Proceedings of the Computational Linguistics-Applications Conference pp. 93–98 ISBN 978-83-60810-47-7

Expanding a Corpus of Closed-World Descriptions by Semantic Unit Selection

Marcus Uneson
Department of Linguistics
Centre for Languages and Literature
Lund University
Sweden
Email: marcus.uneson@ling.lu.se

Peter Juel Henrichsen
Center for Computational Modeling of Language
Copenhagen Business School
Denmark
Email: pjh.isv@cbs.dk

***Abstract*—We present a method for the controlled expansion of a corpus based on "semantic unit selection": units from a speech database are chosen not for closeness to an acoustic or phonetic target, but rather for their semantic content. While unsuited for general speech synthesis, it may be useful for restricted domains.**

We provide an application example from our current line of research: induction of lexical structure (i.e., acoustic, combinatorial, and semantic information) from unanalyzed recordings of informants describing small, closed-world scenarios. Here, semantic unit selection permits existing descriptions to be freely paraphrased and rearranged into new ones. The amount of redundancy can be parameterized, offering a way to control the difficulty to the task.

The method is not dependent on the original scene described but can take a formal description of a new scene as input, or even enumerate all scenes describable by the data (along with descriptions).

***Index Terms*—semantic unit selection, Definite Clause Grammar, Danish, probabilistic generation, simulated speech data**

I. Introduction

ANNOTATED speech data are a prerequisite for most speech research. Unfortunately, the annotation process of the data is tedious, demanding, and expensive. The more complex the annotation, the more demanding and expensive it turns. A skilled typist can produce an orthographic transcription of a recording of duration D in about 4 x D, with no special skills required; time-aligned transcriptions are slightly more expensive. Phonetic transcription, by contrast, is far more time-consuming (on the order of 100 x D or more), requires training, and even so usually produces less consistent results [2].

On a related note, the work needed for the data collection itself may be negligible if you happen to be interested only in widely spoken languages and not too fussy about the actual content, but it can be very demanding if your interest mainly lies in small languages or very specific tasks.

Corpus production is thus work-intensive and work-extensive. Many attempts have been made to remedy this unfortunate situation, ranging from tools for automatic or semi-automatic annotation to initiatives facilitating resource sharing by improving infrastructure and metadata.

In this paper, we will concentrate on yet another strategy: the expansion of an existing corpus by *recombination* of utterances, where we define utterances to be phonetically independent carriers of meaning with respect to some given domain. We have termed the method "semantic unit selection", by analogy with the well-known technique of (phonetic) unit selection, where prerecorded speech units are selected and concatenated based on their calculated acoustic or phonetic distance from a target.

The parallels should not be drawn too far; our approach is much simpler and useful only for restricted domains. Furthermore, recombination will of course not add information absent in the original corpus. "Corpus expansion", as used here, is rather to be thought of as a way of preparing data with a controlled degree of difficulty for learning tasks. As exemplified below, we have successfully used the technique for our current research into the learning of lexical structure from speech (see Section III) using no external resources except a formal description of the (necessarily restricted) domain.

The paper is organized as follows. In Section II, we describe our implementation of semantic unit selection and give examples of its application to (a specific task of) the DanPASS corpus. In Section III we review our method for lexical learning from raw speech data, and present the results of applying it to the recombined descriptions of Section II. Section IV briefly presents SMALLWorlds, a multilingual expansion of the DanPASS corpus for which we believe semantic unit selection will be an important helper technique. Section V concludes with some comments on extensions and possible future uses.

II. Semantic unit selection

Semantic unit selection works as follows. First, the data is segmented into units which are large enough to be rearrangeable without regards to local phonetic phenomena, such as coarticulation and assimilation. With this definition, an automatic segmentation method based on crude measures such as pause durations works well

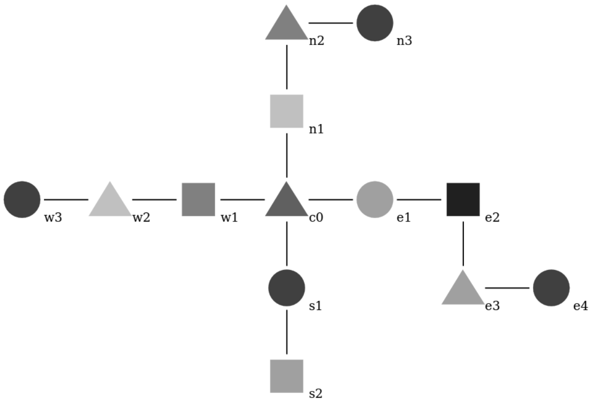

Fig. 1. *The DanPASS geometric network, converted to grayscale and with node IDs added (e1 for first node on Eastern branch, etc). The shapes of the original are in six different colours: yellow (nodes w2, n1), red (w1, n2), green (s1, e4, w3, n3), blue (s2, c1, e3), brown (e2), pink (c0)*

