Contexts for Quantification

Valeria de Paiva
Nuance Communications
1198 Arques Avenue
Sunnyvale, California 94085

Abstract

Logical systems conceived for providing semantics/logical forms for sentences of English abound. From Montague's original Higher-order Intensional Logic in the seventies, to Situation Theory and Discourse Representation Theory (in the eighties) as well as frame languages and Description Logics (in the nineties), to "vanilla" versions of First-order logic (FOL) more recently, the field is rife with possibilities and issues.

I want to describe one such language, the product of several years of development of the NLP-based knowledge representation system Bridge, at Xerox PARC. The design of the language and its logic was historically tied-up to the development of the software system, but the language and its inferential system are of independent interest. I call this logic language TIL for Textual Inference Logic. TIL is one of the systems associated with Natural Logic and it distinguishes itself by the unorthodox treatment of quantification in terms of instantiation of concepts within (McCarthy-style) contexts. I describe this mechanism of quantification, necessitated by a better modeling of negation and other intensional phenomena, ubiquitous in natural language. My goal is to map and relate TIL's expressive power to that of more traditional systems, as for example, the ones described by Moss and Pratt-Hartmann.

Introduction

This work fits within the program of Natural Logic as considered by Moss, who says in [Moss, 2012] "*natural logic* [is] the study of logical systems for linguistic reasoning". We want to show that significant parts of natural language inference can be carried out in easy, mathematically characterized logical systems. The emphasis here is on the methods. Whenever possible, we want to obtain *complete* axiomatizations, because completeness of a logical system provides us with more information about the system modeled, but completeness is not an over-riding goal, while soundness is. This work uses tools from fields like proof theory, categorical logic, (finite) model theory and, eventually, complexity theory.

The work presented here describes the logical system TIL (for textual inference logic), a formalization of the AI software system Bridge developed in PARC to produce logical forms or logical representations of the meaning of sentences English. Previous presentations of TIL described it as a stand alone system, synthesized from what is produced by the software, when processing sentences in English. This is a version of applied logic, in this case logic applied to Natural Language Understanding (NLU).

Much work has gone and still has to go into the *process* of creating sensible logical representations from sentences, but I will not describe this process (or how to improve it) in this note. I want to place TIL in the context of the work of Pratt-Hartmann[Pratt-Hartmann, 2004] and Moss[Moss, 2010], as an extension of traditional syllogistic logic. This means that for the time being we assume that the processing pipeline that produces the logical representations works perfectly: parsing works flawlessly, lexicons are as complete as necessary, disambiguation is not a problem (some oracle knows exactly what the author of the sentence meant and how to write it in this idealized language). In this idealized situation what can we say about the logical system we are operating with?

Which logic?

The idealized system Bridge reads in a sentence s in English and produces a logical knowledge representation r for it. The collection of all representations plus some primitive means of manipulating them forms the system TIL. The logical language TIL was meant to be kept as close as it was possible (and sensible) to FOL (first-order logic) but from the start we knew that to model natural language sentences we wanted intensional 'concepts' and 'contexts'.(cf. [Condoravdi et al., 2001])

Let us start by discussing a simple example: **R1** *"Three boys ate 5 pizzas"* The TIL representation for this sentence is a collection of clauses, some part of the *conceptual structure*, some part of the *contextual structure* and still some part of the *temporal structure* that will not discuss in this note.

R1 Conceptual Structure:
subconcept(boy-1,[List1])
subconcept(eat-3,[List2])

subconcept(pizza-5,[List3])
role(ob,eat-3,pizza-5)
role(sb,eat-3,boy-1)
role(cardinality-restriction,boy-1,3)
role(cardinality-restriction,pizza-5,pl)
Contextual Structure:
context(t)
instantiable(boy-1,t)
instantiable(eat-3,t)
instantiable(pizza-5,t)
top-context(t)

Previous descriptions of TIL [Bobrow et al., 2005; De Paiva et al., 2007] discussed how these clause representations can be achieved using a layered approach and the XLE system[Maxwell and Kaplan, 1996] and why this layers are useful. Previous work also discussed the 'packing' of all the (series of) structures obtained in the several layers, where by 'packing' we mean that instead of disambiguating structures and pruning the less likely ones at each stage of the pipeline, XLE algorithms allow us to keep a condensed representation of all possibilities, effectively avoiding premature pruning of the correct ones. But in this note we are not interested in the process of obtaining representations, nor in the process of disambiguating them. We take as given a collection of representations similar to the one above and we want to relate them to other kinds of representations for natural language semantics.

