work is required in order to argue for an effect that goes beyond the division between very frequent and very infrequent items.

Going back to reconsider the theoretical framework of the Feature Reassembly hypothesis, the experiment did not bring strong evidence in support of feature-based effects of L1 transfer in the locative domain. Taken at a broader perspective, the results are not in line with our interpretation of the FRH. Our focus has been on two main properties of the FRH: initial mapping and reassembly based on contrastive analysis with the L1. We have based our assumptions for both factors on the L1-L2 distribution of forms in our cross-linguistic corpus.

The mapping of corpus data was created as a model of the kind of contrastive analysis and assembly of representations done by learners. We based the analysis of variation purely on the distribution of forms in a corpus which is representative of the L2 input available to advanced learners of English. Our only implicit assumption with respect to semantics comes from the selection of contexts of spatial reference out of all of the corpus occurrences of our target prepositions.

Having put our semantic map into an hypothesis about how Hebrew speakers reassemble lexical items to arrive at target-like configurations of the English prepositions, we attempted to identify testable effects of the reassembly process. We proposed to focus on the kind of contexts where learners make over-production errors, yet we could not convincingly show that over-production follows from L1 transfer.

Several methodological choices warrant further reflection. In terms of the experimental setup, learners were very much aware that they are being tested on prepositional use, which could have distorted their performance. We have also reported an issue with several of the items in the ABSTRACT category, which is possibly tied to optionality in the use of forms. We aimed to use the original corpus sentences in the experiment, because our usage-based view on the conceptual domain provided us with no tools to construct artificial examples. Judgments by native speakers demonstrate the possible limitations of that approach, especially when considering the fact that we used a corpus that is sourced from one text by a single author.

Chapter 5

Conclusions

The goal of the current project has been to examine how far cross-linguistic corpus data can go in predicting the effects of L1 transfer among learners of a second language. We have started from current theoretical approaches to the study of transfer from the generative perspective, and have proposed a method for testing the theory by modeling the process of acquisition. We have also shown how corpus data is used to empirically ground the potential learning that ensues by exposure to L2 input and contrastive analysis with the L1.

While the results of the current project did not achieve conclusive outcomes, we have nonetheless managed to demonstrate the level of sophistication required by empirical studies in SLA, if they are to match the theoretically appealing tenets of feature-based learning. Many of the assumptions taken in the setup of the current study could be reevaluated, first among them is the method of measuring transfer effects and its focus on production errors, yet the high-level goal of integrating SLA theory with an empirical survey of L1-L2 mappings and L2 input should be maintained.

Finally, it is worthwhile to reiterate the fact that while in our experiment we have tested relatively advanced learners, they were still found to produce non target-like prepositions in 27% of contexts. We take this to mean that achieving target-like proficiency in the use of locative prepositions is not a simple task for learners, even in a relatively constrained subset of uses. While we could not convincingly point at L1 transfer as the source of that difficulty, we can conclude that the problem domain of prepositions and spatial reference is relevant for future work in L2 semantics and SLA.

Bibliography

Alonso, R. A., Cadierno, T., and Jarvis, S. (2016). Crosslinguistic influence in the acquisition of spatial prepositions in english as a foreign language. In Alonso, R. A., editor, *Crosslinguistic Influence in Second Language Acquisition*, pages 57–72. Multilingual Matters.

Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.

Bowerman, M. (1996). Learning how to structure space for language: A crosslinguistic perspective. In Bloom, P., Peterson, M. A., Nael, L., and Garrett, M. F., editors, *Language and Space*, pages 385–436. MIT Press, Cambridge, MA.

Cho, J. and Slabakova, R. (2014). Interpreting definiteness in a second language without articles: The case of L2 Russian. *Second Language Research*, 30(2):159–190.

Cho, J. and Slabakova, R. (2017). A feature-based contrastive approach to the L2 acquisition of specificity. *Applied Linguistics*, 38(3):318–339.

