

ADAM PRZEPIÓRKOWSKI

Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland
Institute of Philosophy, University of Warsaw, Poland
adamp@ipipan.waw.pl

TOWARDS A LINGUISTICALLY-ORIENTED TEXTUAL ENTAILMENT TEST-SUITE FOR POLISH BASED ON THE SEMANTIC SYNTAX APPROACH

Abstract

The aim of this programmatic position paper is to show that the semantic syntax tradition of Polish linguistics associated with the name of Stanisław Karolak may be a basis for the development of a taxonomy of entailment types and a corresponding test-suite of entailment examples. The article also puts forward some initial desiderata for such a test-suite.

Keywords: semantic syntax; textual entailment; Polish

1. Introduction

The task of recognising textual entailment (RTE; Dagan, Roth, Sammons, & Zanzotto, 2013) consists in finding out whether the information contained in one text is entailed by that given in another. Let us have a look at an example from Dagan et al., 2013, p. 8:

- (1) **T:** The purchase of Houston-based LexCorp by BMI for \$2Bn prompted widespread sell-offs by traders as they sought to minimize exposure. LexCorp had been an employee-owned concern since 2008.

H1: BMI acquired an American company.

H2: BMI bought employee-owned LexCorp for \$3.4Bn.

H3: BMI is an employee-owned concern.

Given the original text T above, information in hypothesis H1 is entailed by it, information in H2 is contradictory with it, and information in H3 stands in no entailment relation with it.

Textual entailment (TE) corpora contain pairs of sentences together with information about whether they stand in the entailment relation. For example, such

a TE corpus for English may contain the triples $\langle T, H1, \text{yes} \rangle$ (i.e., T does entail H1), $\langle T, H2, \text{no} \rangle$ and $\langle T, H3, \text{no} \rangle$ (i.e., T entails neither H2, nor H3). Instead of this binary entailment classification, some corpora use a tertiary classification, to distinguish pairs such as $\langle T, H2 \rangle$, where the two texts are contradictory, from pairs such as $\langle T, H3 \rangle$, where neither entailment nor contradiction is observed.

Such textual entailment corpora are an increasingly important kind of linguistic resource, as they are used for testing — and, to some extent, training — programs which recognise textual entailment; such programs are important modules in some common Natural Language Processing (NLP) tasks such as Question Answering, Information Extraction and Automatic Summarisation. For example (Dagan et al., 2013, p. 11), given the question *Who painted “The Scream”?* and the following text snippet found via Information Retrieval methods as possibly giving an answer to this question: *Norway’s most famous painting, “The Scream” by Edvard Munch. . .*, an RTE module may be used to verify that this text snippet indeed entails the answer: *Edvard Munch painted “The Scream”*.

Many such application-oriented TE corpora have been created for English since mid-2000s, especially, within the so-called RTE shared tasks.¹ One of these corpora, RTE3 created within the third RTE shared task (Giampiccolo, Magnini, Dagan, & Dolan, 2007), has subsequently been translated into German and Italian, and is currently being translated into Polish within the part of the CLARIN-PL project (<http://clarin-pl.eu/en/>) carried out at the Institute of Computer Science, Polish Academy of Sciences (IPI PAN).² Since such TE corpora are developed with the usefulness for particular applications (such as Question Answering) in mind, it is reasonable to construct equivalent corpora of this kind for multiple languages, as this makes it possible to compare RTE modules and their roles in respective tasks cross-linguistically.

On the other hand, entailment captured in such corpora may require diverse kinds of knowledge and reasoning capabilities, but the standard RTE corpora give no indication of what kind of inference steps are needed to recognise entailment in particular examples. For example, in order to recognise that in (1) above H1 follows from T, one must use some world knowledge (namely, that Houston is situated in America) and some linguistic knowledge (namely, that the noun *purchase* represents the same semantic relation as the verb *acquire*). Moreover, some entailments require purely logical reasoning (as in the classical syllogism in which the conclusion that Socrates is mortal is deductively inferred from the premises that all men are mortal and that Socrates is a man). As these recently developed TE corpora contain no information about the kinds of knowledge and reasoning involved in the entailment, they may be successfully used for a *quantitative* evaluation of RTE modules (the accuracy of the module with respect to the test corpus), but not for the *qualitative* evaluation of encyclopedic, linguistic or logical resources such modules are built on.

There exists an earlier resource of a similar kind, created within the FraCaS project (Cooper et al., 1996),³ which does concentrate on one aspect of inference,

¹See http://aclweb.org/aclwiki/index.php?title=Textual_Entailment_Resource_Pool.

²A similar resource, taking also into account relations other than textual entailment, is also being created within CLARIN-PL at the Wrocław University of Technology.

³The resource was converted to the XML format by Bill MacCartney and made available at

namely, a manually constructed test-suite of inferences verifying semantic properties of natural language words and constructions which correspond to the logical notions of quantification, conjunction, etc., and which represent grammatical phenomena such as anaphora, ellipsis, comparatives, tense, aspect, etc. For example, the following pairs (Cooper et al., 1996, pp. 69, 71) reflect the monotonicity properties of some generalised quantifiers (Mostowski, 1957; Barwise & Cooper, 1981):

(2) entailment:

T: At most ten commissioners spend time at home.

H: At most ten commissioners spend a lot of time at home.

