Correlation between Lexical and Determination Types

Oliver Hellwig
University of Düsseldorf
ohellwig@phil-fak.uni-duesseldorf.de

Wiebke Petersen
University of Düsseldorf
petersen@phil.uni-duesseldorf.de

Abstract

The Paper presents a corpus based study of Löbner's theory of determination types. According to Löbner, nouns prefer a syntactic determination mode that is congruent to their inherent type of lexical determination. The study applies machine learning methods to a large corpus of German texts for detecting the determination modes of nouns automatically and addresses the problem of inducing the lexical determination type of a noun from its distributional fingerprint of occurrences in different determination modes

1 Introduction and previous research

The paper aims at assessing the theory of congruent determination (Löbner, 2011) by applying methods from Computational Linguistics to a corpus of German sentences. Löbner claims that the lexical semantics of a noun influences its syntactic mode of determination. He distinguishes two binary semantic features that result in a fourfold classification of nouns; each class corresponds to a 'natural' mode of determination. The two binary features are uniqueness ($\pm U$) and relationality ($\pm R$). Inherently unique nouns refer to a single, unique referent (e.g., 'sun' or 'steer' in contrast to 'star' or 'wheel'). The reference of an inherently relational noun depends on an additional possessor argument which needs to be specified (e.g., 'surface [of something]', 'brother [of someone]' in contrast to 'sun' or 'stone'). Combining the two binary features results in four lexical determination types of nouns: *Sortal* nouns describe non-unique, non-relational referents ($-U-R$, 'stone', 'chair'); *individual* nouns describe non-relational, but unique referents ($+U-R$, 'sun', 'Peter'), *relational* nouns describe non-unique referents that require a possessor ($-U+R$, 'brother', 'leaf'), and *functional* nouns describe unique, relational referents ($+U+R$, 'steer', 'trunk'). The four basic lexical noun types are summarized in Table 1.

In usage, nouns often occur in determination modes that are not congruent with their underlying lexical types. Löbner (2011) claims that incongruent determination is a marked option and involves a type shift of the noun in the sense of Partee (1986). For example, if a sortal noun is used in a context in which it is uniquely referring for pragmatic reasons, the noun is shifted from $-U$ to $+U$, as in 'the train arrives' when someone waits at the station. The lexical distinction in $\pm U$ nouns and uniqueness shifts are clearly marked in languages with a split article system that distinguish between strong and weak definite articles. In contrast to standard German, for example, the Ripuarian dialect spoken in the Rhineland makes this distinction. The weak definite article 'dr' is used with inherently $+U$ nouns. If a non-unique noun ($-U$) is used in a definite context, the incongruent use is marked by the strong definite article 'dä':

(1) Ripuarian [cf. Löbner (2011)]:

a. Dr Zoch$_{ind.}$ kütt.
DEFARTWEAK parade comes.

b. Dä Zoch$_{sort.}$ kütt.
DEFARTSTRONG train comes.

The sentences (1-a) and (1-b) only differ with respect to the definite article being used. While in (1-b) 'Zoch' is read as a sortal noun ('train'), in (1-a) 'Zoch' refers to the carnival parade, which is conceptualized as a unique individual in cities like Cologne or Düsseldorf.

The lexical $\pm R$ distinction and type shifts along this dimension are less frequently marked in languages. However, languages with an overt morphology for (in)alienability mark shifts from $-R$ to $+R$ (Ortmann, 2015):

Proceedings of the Int. Conference of the German Society for Computational Linguistics and Language Technology,
pages 130–137, University of Duisburg-Essen, Germany, Sep 30–Oct 2 2015.

	non-unique reference [-U]	unique reference [+U]
non-relational [-R]	**sortal noun** tree, stone, woman	**individual noun** pope, universe, Mary
relational [+R]	**proper relational noun** sister [of], student [of], page [of]	**functional noun** mother [of], dean [of], cover [of]

Table 1: The four basic lexical noun types according to Löbner (2011)

(2) Diegueño [after Nichols (1992)]

	a.		b.
	ʔ–ətaly		*ʔ–$ə^n$–ewa*
	1sg–mother		1sg–poss–house
	'my mother'		'my house'