Concepts, contexts and roles

It is traditional for logics of Knowledge Representation to be fragments of first-order logic (FOL). By contrast, it is traditional for logics for natural language semantics to be higher-order intensional logics. Our logic has concepts, which make it look like a "description logic", that is, a fragment of FOL, but it also has contexts, a possible-worlds-like construct that, we hope, is expressive enough for the needs of natural language.

Concepts, the way we conceive them, come from both neo-Davidsonian event semantics and, somewhat independently, from description logics. Some of our reasons for using a concept denoting analysis instead of an individual denoting analysys when mapping noun phrases to logic are discussed in [Condoravdi et al., 2001]. The main reasons are being able to deal with non-existent entities (for example when mapping *"Negotiations prevented a strike"* we do not want to say that there exists negotiations N and there exists a strike S and $prevented(N, S)$, as the prevented strike does not really exist in the actual world) and accounting for downward monotonicity entailments. Instead of constants, variables, predicates and relations symbols like in first-order logic, TIL has concepts and subconcepts, related by `roles` and organized in *contexts*. The basic concepts come from an ontology, which is a parameter of the system. In previous work we used both the Cyc ontology[Lenat et al., 1990] as well as a homegrown version of a judicious merger of WordNet and VerbNet, called the XLE Unified Lexicon [Crouch and King, 2005]. We believe we can also use SUMO[Niles and Pease, 2001] concepts for our representations.

Our initial choice was to use the concepts provided by the biggest common sense knowledge base available CYC, and to take advantage of its reasoning component. Although we used CYC concepts for our logic, we found it useful to map the text to an *abstract* form of knowledge representation (AKR), that could be realized as CYC or SUMO or KM or any other knowledge representation formalism. The logic is called TIL, but the representations are called AKR (for Abstrackt Knowledge Representation) structures.

The design of this abstract knowledge representation language aimed for a sweet spot between ease of mapping from text to a formalism, and tractability of the target formalism. Ease of mapping from that formalism to standard logical representations was also a consideration. A happy development was the realization that the representations we had developed in terms of events, roles and contexts, were already good enough for some important classes of textual inferences that we wanted to concentrate on. In general, the inferences we wanted to concentrate on were immediate, almost simple-minded, but necessary for the understanding of the text. For example, if the text says that "John managed to close the door" then we can safely infer that "John closed the door". Often these inferences were not even recognized as inferences by people. Furthermore, these inferences did not seem to depend crucially on the particular ontology; they were much more dependent on the articulation of inference patterns surrounding the use of particular classes of words which appear quite often in open texts.

The representations are based on the notion of *events* in a neo-Davidsonian style (a good reference is [Lasersohn, 2006]), coupled with the notion of *contexts* based on McCarthy's ideas [McCarthy, 1993]. In our example we need basic concepts for boy, eating and pizza. The concept for `boy-1` is a subconcept of one of the basic concepts provided by our ontology. Similarly for the concepts for `pizza-5` and `eat-3`.

Concepts in TIL are similar to Description Logic concepts in that they correspond pre-theoretically to sets of objects that satisfy a certain property, like predicates in FOL. They are similar to predicates in first-order logic, but are not always unary predicates. Think of concept named `eat-3` above as a collection of 'eating' events, in which other concepts in our domain participate. We have two kinds of concepts, primitive concepts extracted from an idealized version of the chosen ontology and constructed-on-the-fly concepts, which are always sub-concepts of some primitive concept. These second kinds of concepts are dynamic, created by the implemented system Bridge when we feed it English sentences. These dynamic concepts are created and placed in the hierarchy/ontology in use, as best as we can, at run time. We assume that our concepts are as fine or as coarse as the sentences that we deal with require. We also assume that our ontology is not circular or inconsistent. (In practice it is hard to show that this is indeed the case, especially with more expressive ontologies.)

Concepts are related to other concepts via *roles*. In our previous example:
role(ob,eat-3,pizza-5)
role(sb,eat-3,boy-1)

The names of the roles used (ob-ject, sb-subject, agent, patient, theme, etc...) will not matter for us here. Deciding which roles will be used with which concepts is a major problem in computational linguistics. We bypass this problem by assuming that roles are assigned in a consistent, coherent and maximally informative way. We write collections of clauses that correspond to representations of natural language sentences and hence to propositions.

Linguists and computational semanticists are used to a variety of styles of 'event semantics', but perhaps not one exactly like the one here. Most of the linguistic event semantics deal with individuals, while our basic notion is the one of *subconcept* of a concept. This gives us some amount of underspecification of concepts that is useful when dealing with intensional notions, like negation. This system is also similar to description logics, but not exactly one. In particular we are not restricted to unary predicates and roles on these only.