(3) contradiction:

T: Neither commissioner spends time at home.

H: Either commissioner spends a lot of time at home.

(4) no entailment relation:

T: At least three commissioners spend time at home.

H: At least three commissioners spend a lot of time at home.

The research plan outlined in the next section bears some affinity to that of Cooper et al. (1996).

2. Aims and Related Work

The goal of this paper is to put up for discussion a research programme aiming at the development of a linguistically-informed textual entailment test-suite for Polish. We do not call the planned resource a corpus, as — apart from naturally occurring attested sentences — it will contain manually constructed entailment pairs. This is necessitated by the main assumption behind the planned research, namely, that its results should make it possible to evaluate RTE modules *qualitatively*, i.e., that the resulting resource will help identify the kinds of inference phenomena which are not satisfactorily handled by such modules.

For example, one such inference phenomenon is related to nominal hyperonymy: if N_1 is a hyperonym of N_2 (e.g., FRUIT is a hyperonym of APPLE) and V is an intransitive verb, then “an N_2 V.PST”⁴ (e.g., *an apple disappeared*) entails “an N_1 V.PST” (e.g., *a fruit disappeared*), but not the other way round. Conversely, when a is replaced by *all*, “all N_1 s V.PST” (e.g., *all fruits disappeared*) entails “all N_2 s V.PST” (e.g., *all apples disappeared*), but not the other way round. Another phenomenon is diathesis, e.g., passivisation: for any noun phrases NP_1 and NP_2 and any transitive verb V , the passive “ NP_2 was V.PASS by NP_1 ” (e.g., *an apple was eaten by John*) is equivalent to the active “ NP_1 V.PST NP_2 ” (e.g., *John ate an apple*). Hence, in a linguistically-oriented TE test-suite, each pair should be labelled with information whether the reasoning needed to establish (or disprove) entailment

his web page (<http://nlp.stanford.edu/~wcmac/downloads/>).

⁴Morphosyntactic abbreviations used here and in the examples below, such as PST (past), PASS (passive), INS (instrumental) or COMP (complementiser), follow the Leipzig Glossing Rules (<http://www.eva.mpg.de/lingua/resources/glossing-rules.php>) and are written in SMALL CAPITALS. Also lemmata (e.g., FRUIT above) are written in SMALL CAPITALS, wordforms as they occur in texts — in *italics*, and their translations are delimited by ‘single quotes’.

involves the understanding of interaction between hyperonymy and quantification, whether it involves awareness of diathetic equivalences, etc.

This is a relatively novel research task — not just in the context of Polish — and the need for such qualitative RTE evaluation resources has been raised in recent RTE literature, e.g., in Dagan et al. (2013, pp. 23, 161–162; in the section on *Future Directions for Entailment Evaluation* and in the chapter on *Research Directions in RTE*). Few steps have been taken in this direction so far and, to the best of our knowledge, they almost universally concern English. The need for creating specialised RTE corpora for different inference types was expressed in Bentivogli et al. (2010), where a method is proposed of manually distilling such corpora from general RTE corpora.⁵ Similarly, Sammons, Vydiswaran, and Roth (2010) proposed to annotate existing RTE corpora with types of inference steps needed to recognise entailment or lack thereof. Both papers present examples of inference labels, but do not attempt to provide a systematic taxonomy of inference types.

There is also some previous work which concentrates on particular inference types. An early example is Cooper et al. (1996), mentioned in the previous section. A more recent example is Toledo et al. (2012), which reports on the annotation of general RTE corpora (RTE1–4) with occurrences of restrictive, intersective and appositive modification playing a role in textual entailment. A particularly interesting work is that of MacCartney (2009) (see also MacCartney & Manning, 2009) which investigates various monotonicity effects in natural language inference.

The only attempt at providing a preliminary ontology of entailment phenomena that we are aware of is made at the following wiki web page related to Sammons et al. (2010), but containing inference labels revised in January 2011: <https://wiki.cites.illinois.edu/wiki/display/rte/data/Revised+Entailment+Phenomena+Ontology> (last accessed on 6th January 2015). There, five general types of phenomena are listed:

1. Knowledge Domains — contains inference types which occur in RTE corpora particularly often, e.g., lexical relations to do with employment or with killing and injuring,
2. Hypothesis Structures — labels describing structural aspects of the hypothesis (the second element in the entailment pair) relevant to entailment, e.g., the fact that location is provided for the event described there or that one of the semantic relations in the hypothesis is given only implicitly,
3. Inference Phenomena:
 - (a) Syntactic — e.g., categorially different expression of a relation in the text and in the hypothesis (for example, with a verb in the text and with a nominalisation in the hypothesis), or differences in diathesis between the text and the hypothesis (e.g., active vs. passive),
 - (b) Semantic — e.g., various kinds of coreference phenomena, corresponding terms standing in a hyperonymy (meronymy, etc.) relation, the fact that one of the arguments is implicit, etc.

⁵Similar work on Japanese is reported in in Kaneko, Miyao, and Bekki (2013).

4. Negative Entailment Phenomena — labels of this type indicate various phenomena only found in those pairs where entailment does not hold, e.g., when the same relation is expressed in the text and in the hypothesis but with incompatible values of the same argument (e.g., in (1), the purchase relation is present in both T and H2, but the price tags are incompatible: \$2Bn vs. \$3.4Bn),
5. Knowledge Resources — types of inferences involving extra-linguistic knowledge, e.g., spatial knowledge required to infer *George was in France* from *George visited Paris*.