In Diegueño, for example, the word 'ewa' ('house') is conceptualized as a sortal noun ($-R$). In order to be used in a possessive context, it has to be shifted by a possessive marker to $+R$. Other languages like Yucatec indicate by a derelativizing marker shifts from $+R$ to $-R$.[1]

The $\pm R$ distinction also plays a crucial role in formal compositional approaches to the semantics of genitive constructions. In this context, the question arises whether in constructions such as 'Mia's sister' and 'Mia's pen' the possessor 'Mia' is an argument or a modifier of the head noun. Previous research has given three different answers to this question: (1) the possessor is always an argument of the head noun, which is shifted to $+R$ if it is lexically non-relational (Vikner and Jensen, 2002); (2) the possessor is always a modifier and there are no genuine $+R$ nouns in the lexicon; and (3) whether the possessor acts as an argument or as a modifier depends on its underlying lexical $\pm R$ feature (Barker, 1995; Barker, 2011). Partee and Borschev (2003) critically discuss all three options, discard the second and come to the conclusion that the interpretation of genitive constructions is determined by the lexical $\pm R$ feature of a noun.

The discussed examples give evidence for Löbner's fourfold noun classification from a typological perspective. However, Löbner (2011) argues that lexical $\pm U$ and $\pm R$ distinctions and shifts between the determination types are also relevant for languages that do not mark lexical types and shifts overtly. He claims that the 'natural' unshifted mode of determination – that means the determination mode that is congruent with the lexical type of determination – influences the relative frequencies of the different determination modes in which a noun occurs. He reports statistical evidence for his claim from studies for Swedish (Fraurud, 1990) and English (Vieira, 1998; Nissim, 2004; Jensen and Vikner, 2004). However, these studies concentrate only on either the $\pm U$ or the $\pm R$ distinction. A first statistical investigation that combines both features has been carried out in Horn and Kimm (2014) for German using a small hand annotated corpus consisting of two fictional short stories with only 456 noun tokens. The study shows that for both features, $\pm U$ and $\pm R$, the congruent determination is more frequent than the incongruent one.

Det$_{+U}$ determination mode
singular definite ('die Sonne', ['the sun']; 'Mia')
contracted singular definite ('zur Sonne', ['to the sun'])
possessive pronoun ('mein Kopf', ['my head'])
left genitives ('Mias Kopf', ['Mia's head'])
Det$_{-U}$ determination mode
indefinite article ('ein Stein' ['a stone']),
plural ('(die) Steine' ['(the) stones']),
quantifiers ('jeder/kein/einige/beide/zwei Stein(e)', ['every/no/some/both/two stone(s)]')
contrastive demonstrative ('dieser/jener Stein, ['this/that stone']),
interrogative ('welcher Stein', ['which stone'])
Det$_{+R}$ determination mode
possessive pronoun ('mein Kopf', ['my head'])
left genitive ('Mias Kopf', ['Mia's head')
prepositional phrase ('der Kopf von Mia', ['the head of Mia'])
Det$_{-R}$ determination mode
absolute use

Table 2: Investigated modes of determination in German

This paper aims at testing the hypothesis of congruent determination for German with a large amount of data. In parallel to the lexical features $\pm U$ and $\pm R$, we use the features $\mathrm{DET}_{\pm U}$ and $\mathrm{DET}_{\pm R}$ in order to distinguish the determination modes in actual language use. Table 2 shows a selection of different modes of determination in German with their respective features based on Löbner (2011). In order to avoid tedious manual annotations of actual uses, labelers for $\mathrm{DET}_{\pm U}$ and $\mathrm{DET}_{\pm R}$ are developed and applied to a corpus of German news texts. The syntactic determination

[1]An investigation on further typological evidence for type shifts can be found in Ortmann (2015).

mode of each noun token is derived from the results of the labelers. For each noun lemma, its determination modes are cross-tabulated with its lexical semantic determination type. The semantic types are provided by a manual annotation of the $\pm U$ and $\pm R$ features of frequent nouns drawn from a German dictionary (Duden, 1997).