Choi, S. and Bowerman, M. (1991). Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition*, 41(1-3):83–121.

Chomsky, N. (1981). *Lectures on government and binding.* Foris, Dordrecht.

Chomsky, N. (1995). *The minimalist program.* MIT Press, Cambridge, MA.

Davies, M. (2008). The corpus of contemprorary american English (COCA): 560 million words, 1990-present. Available online at https://corpus.byu.edu/coca/.

Dominguez, L., Arche, M. J., and Myles, F. (2011). Testing the predictions of the feature-assembly hypothesis: evidence from the L2 acquisition of Spanish aspect morphology. In *BUCLD 35: Proceedings of the 35th Annual Boston University Conference on Language Development*, volume 35. Cascadilla Press.

Durrant, P. (2014). Corpus frequency and second language learners' knowledge of collocations: A meta-analysis. *International Journal of Corpus Linguistics*, 19(4):443–477.

Ellis, N. C. (2006). Selective attention and transfer phenomena in L2 acquisition: Contingency, cue competition, salience, interference, overshadowing, blocking, and perceptual learning. *Applied Linguistics*, 27(2):164–194.

Feist, M. I. (2000). *On in and on: An investigation into the linguistic encoding of spatial scenes.* PhD thesis, Northwestern Univ. Evanston.

Feist, M. I. and Gentner, D. (2003). Factors involved in the use of in and on. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 25.

Gil, K. and Marsden, H. (2013). Existential quantifiers in second language acquisition: A feature reassembly account. *Linguistic Approaches to Bilingualism*, 3(2):117–149.

Haspelmath, M. (2003). The geometry of grammatical meaning: Semantic maps and cross-linguistic comparison. *The New Psychology of Language*, 2(1976):1–30.

Hawkins, R. (2005). Explaining full and partial success in the acquisition of second language grammatical properties. *Second Language*, 4:7–25.

Ionin, T. and Montrul, S. (2010). The role of L1 transfer in the interpretation of articles with definite plurals in L2 English. *Language Learning*, 60(4):877–925.

Jarvis, S. and Odlin, T. (2000). Morphological type, spatial reference, and language transfer. *Studies in second language acquisition*, 22(4):535–556.

Lardiere, D. (2009). Some thoughts on the contrastive analysis of features in second language acquisition. *Second Language Research*, 25(2):173–227.

Lardiere, D. (2012). Linguistic approaches to second language morphosyntax. In Gass, S. M. and Alison, M., editors, *The Routledge Handbook of Second Language Acquisition.*

Levinson, S. C. and Meira, S. (2003). 'Natural concepts' in the spatial topologial domain–adpositional meanings in crosslinguistic perspective: An exercise in semantic typology. *Language*, 79(3):485–516.

Quick Placement Test (2001). Oxford University Press.

R Core Team (2018). *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria.

Rothman, J. and Slabakova, R. (2017). The generative approach to SLA and its place in modern second language studies. *Studies in Second Language Acquisition*, pages 1–26.

Slobin, D. I. (1996). Two ways to travel: Verbs of motion in English and Spanish. In Shibatani, M. and Thompson, S., editors, *Grammatical constructions: Their form and meaning*, pages 195–219.

Talmy, L. (1975). Figure and ground in complex sentences. In *Annual Meeting of the Berkeley Linguistics Society*, volume 1, pages 419–430.

van der Klis, M., Le Bruyn, B., and de Swart, H. (2017). Mapping the Perfect via translation mining. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 497–502.

Wulff, S., Ellis, N. C., Römer, U., Bardovi-Harlig, K., and Leblanc, C. J. (2009). The acquisition of tense–aspect: Converging evidence from corpora and telicity ratings. *The Modern Language Journal*, 93(3):354–369.

Zwarts, J. (1997). Lexical and functional properties of prepositions. *Lexikalische und grammatische Eigenschaften präpositionaler Elemente. Linguistische Arbeiten*, 371:1–19.