It should be clear that the above classification is very heterogeneous: it is based on widely different criteria and a single phenomenon may, e.g., fit the Knowledge Domains class of phenomena (because it occurs often in RTE corpora) and be an Inference Phenomenon or a Negative Entailment Phenomenon, at the same time. Moreover, this ontology does not cover linguistic inference phenomena in any systematic way. Finally, particular inference labels are described very briefly or sometimes not at all; for example, the label `create` is described in the ontology as “includes create, invent, write, produce, build, born”, and similarly in the equally brief annotation instructions at <https://wiki.cites.illinois.edu/wiki/display/rtedata/Annotation+Instructions> (accessed on 6th January 2015).

The aim of the proposed research is to create a comprehensive and logically coherent taxonomy of linguistic inference phenomena applicable not only to English, but also to Polish and other languages, i.e., taking into consideration a much richer set of phenomena. While this taxonomy should not initially include types of encyclopedic knowledge (e.g., that Paris is a capital of France or that somebody who is alive in 1800 cannot be alive in 2015), it should encompass the more logical types of inference (of the kind discussed in Cooper et al., 1996; MacCartney & Manning, 2009 and Toledo et al., 2012, 2013), related to the meaning of words expressing quantifiers, logical connectives or types of modification. Most importantly, such a taxonomy should be developed by building on linguistic knowledge concerning different ways of expressing semantic relations in natural languages. Hence, unlike the attempts reported in Bentivogli et al. (2010) and Sammons et al. (2010), the taxonomy should ideally reflect all inference types made available by a the system of a natural language (e.g., Polish or English), not just those which happen to occur in a given RTE corpus (especially that such corpora are currently empirically limited, typically to a dozen hundred entailment pairs). As argued in the following section, there is a thread of work in Polish linguistics that is of particular importance in this respect.

3. Methodology

The issue of possible syntactic realisations of various semantic predicates has been extensively studied within the so-called “semantic syntax” approach of Stanisław Karolak (1972, 1984, 2001, 2002), sometimes referred to as the Polish School of Semantic Syntax (Szumska, 2013, p. 13), also by other researchers working in this paradigm (e.g., Grochowski, 1984; Korytkowska & Małdziewa, 2002; Kiklewicz & Korytkowska, 2010, 2012; Szumska, 2013). The main task of this line of research

does not seem to be to provide a taxonomy of inference or equivalence relations holding between natural language constructions, but rather to exhaustively describe syntactic realisations of various types of semantic predicates.

For example, for Kiklewicz and Korytkowska (2010): a three-argument predicate of type $\langle e, \langle e, \langle t, t \rangle \rangle \rangle$, i.e., taking two entities and a truth value, and returning a truth value (such predicates are marked as, e.g., $P(x, y, r)$ in work on semantic syntax), Kiklewicz and Korytkowska (2010) lists 15 general types of syntactic realisations, including: “ VN_x, N_y, V_r ”, “ VN_x, N_y, VI_r ” and “ $VN_{x,y}, N_{ar}, \emptyset_r$ ”. In the first two, the two entity arguments are realised as nominal phrases (N_x, N_y) and the propositional argument is realised as a finite clause (V_r) or an infinitival phrase (VI_r). In the third type, the two entity arguments are realised jointly ($N_{x,y}$) in a reciprocal construction, as in (5) below from Kiklewicz and Korytkowska (2010): §3.2.1.14, and the propositional argument is not realised as such (cf. \emptyset_r), but rather “condensed” to a single nominal entity (N_{ar}) within this proposition (here expressed by *o Andrzej*u):

- (5) Przyjaciele rozmawiali o Andrzeju.
 friends.NOM talk.PL.PST about Andrzej.LOC
 ‘Friends talked about Andrzej.’

Moreover, for each of such types of linguistic realisations, possible surface forms of these types are listed together with lemmata which give rise to such surface constructions. For example, for the type “ VN_x, N_y, V_r ”, the first nominal argument, N_x , is assumed to always occur in the nominative case, but four different surface realisations of N_y are given: N_{dat} (in the dative), N_{acc} (in the accusative), *Praep N_{instr}* (a prepositional phrase with the instrumental NP) and *Praep N_{gen}* (a prepositional phrase with the genitive NP). In all four cases, the surface realisation of the propositional argument V_r is specified as “(*Pron*) *Con* V ”, i.e., a finite clause (V) introduced by a complementiser (*Con*) and an optional pronoun (*Pron*). Two sentences illustrating the type “ VN_x, N_y, V_r ”, with N_y realised as a dative NP or as a PP (prepositional phrase) with a genitive NP, are given below:

- (6) Anna dziękuje Piotrowi (za to), że jej pomógł.
 Anna.NOM thanks.SG Piotr.DAT for this.ACC COMP her.DAT helped.SG
 ‘Anna is thanking Piotr (for the fact) that he helped her.’
- (7) Chcę od ciebie (tego), abyś wyszedł.
 want.1.SG from you.GEN this.GEN COMP.2.SG left
 ‘I want you to leave.’