The rest of the paper is structured as follows. In Section 2, we describe the two labelers for $\mathrm{DET}_{\pm U}$ and $\mathrm{DET}_{\pm R}$, whereby the $\mathrm{DET}_{\pm R}$-labeler is an enhancement of the work presented in Hellwig and Petersen (2014). Section 3 investigates the correlation between the lexical determination type and determination modes in language use for a group of high-frequency lemmata, thereby providing an estimation for the frequency of incongruent determination in German.

2 Automatic annotation of types of determination modes

The automatically annotated modes of determination will be cross-tabulated with the lexical types of determination in order to investigate whether the latter influences the former. Therefore, automatic $\mathrm{DET}_{\pm U}$ and $\mathrm{DET}_{\pm R}$ classifiers are required that have a high precision. Recall is less important for our purpose, as long as the classifiers do not systematically overlook a particular mode of determination.

2.1 Detecting unique reference ($\mathrm{DET}_{\pm U}$)

Due to the German article system, detecting the uniqueness of reference is a comparatively easy task that can be performed with a rule-based approach. The rules used in our $\mathrm{DET}_{\pm U}$ classifier rely on the POS information that is produced by the MATE parser (Bohnet and Nivre, 2012) and the Stanford NLP parser for German (Rafferty and Manning, 2008). Additionally, morphological information about number created by MATE is taken into account. The labeling rules detect the modes of determination as given in Table 2. Nouns in singular number that are directly preceded by (contracted) definite articles, possessive pronouns or another noun in genitive case are labeled as $+U$. Plural nouns or nouns in singular number that are preceded by a member from a finite list of function words ("jeder" ['each'], "solcher" ['such'],"ein" ['a'], "dieser" ['this'],...) are labeled as $-U$. All other nouns are labeled as undetermined with regard to the uniqueness of their referents.

Type	P	R	F
$\mathrm{SVM}^{\mathrm{HMM}}$	92.76	75.7	83.37
CRF	92.23	71.68	80.66
ME	93.97	64.83	76.73
Tree	67.14	84.86	74.97

Table 3: Word-based evaluation by classifier for the class DET_{+R}, 30-fold cross-validation, no additional training data

Type	P	R	F
$\mathrm{SVM}^{\mathrm{HMM}}$	99.06	99.77	99.42
CRF	98.9	99.77	99.33
ME	98.63	99.84	99.23
Tree	99.4	98.38	98.89

Table 4: Word-based evaluation by classifier for the class DET_{-R}; same settings as in Table 3

2.2 Detecting relationality ($\mathrm{DET}_{\pm R}$)

The $\mathrm{DET}_{\pm R}$ classifier builds on the relation detection classifier described in Hellwig and Petersen (2014), which performs a three-class labeling task with classes POR (noun in possessor position), PUM (noun in possessum position), and no-poss. For our purpose, we can reduce the task to a binary classification problem for the features $\mathrm{DET}_{\pm R}$ by merging the classes POR and no-poss to DET_{-R} and identifying the class PUM as DET_{+R}. In Hellwig and Petersen (2014) three statistical classification methods are used that are able to handle sequential data from a categorical scale: Hidden Markov Support Vector Machines ($\mathrm{SVM}^{\mathrm{HMM}}$) (Altun et al., 2003), Conditional Random Fields (CRF) (Lafferty et al., 2001), and Maximum Entropy (ME) (Ratnaparkhi, 1998). Additionally, Hellwig & Petersen experiment with a rule-based tree classifier that shows low precision, but rather good recall scores. Table 3 and Table 4 show the results of the binary classification for 30 cross-validations. Due to the reduced complexity of the task, the results for DET_{+R} are slightly better than those reported for PUM in Hellwig and Petersen (2014). The results in Table 4 demonstrate that reliable results should be expected for DET_{-R} determination.