start: set some shape s to be the current shape c
step: if there is some previously undescribed neighbour n of c:
1) describe n with c as point of departure (i.e., say something about the shape n, its colour and its form; also, if deemed necessary, how to get to it from c)
2) set $c = n$
3) goto step

jump: else if there is some previously described shape s which has undescribed neighbours:
1) set $c = s$
2) goto step

stop: else stop

Fig. 2. *Typical human strategy for describing a network such as that in Figure 1*

enough. Next, each utterance is annotated with respect to its semantic content. Finally, utterances are selected and compositionally combined to larger units.

In this section we first describe one appropriate data set, taken from the DanPASS (Danish Phonetically Annotated Spontaneous Speech) corpus [1].[1] After that, we describe our semantic annotation scheme for this task, expressed in first-order predicate logic: specifically, the Definite Clause Grammar (DCG) formalism of the Prolog language [6], and how these annotations can be used for generation.

A. The DanPASS Geometric Network task

The "Geometric Network" part of the DanPASS corpus was inspired by a similar task used for the study of prosody in spontaneous speech ([7], see also [5]). In the DanPASS version, 18 subjects were asked to describe a small, closed world: a network of coloured shapes, a grayscale rendering of which is shown in Figure 1.[2]

Humans facing such a task typically employ an algorithm similar to that shown in Figure 2 (although often unconsciously). As can be seen, there is non-determinism present in the order the shapes are visited, indicated by "some" present at the choice points: start, jump, step.

In the specific case of DanPASS, non-determinism manifests itself as follows. For the **start** item, the point of departure was fixed to node s2 by the instructions. As for the **jump** item, the network only contains one junction, so the strategy stated in Figure 2 will not introduce any further choices. A variant strategy, possible but less common, omits "previously described"; this permits jumps to any undescribed node. However, for the most part, non-determinism is decided by the **step** item. After having begun with the southern branch according to the rules and reached the center, the description may order the remaining three branches in six ways. If we use SWEN to abbreviate the path [s2, s1, c0, w1, w2, w3, c0, e1, e2, e3, e4, c0, n1, n2, n3]), then the others are SWNE, SEWN, SENW, SNWE, SNEW.

A fourth point of non-determinism, of course, lies in the many possible phrasings of a description: there is a practically infinite number of ways to put into words a vague instruction such as "say something about the colour and shape". However, at the level of abstraction we are dealing with throughout this paper, the exact wording is below the level the horizon of interest. Here lies the by far most fertile source of indeterminacy, yet dealt with in a principled way in our semantic annotation scheme.

B. Semantic annotation of the DanPASS Geometric Network

As it turns out, the strategy depicted in Figure 2 is enough to characterize the DanPASS descriptions in all but two cases. It is not difficult to formalize. In Danish, as in any human language, we would expect to be able to describe a `shape` with respect to `form` and `colour`;[3] to mention a shape as `given` or `new`; and to express a direction or location, `dir`, from some point of reference. It is also useful to be able to express self-corrections, `corr`, and to have some way of representing a *fill*, by which we mean anything without relevance to the description (as

[1] http://www.cphling.dk/˜ng/danpass_webpage/danpass.htm

[2] The scene of the task is somewhat reminiscent of the Blocks World scenarios of early AI research [8] – specifically designed to be describable by a small number of largely connotation-free content words for colours, shapes, and spatial relations.

[3] It is useful to have separate names for the geometric form of an object on the one hand, and this form combined with a colour on the other. In this paper, we use `form` and `shape`, respectively: thus, "square" is a `form`, "brown square" is a `shape`.

formalized above). The only thing we care about in a fill is which of the other categories it embellishes.

All in all, we might get around a dozen DCG categories. Our current set is given and exemplified in Table I. In the table, the placeholder `[this]` appearing on the right-hand side of the DCG rules stands for the terminal at hand, to be filled in with actual language data from the corpus – be it as orthographic transcription (as shown in the table), phonetic transcription, sound, or some other representation. As can be seen, terminals may be recursively combined with non-terminals. The first argument to the DCG rules maintains a semantic representation in typical Prolog-style [6]; the second, when present, may be used in generation (see below).