It should be clear that the above surface syntactic specifications are not fully explicit: the form of the complementiser (*Con*) is not specified (two different complementisers are needed in the two examples above), and neither is the form of the preposition (*Praep*) or the optional pronoun (*Pron*; in fact it is introduced by a preposition in (6) — a possibility not mentioned in the schema at all). While such information is present in some earlier work, notably in Korytkowska and Małdziewa (2002), other syntactic distinctions commonly assumed in contemporary linguistics

are not handled in the semantic syntax approach, including the semantically potent distinction between subject control and object control (cf., e.g., Rosenbaum, 1967 and Landau, 2013, and — in the context of Polish — e.g. Przepiórkowski, 2004 and Witkoś, 2007).⁶ Also the quasi-formal notation used in this paradigm leaves much to be wished for. This includes the use of various — often misleading — conventions instead of mechanisms standard in contemporary formal semantics such as lambda calculus and explicit semantic types (here: *e* and *t*). One such convention is the use of the same symbols with different meanings (e.g., *V* indicating the described predicate in some places and a finite clause in other), another — the use of specific variable names for signalling semantic types.⁷

Nevertheless, despite these deficiencies, this thread of work in Polish linguistics remains a rich source of information on different lexical and syntactic ways of expressing the same semantic relations. For example, Karolak (1984, p. 94) discusses the (stems of the) lemmata *NALEŻEĆ* ‘belong’, *MIEĆ* ‘have’, *WŁASNOŚĆ* ‘property’ and *WŁAŚCICIEL* ‘owner’. While all of them express the 2-argument ownership relation, the first two realise the two arguments differently (the subject of *NALEŻEĆ* corresponds to the non-subject argument of *MIEĆ*, and conversely for the other argument of *NALEŻEĆ*), and similarly for *WŁASNOŚĆ* and *WŁAŚCICIEL*. Awareness of these facts makes it possible to recognise that the following four sentences are semantically equivalent (Karolak, 1984, p. 94):

- (8) Zastawa należy do Piotra.
china set.NOM belongs to Piotr.GEN
‘The china set belongs to Piotr.’
- (9) Piotr ma zastawę.
Piotr.NOM has china set.ACC
‘Piotr has a china set.’
- (10) Zastawa jest własnością Piotra.
china set.NOM is property.INS Piotr.GEN
‘The china set is the property of Piotr.’
- (11) Piotr jest właścicielem zastawy.
Piotr.NOM is owner.INS china set.GEN
‘Piotr is the owner of the china set.’

⁶Such detailed morphosyntactic information is explicitly given in the largest Polish valence dictionary, *Walenty*, developed at IPI PAN (Przepiórkowski et al., 2014a, 2014b; Hajnicz, Nitoń, Patejuk, Przepiórkowski, & Woliński, in press). See <http://zil.ipipan.waw.pl/Walenty> for description, publications and textual snapshots of the dictionary, and <http://walenty.ipipan.waw.pl/> for a web interface to the current state of *Walenty*.

⁷This latter convention is incorrectly assumed to be a necessary property of the underlying logic, cf. Kiklewicz and Korytkowska (2012, p. 62). On the other hand, the notation used in such recent semantic syntax work is certainly more clear than the original notation of Karolak (1984), where, e.g., in $M\{T, L\{\phi[x, y, z, \phi[x \dots n, f(x \dots n)]]\}$ (on page 73), multiple occurrences of the same unbound variables *x* and *y* should actually be understood as different and unrelated variables, the two occurrences of ϕ refer to different predicates, the notation $x \dots n$ is never explained (but the two occurrences of *n* seem to indicate the — possibly different — numbers of arguments of corresponding predicates), and the semantics of different types of brackets is unclear.

Another important phenomenon extensively discussed within this thread of work is the suppression of some semantic arguments, as in case of MEŻATKA ‘married woman’, where only one of the two arguments of the relation also expressed by OŻENIĆ SIĘ ‘marry’ may be realised (Karolak, 1984, p. 63):

- (12) Roman ożenił się z Marią.
 Roman.NOM marry.PST REFL with Maria.INS
 ‘Roman married Maria.’
- (13) Maria została mężatką.
 Maria.NOM become.PST married woman.INS
 ‘Maria got married.’

Also the possibility to “condense” a propositional argument to an entity within it gets a fair treatment in the semantic syntax approach, as already illustrated in (5), where *o Andrzeju* ‘about Andrzej’ may represent a proposition like ‘about what Andrzej did’ or ‘about what Andrzej is like’.

Obviously, as in case of any research based on previous work, it is necessary to maintain a critical approach to prior claims, and no exception should be made here: the characterisation of some apparent equivalences discussed in semantic syntax turns out to be imperfect on closer scrutiny. For example, the claim that the following two sentences, each involving the negated trust relation, are equivalent (Karolak, 1984, p. 50) does not seem to be correct: if Piotr is *nieufny* ‘distrustful’, that does not necessarily imply that he does not trust anybody, but may mean that he takes longer time to start trusting people he does not know:

- (14) Piotr jest nieufny.
 Piotr.NOM is distrustful.NOM
 ‘Piotr is distrustful.’
- (15) Piotr nie ufa nikomu.
 Piotr.NOM NEG trusts nobody.INS
 ‘Piotr does not trust anybody.’