We test three strategies to improve over the baseline results for DET_{+R} that are shown in Table 3. In the first approach, we postprocess the output of the four basic classifiers with another nonlinear classifier. For this sake, the symbolic outputs of $\mathrm{SVM}^{\mathrm{HMM}}$ and of the tree classifier are transformed into the pseudo-confidence values of 1.0

Meta-classifier	P	R	F
Logistic regression	94.70	74.87	83.63
Neural network	94.46	75.76	84.08

Table 5: Results of post-processing the classifications with meta-learners; same settings as in Table 3

Voting	P	R	F
3-majority	95.23	73.48	82.95
4-majority	99.27	51.78	68.06

Table 6: Results of combining single classifiers by majority voting; same settings as in Table 3

for the class DET_{+R} and 0.0 for the class DET_{-R}, and the confidence values generated by CRF and ME are added to this feature vector. The resulting four-dimensional vectors, which consist of pseudo-confidences for SVM^{HMM} and the tree classifier and of true confidence values for CRF and ME, are used for training (a) a logistic regression and (b) a neural network with one hidden layer of dimension 2. Table 5 records the results of applying these two meta-learners. Both sets of results don't differ significantly from the best single classifier (refer to Table 3).

In the second improvement approach, we apply semi-supervised learning to the binary classification problem. First, we train all classifiers with the full set of gold data used in Hellwig and Petersen (2014), which consists of approximately 1.100 manually annotated sentences. Next, we use the trained classifiers to annotate a holdout-part of the new unannotated corpus (1M sentence extract from the Leipzig corpora, 2005) without supervision. Because SVM^{HMM}, CRF, ME, and the tree classifier can provide different decisions for each word in the holdout set, we merge their decisions for each word using majority voting. The majority voter is expected to produce accurate results, and not to be biased towards certain types of relational constructions. We test majority voting with a winning majority of three and of four classifiers (results in Table 6). While the result for 3-majority is close to the baseline values shown in Table 3, the 4-majority produces a high accuracy along with a rather low recall. The high precision of the 4-majority is basically desirable for our purposes, as long as the low recall of approximately 52% does not point to a bias in the voting process. Therefore, we study the effect of majority voting on the distribution of the different relational determination modes (see Table 2) of DET_{+R} records from the gold data. When comparing the frequencies of relational types in these samples with their frequencies in the full gold data, a χ^2 test yields a p value of 0.3533 for 3-majority, but of $p = 0.0000056$ for 4-majority voting with notable frequency differences for the determination modes 'possessive pronoun' and 'prepositional phrase with *von*'. To avoid the bias inherent in the 4-majority voting process, we build a new additional training set T^*. This set consists only of those sentences from the holdout set in which 3-majority votes were obtained for all words. When retraining the classifiers, the training set for each fold n consists of the respective $(n-1)$ parts of the gold data and the full set T^*. Table 7 shows that the retraining slightly improves the F score of SVM^{HMM}, but produces a lower recall for ME when compared with the data in Table 3.

Type	P	R	F
SVM^{HMM}	93.26	76.97	84.34
CRF	92.19	71.8	80.73
ME	95.17	53.56	68.54

Table 7: Performance of the classifiers after semi-supervised retraining with the gold data and with 15.326 silver-annotated sentences (10.330 PUM annotations, 3-majority); 30 cross-validations. Cmp. with Table 3

In the third improvement approach, we use neural embeddings of the context words instead of their sparse 1-of-n encodings, and train a neural network on the given binary classification task. We test this approach because the possessive constructions studied in this paper can be interpreted as a kind of frame semantic relation, though on a rather abstract level. Mesnil et al. (2015) report about significant improvements when a combination of recurrent neural networks and pretrained neural word embeddings is applied to a simple form of frame semantic labeling (i.e., slot filling). We use a 750 MB corpus of German newspaper texts to train the neural embeddings. The newspaper corpus is lemmatized using MATE in order to reduce the size of the vocabulary, and then fed into the word2vec software (Mikolov et al., 2013).[2] The input features for the neural network are constructed as follows. For each word w_j, we build

[2]Parameters: hidden size: 300, window size: 5, training iterations: 10.