The annotation scheme described covers most data, but not all. For one thing, it cannot easily represent utterances describing several shapes at a time, such as *til en gul trekant og igen et skridt til venstre til en grøn cirkel* 'to a yellow triangle and again a step to the left to a green circle'. Such utterances are currently ignored. However, for the self-paced monologues of the task at hand, they are unusual (e.g, in the 959 utterances of the Geometric Network, there are only two instances). More generally, some speakers and some speaking styles tend to offer fewer recombination units than others, which decreases compositionality and makes them less useful for semantic unit selection.

C. Generation

If we now combine session-level annotations as described with a few top level categories for descriptions, we can generate novel descriptions. We can either provide the semantics argument as input (if we already have a ground formalization of a network for which we also want linguistic descriptions), or we can have it instantiated from data (if we want to generate language along with formalizations of scenes). In any case, the output will be a string of terminals, each representing (part of) an utterance of the original data.

Some of the additional categories will be the same for all descriptions, while others are language-specific, encoding typological properties such as word order and expression of information structure. We will omit the former for lack of space, but the current language-specific ruleset for Danish is given below in its entirety. Basically, it encodes the well-known fact that in Danish nominal phrases has Adjective-Noun word order and that information structure is primarily expressed by the definite/indefinite distinction (at least as far as this particular corpus is concerned).

```
new(S) --> shape(S, [indef]).
given(S) --> shape(S, [def]).
shape(shape(F,C), [Species]) -->
  colour(C, [Species]), form(F, [Species]).
```

1) On ungrammaticality: The reader may wonder how such a crude representation treats syntactic constraints at the concatenation points. Surely the concatenated utterances will be ungrammatical or contain repetitions? There are two answers to this question. One (arguably the traditional one) is yes, but the semantic annotation can easily be complemented by an argument which remedies that. For instance, it could specify local syntactic constraints for agreement in smaller phrases. This is one of the possible (parallel) uses of the second argument to `form, colour, shape` shown in Table I; higher-level categories can be treated analogously.

Our main answer, however, is another: yes, they will be ungrammatical in the Chomskyan sense, but not noticeably more so than were the originals. The data – elicited, self-paced monologue as it is – contain many examples of disfluencies, hesitations, repetitions, self-corrections, and lacking agreement. All of them are represented in Example 1 (original data) and 2 (recombined); commas mark utterance boundaries.

(1) *så kommer der en blå cirkel og så*
then comes (there) a blue circle and then
kommer der en, dernæst kommer der
comes (there) a, after that comes (there)
så en lilla, nej en brun, en brun firkant
then a pink, no a brown, a brown square

(2) *og til, og i lige linje, og så*
and to, and in (a) straight line, and then
kommer der en rød, grøn, cirkel, ja
comes (there) a red, green, circle, yes

As can be seen, the syntax is highly relaxed. Incoherent as these utterances may seem in writing, the auditory impression is actually not peculiar at all (for neither of them). The slightly hesitating style, present in the original recording and inherited in the recombined version, is well-known to the ear, albeit not to the eye.

In our experience, any impression of unnaturalness does not primarily come from agreement violations (and certainly not from repetitions). More obvious are errors coming from inappropriate modeling of new and given information, especially when resulting in prosodic mismatches. This is another and more interesting type of congruence, and the main use of the second argument of `shape, colour, form` in Table I.

2) Parameterizing redundancy: Usually, there will be many ways to expand a non-terminal. So far, we have said nothing about how to select between competing expansions – conceptually, we have flipped a coin whenever a choice was needed.

More interesting, however, is to guide the generation process in a systematical way. In particular, by weighting rules which involve fills, we get a principled means of controlling redundancy. The current implementation employs an intermediate unit `line` which is to expand the categories listed in Table I. It takes a weight as its first argument.[4]

```
line(Wstart, start(Start)) --> start(Start).
```

[4]In the interest of reducing clutter, we ignore here (and in Table I) that some fills may only go before or after the phrase they embellish; the extension to cater for that is straightforward.