Zwarts, J. (2010). Semantic map geometry: Two approaches. *Linguistic Discovery*, 8(1):377–395.

Zwarts, J. (2017). Spatial semantics: Modeling the meaning of prepositions. *Language and Linguistics Compass*, 11(5).

Zwarts, J. and Winter, Y. (2000). Vector space semantics: A model-theoretic analysis of locative prepositions. *Journal of Logic, Language and Information*, 9(2):169–211.

Appendix A

Experimental Items

Fragment	P	Category	V	N	V-P Frequency	P-N Frequency	Score	
But he did seem to realise he was being watched, because he looked up suddenly __ the cat , which was still staring at him from the other end of the street .	at	proximity	look	cat	5.32	11.48	11.48	high
Hermione was stomping up the stairs looking disapprovingly __ the package in Harry’s hand .	at	proximity	look	package	3.85	11.48	11.48	high
Ron pointed: __ the end of a passage to the left, something huge was moving towards them.	at	proximity	-	end	9.43	NA	9.43	high
He clicked it once and twelve balls of light sped back to their street lamps so that Privet Drive glowed suddenly orange and he could make out a tabby cat slinking around the corner __ the other end of the street.	at	proximity	-	end	9.43	NA	9.43	high
Mr Dursley blinked and stared __ the cat	at	proximity	stare	cat	5.32	10.34	10.34	high
Ron, __ the next table , wasn’t having much more luck.	at	proximity	-	table	9.20	NA	9.20	low
”What did you say you’ve got __ home, Malfoy, a Comet Two Sixty?”	at	proximity	-	home	9.01	NA	9.01	low
At once, the black ball rose high in the air and then pelted straight __ Harry’s face.	at	proximity	pelt	face	8.11	0.69	8.11	low
He glared __ them all as if to say, ”Or else”.	at	proximity	glare	them	NA	8.16	8.16	low
He dashed back across the road, hurried up to his office, snapped __ his secretary not to disturb him, seized his telephone and had almost finished dialling his home number when he changed his mind.	at	proximity	snap	secretary	NA	6.70	6.70	low
Harry looked __ the floor.	at	target	look	floor	6.86	11.48	11.48	high
A murmur ran through the crowd as Adrian Pucey dropped the Quaffle, too busy looking over his shoulder __ the flash of gold that had passed his left ear .	at	target	look	flash	3.93	11.48	11.48	high
Harry, Ron and Hermione looked __ each other, wondering what to tell him.	at	target	look	-	NA	11.48	11.48	high
Professor McGonagall was looking __ Ron and Harry.	at	target	look	PROPER	NA	11.48	11.48	high
At these words, Hermione seized Hagrid’s binoculars, but instead of looking up __ Harry , she started looking frantically at the crowd.	at	target	look	PROPER	NA	11.48	11.48	high
”Well”, said Professor McGonagall, staring __ the three of them.	at	target	stare	them	NA	10.34	10.34	low
”You don’t mean - you can’t mean the people who live here?” cried Professor McGonagall, pointing __ number four.	at	target	point	number	6.08	8.65	8.65	low
He pointed __ the three balls left inside the box.	at	target	point	ball	5.70	8.65	8.65	low
Suddenly, people were pointing up __ Harry all over the stands .	at	target	point	PROPER	NA	8.65	8.65	low
”Oy, pea-brain!” yelled Ron from the other side of the chamber, and he threw a metal pipe __ the	at	target	throw	it	NA	7.66	7.66	low