a vector of subfeatures by concatenating (1) its neural embedding of size 300 that is generated by word2vec, (2) its case information according to MATE, (3) its POS tag according to MATE, and (4) a Boolean flag indicating if the word is terminated by the letter s (which is a frequent genitive ending in German). Features (2) and (3) are added in 1-of-n encodings after removing POS tags that occur less than 100 times in the training corpus. For each word w_i to be classified, these features are collected for w_{i-1}, w_i, and w_{i+1}, and concatenated to form the input vector of the neural network. The resulting concatenated vector is fed into an Elman network.[3] A cross-validation with 30 folds using the same data as for Table 3 and Table 7 produces $P = 79.29, R = 16.67, F = 27.54$. Obviously, recurrent neural networks combined with pretrained neural word embeddings are not really helpful for the given task. This conclusion is supported by a closer inspection of the parameters learned by the CRF model. The highest scoring parameters are induced by rules that operate almost exclusively on POS tags, case information for different context words, or combinations of POS and case information.[4] A full deep-learning pipeline will certainly generate different and better results. However, the problem at hand may lend itself to a shallow solution that is focused on morpho-syntactic features.

None of the three improvement strategies examined in this section substantially improves the performance of the $\text{DET}_{\pm R}$ labeler, while they increase the complexity of the classification workflow at the same time. Therefore, we decided to use the baseline system with binary output for the labeling task described in the next section.

3 Results

The paper aims at examining if and how the syntactic determination modes of nouns in language use are correlated with their semantic lexical determination types. The preceding section has described the classifiers that are used to detect the different syntactic determination modes automatically. This section introduces the data source for the lexical determination types. In addition, it determines the frequencies of congruent and incongruent determination uses for frequent nouns by contrasting the syntactic determination modes that were detected using the automatic labelers with those of the manually annotated lexical determination types.

3.1 Data for lexical determination types

For information about the lexical determination types we had access to manually annotated lexical data that was originally created for internal use in SFB 991 and that contains approximately 23.000 tokens annotated with concept types. The data consists of noun lemmata that occur at least ten times in a corpus consisting of German news, narrative, and scientific texts. Each noun occurrence is disambiguated with respect to the readings given in the standard German dictionary Duden (1997). Finally, each lexical reading that occurs at least once is classified with respect to its lexical type of determination (see Table 1).

To obtain clearer effects, we have selected only those noun lemmata (1) to which only a single lexical determination type had been assigned, and (2) that occured at least 100 times in our corpus (1M sentence extract from the Leipzig corpora, 2005). This filtering process produced a candidate list L_{cand} of 217 noun lemmata that contained 24 functional, 48 individual, 53 relational, and 92 sortal nouns.

3.2 Congruency of determination

The German 1M sentence extract of the Leipzig corpora from 2005 was processed with the $\text{DET}_{\pm U}$ and $\text{DET}_{\pm R}$ labelers from Section 2, skipping all sentences for which MATE could not detect a single root node. Only those records were retained for which the $\text{DET}_{\pm U}$ labeler assigned either DET_{-U} or DET_{+U}, and for which at least three members of the $\text{DET}_{\pm R}$ classifier ensemble agreed (3-majority). Furthermore, all records whose lemmata were not found in L_{cand} were removed. The resulting list L_{res} contained 433.230 noun tokens for the 217 lemmata from L_{cand}.

Table 8 displays the absolute frequencies from L_{res} for the lexical (rows) determination types and the syntactic (columns) determination modes. The data shows on one hand that within one lexical determination type the congruent determination mode is dominant only for sortal $(-U-R)$ and individual $(+U-R)$ nouns. Relational $(-U+R)$ and functional $(+U+R)$ nouns occur more often in DET_{-R} contexts than in DET_{+R}. A first shallow analysis

[3] Architecture: one hidden layer with 100 units, output layer with one unit. Activation functions: Sigmoid for the hidden layer, softmax for the output layer. Sequential gradient descent learning, initial learning rate: 0.005. A validation set was used for early interruption of the learning process.

[4] The five highest scoring rules of the CRF were: pos=N:PUM, m1-pos=PRPOSS:PUM, p1-word=der:PUM, pos=PREP:no-poss, and p1-case=gen:PUM.