TABLE I

Sample semantic annotations of the DanPASS Geometric Network task. DCG abbreviations: fm:form (tr:triangle:, sq:square:, cl:circle); col:colour (br:brown, bl:blue, pk:pink, gr:green); [sS]h:[sS]hape; ind:indefinite; def:definite; Gn:Given; rt:right

DCG annotation	Danish examples	English translation
`start(sh(fm(sq),col(bl))) --> [this].`	man starter forneden med en blå firkant	you start at the bottom with a blue square
`start(Sh) --> [this], shape(Sh, [ind]).`	du starter nederst på papiret med en	you start at the bottom of the paper with a
`form(fm(tr), [ind]) --> [this].`	trekant	triangle
`colour(col(gr), [ind]) --> [this].`	grøn	green
`shape(sh(fm(sq), col(br)), [ind]) --> [this].`	brun firkant	brown square
`new(Sh) --> [this], shape(Sh, [ind])`	lægger du en	you put
`given(sh(F,col(pk))) --> [this],form(F,[def]).`	den lilla	the pink
`dir(rt) --> [this].`	her drejer stien til højre; vandret til højre; så lægger du til højre for den	here the path turns right; horizontally to the right; then you put to the right of the
`step(step(_G, dir(up), sh(fm(cl),col(gr)))) --> [this].`	oven over har vi en grøn cirkel	above we have a green circle
`step(step(sh(fm(tr),col(bl)), dir(rt), S --> [this],shape(S, [ind]).`	og til højre for den blå trekant er der en	and to the right of the blue square there is a
`jump(Gn) --> [this], shape(Gn, [def]).`	og man hopper tilbage til den	and you jump back to the
`stop --> [this].`	nu er jeg færdig; ja ja det var vist det; ja; sådan tror jeg jeg vil beskrive det	now I am finished; yeah that's about it; yeah; that's how I would describe it
`fill(start) --> [this].`	nå; jeg er klar nu; jeg skal beskrive et netværk af brikker her	well; i am ready now; I will describe a network of pieces here
`fill(step) --> [this].`	okay; det er jeg lidt i tvivl om	okay; I wouldn't know about that
`fill(stop) --> [this].`	hele netværket har altså form som en slags kors	the entire network thus has the shape of a kind of cross
`fill(_) --> [this].`	ehm; hvad hedder det	ehm; what's it called
`corr(_, sh(f(tr),col(pk))) --> [this].`	undskyld til den lilla trekant	excuse me to the pink triangle

```
line(WstartF, start(Start)) --> fill(start), start(Start).
line(Wstep, step(Step)) --> step(Step).
line(WstepF, step(Step)) --> fill(step), step(Step).
...
```

Greater weights assigned to fills, `WxxxF`, will now generate wordier descriptions. Similarly (although we have not yet tested that in practice), another parameter can control the amount of self-corrections.

Since it is easier to remove redundancy and corrections than to add them, data from rambling speakers are actually better suited for this purpose. In one way, this technique is a small step towards a parameterization of speaker personality. At any rate, it is very useful for preparing data with controlled degree of difficulty for learning tasks.

III. Learning lexical structure

The current work is an offspring of a long-term project on language learning initiated at CBS, Copenhagen. The embedding project is aimed at developing computational methods for induction of lexical structure from unanalyzed speech recordings, simulating human language acquisition. By "lexical structure", we refer to a totality of acoustic, combinatorial, and semantic information. The experiments reported in the following take phonetic transcriptions as input; however, our final goal is learning from unannotated speech recordings.

Among other corpus materials, the project has used DanPASS extensively due to its attractive mixture of spontaneity and control (cf. [7]). The informants know exactly what task to solve, yet they are free to express themselves in their natural style often including self-corrections, incomplete constituents, irrelevant information, and even meta-language (see Section II-C1 and Example 3 below). Arguably, any lexical induction model with an ambition of psychological realism must be able to cope with naturalistic speech. In this line of work, we have developed semantic unit selection as a partial solution to a data preparation problem: our learning algorithms seemed to work well when applied to the 18 DanPASS sessions (deducing correctly the Danish names for the colours and shape terms in Figure 1), but to establish a solid verification, we needed more fine-grained control of the difficulty of the learning task.

In the following, we give a very compressed overview of the learner; for further details, we refer to [3]. Data from DanPASS are presented to the learning algorithm in the form of phonetic transcripts, as in Example 3.