Similar doubts may be raised about another pair of derivationally related lexemes discussed there: BOJAŻLIWY ‘fearful, timid’ and BĄĆ SIĘ ‘fear, be afraid’: when somebody is *bojaźliwy*, that does not necessarily mean that he or she fears everything, but may simply mean that he or she fears more things than usual or fears the usual things more than other people do.

While much semantic syntax work is concerned with different realisations (or suppression) of arguments, Grochowski (1984) discusses ways of combining semantic predicates in modification constructions, as in the first of the following two sentences (from Karolak, 1972, p. 152), equivalent to the second sentence, where the purpose relation is expressed more explicitly (Grochowski, 1984, p. 266):

- (16) Złożył wizytę, żeby się oświadczyć.
 paid visit COMP REFL propose.INF
 ‘He paid a visit in order to propose (to her).’

- (17) Celem jego wizyty były oświadczyzny.
 purpose.INS his visit.GEN were proposal.PL.NOM
 ‘The purpose of his visit was to propose (to her).’

This work is also a good source of information about equivalent ways of expressing various logical relations, e.g., the relation expressed by *PONIEWAŻ* ‘because’ (Grochowski, 1984, pp. 288–290), e.g., using forms of subordinate conjunctions *ALBOWIEM*, *BO*, *GDYŻ*, etc., or complex prepositions *Z POWODU*, *W WYNIKU*, *Z RACJI*, etc.

Note that observations made within the semantic syntax school concern not only the symmetrical relation of equivalence, but also the asymmetrical relation of entailment, especially, in cases of “condensing” as in (5) above or (18)–(19) below (Grochowski, 1984, p. 267):

- (18) Jan idzie do delikatesów po kawę.
 Jan.NOM goes to delicatessen.GEN for coffee.ACC
 ‘Jan goes to the delicatessen for coffee.’
- (19) a. Jan idzie do delikatesów, aby kupić kawę.
 Jan.NOM goes to delicatessen.GEN COMP buy.INF coffee.ACC
 ‘Jan goes to the delicatessen to buy coffee.’
- b. Jan idzie do delikatesów, aby ukraść kawę.
 Jan.NOM goes to delicatessen.GEN COMP steal.INF coffee.ACC
 ‘Jan goes to the delicatessen to steal coffee.’

As, without the full context, it is not clear what proposition is “condensed” to *po kawę* ‘for coffee’ in (18), it is entailed by (19a), (19b) and similar such sentences, but — strictly speaking — entails neither of them (although native speakers will probably often infer (19a) from (18)).

In summary, we claim that the semantic syntax tradition associated with the name of Stanisław Karolak may be a reasonable starting point when devising a linguistically-oriented taxonomy of entailment (and, in particular, equivalence) phenomena. Initial steps towards creating such a taxonomy are made in the following section.

4. Towards a Taxonomy

While the development of a taxonomy of phenomena and kinds of knowledge determining the process of entailment is a research programme requiring much deeper studies of both the linguistic (esp., semantic syntax) literature and of the available entailment corpora, we will boldly attempt to sketch here some desiderata for such a taxonomy.

First of all, as already indicated above, the creation of the taxonomy will first concentrate on linguistic phenomena rather than on world knowledge. As is well known, the issue of distinguishing knowledge about language from knowledge about the world is vexed, and many linguists have for a long time remained sceptical about the possibility of making such a strict distinction, as illustrated by the following quote from Bloomfield (1933, p. 139; cited after Hobbs, 2011, p. 756): *In order to*

give a scientifically accurate definition of meaning of every form of a language, we should have to have a scientifically accurate knowledge of everything in the speakers' world. More recent discussions of this issue may be found in Hobbs (2011, pp. 755–760) and Ovchinnikova (2012, pp. 31–33), with both authors concluding that it is not clear that a border between these two kinds of knowledge may be drawn. We will not assume such a clear boundary either, but — as a methodological decision — will start with phenomena which are least controversially purely linguistic (e.g., concerning diathesis), gradually moving towards phenomena bordering on world knowledge, e.g.: hyperonymy, meronymy and other relations defined in wordnets (Miller, Beckwith, Fellbaum, Gross, & Miller, 1990; Piasecki, Szpakowicz, & Broda, 2009); the kinds of information found in generative lexicons (Pustejovsky, 1995), e.g., making it possible to (defeasibly) infer *John finished smoking a cigarette* from the shorter *John finished a cigarette*; etc.

Second, the taxonomy will initially be constructed on the basis of Polish and English, but in a way that will make it possible to extend it to other languages. Hence, the top (most general) categories will be maximally language-independent, and only the lower (more specific) categories will perhaps indicate more language-dependent phenomena. For example, one category may be concerned with diathesis phenomena, common across natural languages, with subcategories such as: impersonal, passive, dative alternation, locative alternation, causative alternation, etc., which vary across languages considerably. Similarly, there may be a category (perhaps a subcategory of a more general class of various entailment phenomena related to derivative morphology) encompassing diminutives and augmentatives, with a subcategory of depreciative forms such as Polish *profesory* 'professors' (Saloni, 1988), which have no direct equivalent in English and many other languages.