	DET_{-U} DET_{-R}	DET_{+U} DET_{-R}	DET_{-U} DET_{+R}	DET_{+U} DET_{+R}
-U-R	**140.706**	87.476	4.716	7.042
+U-R	531	**66.773**	8	58
-U+R	43.504	17.999	**10.243**	5.457
+U+R	15.239	21.264	3.714	**7.674**

Table 8: Absolute frequencies of congruent and incongruent determination types

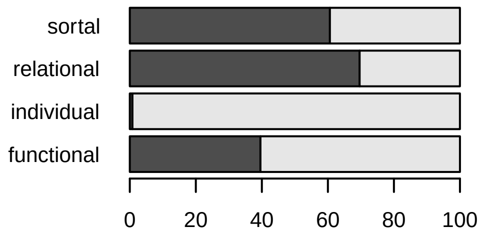

Figure 1: Uses of $DET_{\pm U}$, split up by lexical types. Light grey: Unique uses, dark grey: non-unique uses; x-axis in percent

of the incongruent DET_{-R} uses indicates that the relational argument is frequently implicitly fixed by the context, but not overtly expressed within the noun phrase. On the other hand, the lexical determination type that is congruent with a given determination mode is preferably chosen. There is only one exception to this rule: Among the noun tokens which are used in the DET_{+U-R} determination mode the dominant group is formed by sortal nouns and not by individual nouns with congruent determination type. However, this is due to the fact that sortal nouns are much more frequent than individual nouns.

In addition to the absolute frequencies given in Table 8, the bar plots in Figure 1 and Figure 2 show the percentage distribution of the $DET_{\pm U}$ and $DET_{\pm R}$ features separately for the four lexical determination types. The plots indicate that on average, if one looks at the accumulated frequencies of all nouns belonging to one lexical type, $+U$ nouns show a stronger tendency to occur in DET_{+U} contexts than $-U$ nouns and, $+R$ nouns show a stronger tendency to occur in DET_{+R} contexts than $-R$ nouns. Similar results have already been reported in the smaller manual study (Horn and Kimm, 2014). Thus, there is statistical evidence for the hypothesis that the lexical determination type of a noun influences its use in actual determination modes with a tendency towards congruent determinations.

However, this influence is not strong enough to determine the lexical determination type of a single noun lemma by its frequency distribution over the different types of determination modes in actual language use. Figure 3 shows all 217 lemmata from L_{cand} grouped by their lexical determination types and placed by their individual distributions in $DET_{\pm U}$ / $DET_{\pm R}$ contexts. The plot demonstrates that, apart from the individual nouns, which mainly consist of named entities, all other noun types spread over a rather big region instead of neatly gathering in one corner.

The picture is especially blurred for relational and functional nouns, which occur less frequently in relational determination mode than expected from their underlying lexical types. This result, which partially coincides with the findings reported in Horn and Kimm (2014), points to the influence of other linguistic factors such as anaphoric uses that will be examined in a follow-up study.

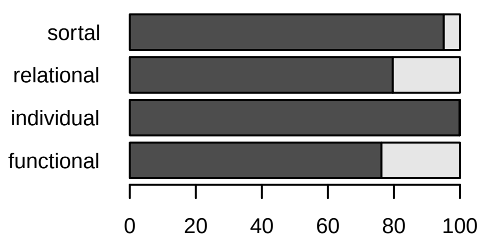

Figure 2: Uses of $DET_{\pm R}$, split up by lexical types. Light grey: Relational uses, dark grey: non-relational uses

Acknowledgement

Research for this paper was funded by a grant of the DFG (SFB 991, project C02).

References

Yasemin Altun, Ioannis Tsochantaridis, and Thomas Hofmann. 2003. Hidden Markov Support Vector Machines. In *Proceedings of the Twentieth International Conference on Machine Learning*.

Chris Barker. 1995. *Possessive Descriptions*. CSLI Publications, Stanford.

Chris Barker. 2011. Possessives and relational nouns. In *Semantics: An International Handbook of Natural Language Meaning*, chapter 45, pages 1108–1129. de Gruyter.