(3) T: jaβəˈgønʔɐ
D: jeg begynder
E: I begin

#93 n-gram analysed: [t ʀ æ k a n ʔ d]
1.000000 [t ʀ æ k a n ʔ d]
0.837132 [f i ʁ k a n ʔ d]
0.727861 [l e l a t ʀ æ k a n ʔ d]
0.646050 [s i ʁ g l]
0.629778 [t ʀ æ k a n ʔ t]
0.625563 [d ɛ n ʔ]
0.614339 [f i ʁ k a n ʔ t]

Fig. 3. *Sample from Siblings & Cousins log (stage 1 and 2, see text). Analysed n-gram #93:* [t ʀ æ k a n ʔ d] *(the Danish word for triangle). Listed: high-scoring n-grams, sorted by context selectional similarity with* [t ʀ æ k a n ʔ d]. *Male speaker, session ID m_014g*

triangle: [t ʀ æ k a n d]
square: [f i ʁ k a n d]
circle: [s i ʁ g l̩]
blue: [b l ɔ ʔ]
green: [g ʀ ɶ n ʔ]
red: [e n ʀ œ ð ʔ]
yellow: [g u l]
PATH: [s2 s1 c0 w1 w2 w3 e1 e2 e3 e4 n1 n2 n3]

Fig. 4. *Sample from learning log (stage 3, see text; cf. also Fig. 3), showing the deduced lexemes and the associated path*

'neːði'bɔnʔn̩ɛfi'guʔɐnmɛn̩'blɔʔ T: 'fiʁkanʔd
nede i bunden af figuren med en blå D: firkant
in the bottom of the figure with a blue E: square
sʌgɒjad'sgʀid'ʌp ʌ'tɑʔeŋ'gʀɶnʔ T:
så går jeg et skridt op og tager en grøn D:
and I take a step upwards and take a green E:
ʌdɐ'kʌmʔɐŋ'gʀɶnʔ'siʁgl̩
og der kommer en grøn cirkel
and there comes a green circle

For the convenience of the reader the transcription (T) has spaces added corresponding to pauses, and is given also in Danish orthography (D) and glossed in English (E). However, the learner only sees the unsegmented and unnormalized stream of phones, including any syntactic irregularities, self-negotiations and self-corrections, hesitations, and even factual description errors.

When used as experimental data, the string of phones (excluding marking of stress and vowel length) is processed in three stages:

1) all frequently occurring n-grams, such as [ililinjɪ], are identified;
2) the n-grams are arranged in sets of three based on distributional similarity, such as [[f i ʁ k a], [i l i l i n j ɪ], [t ʀ æ k a n d]];
3) the triplets are piped to the inference module as lexical hypotheses.

For identification and arrangement of n-grams (steps 1 and 2), we employed the algorithm Siblings & Cousins ([4]). This algorithm exploits the fact that two words with complementary semantics (like two distinct color terms, say *blue* and *green*) tend to prefer similar contexts at their right and left edges. For instance, consider a pair of corpus instances *a blue circle* and *a green circle*, both quite frequent in DanPASS. In this case, of course the color terms share the context *a _ circle*. Quantifying over all n-gram candidates and all their respective context functions, the Siblings & Cousins algorithm produces analyses as the one shown in Figure 3.

Based on their left and right context selection functions, the n-grams [t ʀ æ k a n ʔ d] and [t ʀ æ k a n ʔ d] are thus – not surprisingly – judged to be similar to a degree of 100%. More interestingly, the n-grams [t ʀ æ k a n ʔ d] and [f i ʁ k a n ʔ d] score 83.7%, meaning that these two n-grams do indeed prefer the same contexts to a high degree. This is satisfactory, the two n-grams representing the Danish words for triangle and square, respectively. As can be seen, some semantically neutral variations in pronunciation is also detected (e.g. the t/d allophones).

Note in Figure 3 the n-gram [l e l a t ʀ æ k a n ʔ d]. This n-gram corresponds to a compound expression (purple triangle), but was nevertheless picked by the algorithm as a possible semantic atom (based on its contextual similarity with a shape term proper). This judgment is actually not surprising, given the fact that color [lela], purple, is represented in the layout (Figure 1) by one object only, creating a strong cohesion effect. This interplay between atomic and compositional semantic readings of compound words is of course well-known in human discourse too: the term *red herring* may occasionally be used to refer to a herring which happens to be red; but usually its meaning is atomic.

Turning now to step 3: for each triplet, the inference engine searches for a division of the entire transcription into 13 subsections (corresponding to the 13 objects in Figure 1), each containing a triplet element (the one in the example would thus be rejected, [ililinjɪ] not being a shape name). On success, the 13-section is checked for consistency with the human description strategies (Figure 2), and a corresponding colour mapping is deduced (Figure 4).