Third, each maximally specific category will be illustrated with examples demonstrating the impact of phenomena in this category on entailment. For example, in case of nominalisations expressing propositions, the following examples (Korytkowska & Małdzińska, 2002, p. 26) may be used to construct four entailment triples: $\langle(21a), (20), \text{yes}\rangle$, $\langle(21b), (20), \text{yes}\rangle$, $\langle(20), (21a), \text{no}\rangle$ and $\langle(20), (21b), \text{no}\rangle$ — i.e., either of (21a–b) entails (20), but not the other way round.

(20) To spowodowało u Ani ból głowy.
 this.NOM caused at Ania.GEN pain.ACC head.GEN
 'This caused Ania's headache.'

(21) a. To spowodowało, że Anię boli głowa.
 this.NOM caused COMP Ania.ACC ache.PRS head.NOM
 'The effect of this is that Ania has a headache.'

b. To spowodowało, że Anię bolała głowa.
 this.NOM caused COMP Ania.ACC ache.PST head.NOM
 'The effect of this was that Ania had a headache.'

The artificially constructed entailment examples in the test-suite will be minimal in the sense that each such an example (e.g., each of the four triples given above) should illustrate a very small number of entailment phenomena, often just one.

This should be contrasted with entailment examples in typical RTE corpora, as in (1) above, where usually multiple entailment steps of very different nature must be made. However, annotating such more realistic entailment pairs with labels from the taxonomy is also planned, as the two kinds of resources — a test-suite with minimal pairs and a realistic RTE corpus — will serve to evaluate different aspects of RTE systems.

Moreover, the manually constructed test-suite, just like typical RTE corpora, should also contain examples of non-entailment, as already implied above. The assumption is that, in such examples, it is possible to identify one or two phenomena which break a chain of entailment steps (cf. Negative Entailment Phenomena in §). For example, both of the triples $\langle(20), (21a), \text{no}\rangle$ and $\langle(20), (21b), \text{no}\rangle$, should be marked with an appropriate nominalisation label as an entailment step that would have to be made to infer (21a–b) from (20),⁸ but they should also be annotated with a label indicating that the hypothesis contains additional temporal information about the time of the headache that is not present in the premise.

As a starting point for the development of the taxonomy, let us discuss just a few possible categories of entailment steps:

(22) logical constructions:

- a. connectives
- b. quantifiers
- c. negation
- d. collectivity and distributivity

(23) ...

(24) different expressions of the same lexical semantic predicates:

- a. lemma-preserving diathesis (includes obligatory argument suppression, as in some impersonal constructions):
 - i. impersonal constructions
 - ii. passivisation
 - iii. dative inversion
 - iv. ...
- b. other grammatical-class-preserving diathesis
- c. diathesis across grammatical classes:
 - i. verbal–nominal
 - ii. verbal–adjectival
 - iii. ...
- d. condensation of a propositional dependent
- e. ...

⁸Note that nominalisation is not understood directionally here (from verbal to the nominal), but rather as a step that can be made in either direction.

(25) hyperonymy relations

(26) ...

(27) world knowledge

For brevity, we will illustrate these preliminary entailment categories with mostly English examples.

The category (22) contains entailment steps analogous to those involving logical connectives and quantifiers in formal logic, similar to those discussed in Cooper et al. (1996) and MacCartney and Manning (2009). For example, the entailment pair *⟨John is eating and drinking., John is eating.⟩* — and perhaps also the non-entailment pair *⟨John is eating or drinking., John is eating.⟩*, with *or* instead of *and* — would be labelled with (22a), as the entailment involves understanding how natural languages express logical connectives. As mentioned above, in the context of semantic syntax, Grochowski (1984) is a good source of information on different ways of expressing such connectives in Polish. Similarly, the entailment pair *⟨Many people came., Somebody came.⟩* would be marked with (22b), as it involves the understanding of words (here, *many* and *somebody*) expressing logical quantifiers. Moreover, the pair *⟨Every person came., Every woman came.⟩* — and maybe also the non-entailment pair *⟨Every woman came., Every person came.⟩* — would be marked with both (22b) and with (25), as the entailment combines the understanding of the monotonicity properties of the quantifier expressed by *every* and the fact that PERSON is a hyperonym of WOMAN. Similarly, the pair *⟨John didn't buy vegetables., John didn't buy Brussels sprouts⟩* should be marked with (22c) and (25). Another subcategory, related to quantification, is concerned with the issues of collectivity and distributivity, e.g., the fact that *John ate an apple.* does not follow from *John and Mary ate an apple.* (they could have eaten half an apple each), but it does follow from *John and Mary ate an apple each.* While in English the distributive element EACH has the same form as the quantifier EACH (as in *Each boy ate an apple.*), a distinct preposition-like distributive element is observed in Polish and other Slavic languages, PO (cf. Przepiórkowski, 2014 and references therein).