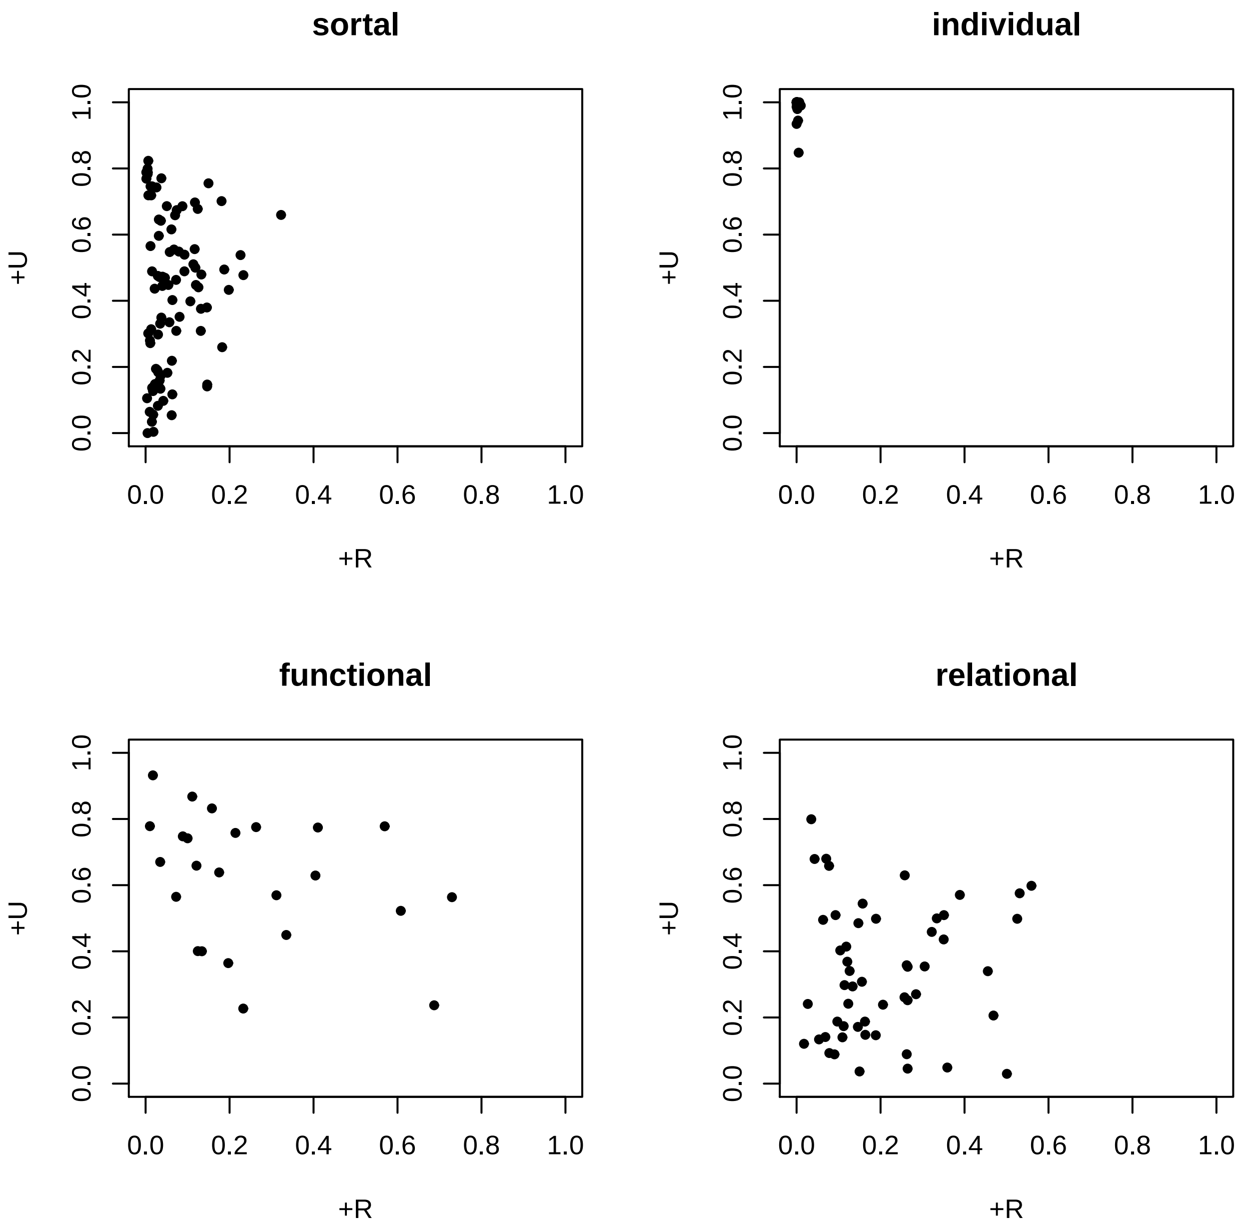

Figure 3: $\mathrm{DET}_{\pm U}/\mathrm{DET}_{\pm R}$ distribution plots for 217 lemmata grouped by their lexical determination