As the deduction regime is based on streams of phones rather than delimited tokens, the deduced lexicons often contain unusual segmentations. Notice thus that the color name for red is rendered as [e n ʀ œ ð ʔ] (including the indefinite article [en]). We have performed 30 test runs using a variety of recombined DanPASS sessions. The correctness of the learning algorithm was verified in all cases; however, the delimitations of each lexeme vary quite a lot. The deduced terms for red thus include: [ʀ œ ð ʔ], [n ʀ œ ð ʔ], [e n ʀ œ ð ʔ], [ɑ e n ʀ œ ð ʔ], [ɾ ɑ e n ʀ œ ð ʔ], and even [ʌ w s ə k ʌ m ʔ ɐ ɾ ɑ e n ʀ œ ð ʔ]. All are of course equally well-formed, given our data-driven approach.

IV. The SMALLWorlds corpus

For the comparison, the experiments reported here have generated and processed files which describe the same network as do the DanPASS originals. Of course, neither Danish nor the network itself are universal constants. As for the language, a cross-linguistic perspective on the learning task is clearly valuable. We are currently in the process of collecting similar descriptions in a large selection of languages other than Danish into a corpus called SMALLWorlds (Spoken Multilingual Accounts of Logically Limited Worlds). At the time of writing, it comprises more than 200 speakers in around 50 languages from several major language families.

As for the network, there is really nothing in the corpus expansion method which requires the recombined descriptions to refer to the same figure that was used to elicit the original data – the network to be described could just as well be given as input. This adds another important dimension along which the difficulty of the learning task may be parameterized.

We could even have the network only partially specified, in which case we will receive an enumeration of all paraphrases of all descriptions of all consistent networks that the original data permits, together with their formal descriptions. By extrapolation, we could (at least in principle) enumerate the transitive closure of the input description: all paraphrases of all descriptions of all networks which actually can be described from the original data, along with the associated formal descriptions and sound files. At any rate, we expect that semantic unit selection will be an important tool for exploiting the SMALLWorlds corpus.

V. Conclusion and outlook

We have demonstrated how recombination by semantic unit selection allows a principled way of manipulating original data, which may create new descriptions of formally describable scenes, with a controlled measure of redundancy. Expanding on this perspective, we suggest that the recombination regime be explored in other settings where generated, yet naturally sounding discourse is in demand. An obvious case in point are computer games. Modern game interfaces often provide an impressive visual experience. The auditory scene, by contrast, is usually restricted to rather dull collections of canned utterances, non-verbal sound tokens, and musical tapestry. We speculate that semantic unit selection principles may provide the base for much more realistic speech.

As is well known from the speech synthesis industry, exact repetitions of reproduced speech tokens make people experience artificial voices as 'highly unnatural'. Of course, no natural voice would, or could, repeat itself exactly. One answer to the naturalness issue is to simply apply a small amount of random variation to the synthetic speech production, often with a significant improvement in experienced naturalness. We believe that an even greater improvement can be achieved in a semantic unit selection regime, allowing simple-minded variation techniques to be replaced by intelligent speech design.

References

[1] N. Grønnum, "A Danish phonetically annotated spontaneous speech corpus (DanPASS)," *Speech Communication*, vol. 51, pp. 594–603, 2009.
[2] M. Hasegawa-Johnson, Lecture notes from minicourse on Speech tools, http://www.isle.illinois.edu/sst/courses/minicourse/2009/, 2009.
[3] P. J. Henrichsen, "Den lilla trekant – learning Danish shape and color terms from scratch," in *Linguistic theory and raw sound*, ser. Copenhagen Studies in Language, P. J. Henrichsen, Ed., 2010, no. 40, pp. 27–44.
[4] P. J. Henrichsen, "Fishing in a speech stream, angling for a lexicon," in *NODALIDA-2011*, May 2011, Riga University.
[5] W. Levelt, *Speaking: From Intention to Articulation*, MIT Press, Cambridge, MA., 1989.
[6] F. Pereira and D. Warren, "Definite clause grammars for language analysis–a survey of the formalism and a comparison with augmented transition networks," *Artificial intelligence*, vol. 13, no. 3, pp. 231–278, 1980.
[7] M. Swerts and R. Collier, "On the controlled elicitation of spontaneous speech," *Speech Communication*, vol. 11, no. 4-5, pp. 463–468, 1992.
[8] T. Winograd, "Procedures as a representation for data in a computer program for understanding natural language," Ph.D. dissertation, MIT, 1971, Technical Report AITR-35.