Another class, not listed above, should be concerned with grammatical phenomena discussed in Toledo et al. (2012, 2013), namely, restrictive and intersective modification, apposition, copular constructions, etc. Other categories, also not explicated here, should represent phenomena extensively discussed within semantic syntax and within the related work on Bulgarian-Polish contrastive grammar summarised in (Koseska-Toszewa, Korytkowska, & Roszko, 2007), namely, definiteness, modality, tense and aspect, and perhaps also spatial (locative) relations. Another class related to intensive work within the semantic syntax paradigm, already mentioned in §, is dedicated to various ways of expressing a given semantic predicate and its arguments on the surface. For example, the equivalence between *John gave a book to Mary.* and *John gave Mary a book.*, involving the same lemma GIVE, would be marked with (24a.iii), the equivalence between (8) and (9), involving two different verbal lemmata NALEŻEĆ ‘belong’ and MIEĆ ‘have’, would be marked with (24b) (and similarly for (10) and (11), involving two nominal lemmata WŁASNOŚĆ

‘property’ and WŁAŚCICIEL ‘owner’), while the equivalence between, e.g., (8) and (10), as well as (12)–(13), where the same semantic predicate is expressed by lexical items belonging to different grammatical classes, verbal and nominal, would be labelled with (24c). Similarly, the pairs ⟨(19a), (18)⟩ and ⟨(19b), (18)⟩, involving “condensation” of a propositional argument, would be labelled with (24d). Obviously, this category of entailment (and equivalence) relations would contain many more subcategories than shown in (24) — this is signalled by multiple occurrences of ellipses ‘...’.

5. Conclusion

In this position paper we tried to tie two threads of linguistic and computational linguistic research which — to the best of our knowledge — have never met before, namely, work on textual entailment developed so far mainly in the context of English and a few other non-Slavic languages,⁹ and work on semantic syntax, carried out in the context of Polish and other Slavic languages. We argued that the latter thread may constitute a good starting point for the development of a linguistically-oriented taxonomy of entailment types, as well as a test-suite of entailment pairs labelled with elements of this taxonomy. While the paper is admittedly programmatic, the research direction it proposes seems sufficiently novel and risky to put it forward for discussion — and critique from both: computational linguists and semantic syntax researchers — at this very early stage.

References

- Barwise, J. & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4 (2), 159–219. <http://doi.org/10.1007/BF00350139>
- Bentivogli, L., Cabrio, E., Dagan, I., Giampiccolo, D., Leggio, M. L., & Magnini, B. (2010). Building textual entailment specialized data sets: A methodology for isolating linguistic phenomena relevant to inference. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation, LREC 2010* (pp. 3542–3549). Valletta, Malta: ELRA.
- Bloomfield, L. (1933). *Language*. New York: Holt.
- Cooper, R., Crouch, D., van Eijck, J., Fox, C., van Genabith, J., Jaspars, J., Kamp, H., Pinkal, M., Milward, D., Poesio, M., & Pulman, S. (1996). *FraCaS: Using the framework*. Deliverable D16, University of Edinburgh.
- Dagan, I., Roth, D., Sammons, M., & Zanzotto, F. M. (2013). *Recognizing textual entailment: Models and applications*. San Rafael, Calif.: Morgan & Claypool.
- Giampiccolo, D., Magnini, B., Dagan, I., & Dolan, B. (2007). The third PASCAL Recognizing Textual Entailment Challenge. In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague: Association for Computational Linguistics. <http://doi.org/10.3115/1654536.1654538>
- Grochowski, M. (1984). Składnia wyrażen polipredykatywnych (zarys problematyki). In Z. Topolińska (Ed.), *Gramatyka współczesnego języka polskiego: Składnia* (pp. 213–299). Warszawa: Wydawnictwo Naukowe PWN.
- Hajnicz, E., Nitoń, B., Patejuk, A., Przepiórkowski, A., & Woliński, M. (in press). Internetowy słownik walencyjny języka polskiego oparty na danych korpusowych. *Prace Filologiczne*, 65.

⁹But some initial work on Czech has recently been reported (Nevřilová, 2014).