Bernd Bohnet and Joakim Nivre. 2012. A transition-based system for joint part-of-speech tagging and labeled non-projective dependency parsing. In *EMNLP-CoNLL*, pages 1455–1465.

Duden. 1997. *Duden Universalwörterbuch A-Z*. Bibliographisches Institut und Brockhaus, Mannheim, 3 edition. Electronic version.

Kari Fraurud. 1990. Definiteness and the processing of noun phrases in natural discourse. *Journal of Semantics*, (7):395–433.

Oliver Hellwig and Wiebke Petersen. 2014. Detecting relational constructions in German texts automatically. In Josef Ruppenhofer and Gertrud Faaß, editors, *Proceedings of the 12th edition of the KONVENS conference*, pages 40–47.

Christian Horn and Nicolas Kimm. 2014. Nominal concept types in German fictional texts. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald, and Wiebke Petersen, editors, *Frames and Concept Types. Applications in Language and Philosophy*, volume 94 of *Studies in Linguistics and Philosophy*, pages 343–362. Springer Verlag.

Per Anker Jensen and Carl Vikner. 2004. The English pre-nominal genitive and lexical semantics. In J.Y Kim, Y.A. Lander, and B.H. Partee, editors, *Possessives and beyond: semantics and syntax*, number 29 in University of Massachusetts Occasional Papers in Linguistics. GLSA Publisher, Amherst.

John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional Random Fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, pages 282–289.

Sebastian Löbner. 2011. Concept types and determination. *Journal of Semantics*, 28:1–55.

Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, and Geoffrey Zweig. 2015. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.

Johanna Nichols. 1992. *Linguistic Diversity in Space and Time*. The University of Chicago Press, Chicago.

Malvina Nissim. 2004. Lexical information and choice of determiners. In J.Y Kim, Y.A. Lander, and B.H. Partee, editors, *Possessives and beyond: semantics and syntax*, number 29 in University of Massachusetts Occasional Papers in Linguistics, pages 133–152. GLSA Publisher, Amherst.

Albert Ortmann. 2015. Uniqueness and possession: Typological evidence for type shifts in nominal determination. In Martin Aher, Daniel Hole, Emil Jeřábek, and Clemens Kupke, editors, *Logic, Language, and Computation*, volume 8984 of *Theoretical Computer Science and General Issues*, pages 234–256. Springer, Berlin, Heidelberg.

Barbara Partee and Vladimir Borschev. 2003. Genitives, relational nouns, and the argument-modifier ambiguitiy. In Ewald Lang, Claudia Maienborn, and Cathrine Fabricius-Hansen, editors, *Modifying Adjuncts*, pages 67–112. Walter de Gruyter, Berlin.

Barbara Partee. 1986. Noun phrase interpretation and type-shifting principles. In J. Groenendijk, D. de Jongh, and M. Stokhof, editors, *Foundations of pragmatics and lexical semantics*, pages 115–143. Foris, Dordrecht.

Anna N. Rafferty and Christopher D. Manning. 2008. Parsing three German treebanks: Lexicalized and unlexicalized baselines. In *Proceedings of the Workshop on Parsing German*, pages 40–46, Stroudsburg, PA, USA. Association for Computational Linguistics.

Adwait Ratnaparkhi. 1998. *Maximum Entropy Models for Natural Language Ambiguity Resolution*. Ph.D. thesis, University of Pennsylvania.

Renata Vieira. 1998. *Definite description processing in unrestricted text*. Ph.D. thesis, University of Edinburgh.

C. Vikner and P. A. Jensen. 2002. A semantic analysis of the English genitive. Interaction of lexical and formal semantics. *Studia Linguistica*, 56(2):191–226.