Continued on next page

Continued from previous page

Fragment	P	Category	V	N	V-P Frequency	P-N Frequency	Score	
troll.								
Mr Dursley sat frozen __ his armchair.	in	containment	sit	armchair	6.13	10.00	10.00	high
As he sat __ the usual morning traffic jam, he couldn't help noticing that there seemed to be a lot of strangely dressed people about.	in	containment	sit	jam	NA	10.00	10.00	high
Harry's wand had still been __ his hand when he'd jumped	in	containment	be	hand	9.89	NA	9.89	high
"There will be books written about Harry - every child __ our world will know his name!"	in	containment	-	world	9.48	NA	9.48	high
"You'd better get off to Gryffindor Tower. Students are finishing the feast __ their houses."	in	containment	finish	house	9.41	6.64	9.41	high
Hermione was crying __ the girls toilets and wanted to be left alone	in	containment	cry	toilet	5.77	7.21	7.21	low
He was busy rummaging __ his cloak, looking for something.	in	containment	rummage	cloak	5.48	5.95	5.95	low
"The key is __ the lock", Harry muttered. "We could lock it."	in	containment	be	lock	6.12	NA	6.12	low
"You're lucky you weren't killed. Why aren't you __ your dormitory?"	in	containment	be	dormitory	4.80	NA	4.80	low
"Why isn't he down __ the dungeons with the rest of the teachers?"	in	containment	be	dungeon	4.63	NA	4.63	low
Mr Dursley always sat with his back to the window in his office __ the ninth floor.	on	abstract	-	floor	9.66	NA	9.66	high
There was a tabby cat standing __ the corner of Privet Drive , but there wasn't a map in sight.	on	abstract	stand	corner	7.44	9.01	9.01	high
It was __ his way back past them that he caught a few words of what they were saying.	on	abstract	-	way	9.45	NA	9.45	high
"I must have passed a dozen feasts and parties __ my way here."	on	abstract	-	way	9.45	NA	9.45	high
But __ the edge of town, drills were driven out of his mind by something else.	on	abstract	-	edge	8.39	NA	8.39	high
Around the Quidditch pitch were six golden poles with hoops __ the end.	on	abstract	-	end	8.03	NA	8.03	low
"Oliver Wood will meet you tonight __ the Quidditch pitch at seven o'clock for your first training session."	on	abstract	meet	pitch	4.14	7.04	7.04	low
It was __ the corner of the street that he noticed the first sign of something peculiar - a cat reading a map.	on	abstract	be	corner	7.44	NA	7.44	low
Dumbledore turned and walked back down the street. __ the corner he stopped and took out the silver Put-Outer.	on	abstract	-	corner	7.44	NA	7.44	low
"Funny stuff __ the news", Mr Dursley mumbled in front of the TV.	on	abstract	-	news	6.66	NA	6.66	low
It was now sitting __ his garden wall.	on	surface	sit	wall	8.80	9.98	9.98	high
Quirrell saw the troll, let out a faint whimper and sat quickly down __ a toilet , clutching his heart .	on	surface	sit	toilet	5.98	9.98	9.98	high
He laid Harry gently __ the doorstep, took a letter out of his cloak , tucked it inside Harry's blankets and then came back to the other two.	on	surface	lay	doorstep	6.21	8.99	8.99	high
Harry and Seamus swished and flicked, but the feather they were supposed to be sending skywards just lay __ the desktop.	on	surface	lay	desktop	4.88	8.99	8.99	high
Dumbledore reached out and patted her __ the shoulder.	on	surface	pat	shoulder	8.88	6.94	8.88	high
Mr Dursley picked up his briefcase, pecked Mrs Dursley __ the cheek and tried to kiss Dudley goodbye but missed	on	surface	peck	cheek	7.94	4.54	7.94	low
He could just see the bundle of blankets __ the step of number four .	on	surface	-	step	7.71	NA	7.71	low
He wiped it __ the troll's trousers.	on	surface	wipe	trousers	4.77	7.03	7.03	low
If anyone looked out of their window now, they wouldn't be able to see anything that was happening down __ the pavement.	on	surface	happen	pavement	6.57	6.98	6.98	low
He drummed his fingers __ the steering wheel and his eyes fell on a huddle of these weirdos standing quite close by.	on	surface	drum	wheel	5.87	5.82	5.87	low