- Hobbs, J. R. (2011). Word meaning and world knowledge. In C. Maienborn, K. von Heusinger, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (Vol. 1, pp. 740–761). Berlin: De Gruyter Mouton.
- Kaneko, K., Miyao, Y., & Bekki, D. (2013). Building Japanese textual entailment specialized data sets for inference of basic sentence relations. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics* (Vol. 2: *Short Papers*, pp. 273–277). Sofia, Bulgaria.
- Karolak, S. (1972). *Zagadnienia składni ogólnej*. Warszawa: Wydawnictwo Naukowe PWN.
- Karolak, S. (1984). Składnia wyrażen predykatywnych. In Z. Topolińska (Ed.), *Gramatyka współczesnego języka polskiego: Składnia* (pp. 11–211). Warszawa: Wydawnictwo Naukowe PWN.
- Karolak, S. (2001). *Od semantyki do gramatyki: Wybór rozpraw*. Warszawa: Slawistyczny Ośrodek Wydawniczy PAN.
- Karolak, S. (2002). *Podstawowe struktury składniowe języka polskiego*. Warszawa: Slawistyczny Ośrodek Wydawniczy PAN.
- Kiklewicz, A. & Korytkowska, M. (Eds.). (2010). *Podstawowe struktury zdaniowe współczesnych języków słowiańskich: Białoruski, bułgarski, polski*. Olsztyn: Centrum Badań Europy Wschodniej Uniwersytetu Warmińsko-Mazurskiego w Olsztynie.
- Kiklewicz, A. & Korytkowska, M. (2012). Modelowanie płaszczyzny syntaktycznej a segmentacja hasła słownikowego (na przykładzie języków słowiańskich). *Biuletyn Polskiego Towarzystwa Językoznawczego*, 68, 49–68.
- Korytkowska, M. & Małdziewa, W. (2002). *Od zdania złożonego do zdania pojedynczego: Nominalizacja argumentu propozycjonalnego w języku polskim i bułgarskim*. Toruń: Wydawnictwo Uniwersytetu Mikołaja Kopernika.
- Koseska-Toszewa, V., Korytkowska, M., & Roszko, R. (2007). *Polsko-bułgarska gramatyka konfrontatywna*. Warszawa: Wydawnictwo Akademickie Dialog.
- Landau, I. (2013). *Control in generative grammar: A research companion*. Cambridge: Cambridge University Press.
- MacCartney, B. (2009). *Natural language inference* (Unpublished Ph.D. dissertation). Stanford University.
- MacCartney, B. & Manning, C. D. (2009). An extended model of natural logic. In *The Eighth International Conference on Computational Semantics (IWCS-8)* (pp. 140–156). Association for Computational Linguistics. <http://doi.org/10.3115/1693756.1693772>
- Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., & Miller, K. J. (1990). Introduction to WordNet: An online lexical database. *International Journal of Lexicography*, 3(4), 235–244. <http://doi.org/10.1093/ijl/3.4.235>
- Mostowski, A. (1957). On a generalization of quantifiers. *Fundamenta Mathematicae*, 44, 12–36.
- Nevěřilová, Z. (2014). *Paraphrase and textual entailment generation in Czech* (Unpublished Ph.D. thesis). Masaryk University, Brno.
- Ovchinnikova, E. (2012). *Integration of world knowledge for natural language understanding*. Paris: Atlantis Press.
- Piasecki, M., Szpakowicz, S., & Broda, B. (2009). *A wordnet from the ground up*. Wrocław: Oficyna Wydawnicza Politechniki Wrocławskiej.
- Przepiórkowski, A. (2004). On case transmission in Polish control and raising constructions. *Poznań Studies in Contemporary Linguistics*, 39, 103–123.
- Przepiórkowski, A. (2014). Syntactic and semantic constraints in a Glue Semantics approach to distance distributivity. In M. Butt & T. H. King (Eds.), *The Proceedings*

- of the LFG'14 Conference (pp. 482–502). Stanford, CA: CSLI Publications.
- Przepiórkowski, A., Skwarski, F., Hajnicz, E., Patejuk, A., Świdziński, M., & Woliński, M. (2014a). Modelowanie własności składniowych czasowników w nowym słowniku walencyjnym języka polskiego. *Polonica*, 33, 159–178.
- Przepiórkowski, A., Hajnicz, E., Patejuk, A., Woliński, M., Skwarski, F., & Świdziński, M. (2014b). Walenty: Towards a comprehensive valence dictionary of Polish. In N. Calzolari, K. Choukri, T. Declerck, H. Loftsson, B. Maegaard, J. Mariani, A. Moreno, J. Odijk, & S. Piperidis (Eds.), *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014* (pp. 2785–2792). Reykjavík, Iceland: ELRA.
- Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: The MIT Press.
- Rosenbaum, P. (1967). *The grammar of English predicate complement constructions*. Cambridge, MA: The MIT Press.
- Saloni, Z. (1988). O tzw. formach nieosobowych [rzeczowników] męskoosobowych we współczesnej polszczyźnie. *Biuletyn Polskiego Towarzystwa Językoznawczego*, 41, 155–166.
- Sammons, M., Vydiswaran, V., & Roth, D. (2010). “Ask not what textual entailment can do for you...”. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics* (pp. 1199–1208). Uppsala, Sweden.
- Szumska, D. (2013). *The adjective as an adjunctive predicative expression*. Frankfurt am Main: Peter Lang.
- The Leipzig Glossing Rules: Conventions for interlinear morpheme-by-morpheme glosses. (2015).
- Toledo, A., Katrenko, S., Alexandropoulou, S., Klockmann, H., Stern, A., Dagan, I., & Winter, Y. (2012). Semantic annotation for textual entailment recognition. In I. Batyrshin & M. G. Mendoza (Eds.), *Advances in Computational Intelligence: 11th Mexican International Conference on Artificial Intelligence, MICAI 2012 San Luis Potosí, Mexico, October 27 – November 4, 2012* (pp. 12–25). Berlin: Springer-Verlag. (*Lecture Notes in Artificial Intelligence*, 7630).
- Toledo, A., Alexandropoulou, S., Katrenko, S., Klockmann, H., Kokke, P., & Winter, Y. (2013). Semantic annotation of textual entailment. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) — Long Papers* (pp. 240–251). Potsdam: Association for Computational Linguistics.
- Topolińska, Z. (Ed.). (1984). *Gramatyka współczesnego języka polskiego: Składnia*. Warszawa: Wydawnictwo Naukowe PWN.

Acknowledgment

Work reported here has been partially financed by the Polish Ministry of Science and Higher Education within the CLARIN ERIC programme 2015–2016 (<http://clarin.eu/>).

The author declares that he has no competing interests.

This is an Open Access article distributed under the terms of the Creative Commons Attribution 3.0 PL License (<http://creativecommons.org/licenses/by/3.0/pl/>), which permits redistribution, commercial and non-commercial, provided that the article is properly cited.

© The Author 2015

Publisher: Institute of Slavic Studies, PAS, University of Silesia & The Slavic Foundation