One crucial feature of the system TIL is its use of *contexts* and its approach to modelling negation, implication and quantification as well as propositional attitudes and other intensional phenomena.

There is a first initial context (written as `t`) that corresponds roughly to what the author of the sentence takes the world to be like. More precisely, in an interpretation of a sentence, the top context corresponds to what the author of the sentence is committed to, with respect to what the world she is describing is like. But since this circumlocution is awkward, we will usually talk about this top level context as the 'true context'. From a practical perspective, contexts in our logic were conceived as syntactic ways of dealing with intensional phenomena, including negation and non-existent entities. They support making existential statements about the existence and non-existence in specified possible worlds of entities that satisfy the intensional descriptions specified by our concepts. It is clear that intensional notions are required when dealing with the representation in logic of propositional attitudes.

Traditional *propositional attitudes* predicates relate contexts and concepts in our logic. Thus a concept like 'knowing' or 'believing' or 'saying' introduces a context that represents the proposition that is known, believed or said. For example, if we want to represent the sentence *Ed denied that the diplomat arrived*, we will need concepts for the arriving event, for the denying event, for the diplomat and for Ed. And we will need roles that describe how these concepts relate to each other. Thus we need to say who did the 'denying' and 'what was denied' and who did the arriving. The content of what was denied in the denying event is the proposition corresponding to the contents of *The diplomat arrived*. The role corresponding to 'what was denied' relates a dynamic concept, the concept of the denying event, to (the contents of) a new context. This new context could be called simply c_1, but we use instead the 'context head' to name it.

While it may seem uncontroversial that "propositional attitude predicates" require intensional notions, we use contexts also for negations, disjunctions and conditionals, which may be considered unnecessarily complex. For example, for the sentence: **R2** *No boys slept*.

Conceptual Structure:
role(cardinality-restriction,boy-5,no)
role(sb,sleep-6,boy-5)
subconcept(boy-5,[List1])
subconcept(sleep-6,[List2])
Contextual Structure:
context(ctx(sleep-6)),
context(t)
context-lifting-relation(antiveridical,t,ctx(sleep-6))
context-relation(t,ctx(sleep-6),not-29)
instantiable(boy-5,ctx(sleep-6))
instantiable(hop-6,ctx(sleep-6))
top-context(t)
uninstantiable(sleep-6,t)

Contexts here are similar to the formal objects with the same name, discussed by J. McCarthy. Contexts are used to 'fence-off' concepts and corresponding roles. Contexts allow us to localize reasoning: the existence of the denying event and of Ed are supposed to happen in the true world, but the existence of the arrival of the diplomat is only supposed to happen in the world of the things denied by Ed. In particular the arrival event should be considered as not happening, if Ed is a reliable source. (The system takes no position as to the instantiability or not of the arrival event in the top context, the instantiability of the arriving is only stated in the context of the things denied by Ed, $ctx(arrive : 4)$.) In some cases (for example if the sentence was *Ed knew that the diplomat arrived*) we can percolate up the truth of assertions in inner contexts up to the outside context. In many cases we cannot. The happening or not of events is dealt with by the *instantiability/uninstantiability* predicate that relates concepts and contexts.

While we may be prepared to make the simplifying assumption that if 'X is known' than 'X is true', we certainly do not want to make the assumption that if 'X is said' than 'X is true'. We say that the context introduced by a knowing event is *veridical* with respect to the initial context t, while the context introduced by a saying event is *averidical* with respect to the initial context. Negation introduces a context that is *anti-veridical* with respect to the original context. Having introduced notions of veridicality, antiveridicality and averidicality between contexts we have expanded the expressive power of our language of representations considerably. Thus we have a fairly general mechanism of contexts (these can clearly be iterated), which can represent some positive and some negative information. Similarly to McCarthy's logic we also have 'context lifting rules' that allow us to transfer veridicality statements between contexts, in a recursive way. A precise description of the algorithm explaining how these context lifting rules work for specific classes of verbs (marked in the lexicon) can be found in [Nairn, Condoravdi, and Kartunnen., 2006].

Given that our logical representations are so intimately connected to the underlying ontology, one might expect that the change of ontology from CYC to Wordnet/Verbnet or any other ontology would necessitate a total reworking of the system Bridge. This turned out not to be case, the re-architecture of the system was surprisingly easy (from CYC to XLE-UL) and now I expect it to be almost trouble-free for a different ontology. The achievement was making the ontology a parameter of the system that can be exchanged at will, given that other ontologies will be mapped to via common words in any case.

Inference in TIL

The reason for introducing event concepts was the fact that they make some inferences that can be complicated in other semantical traditions very easy. For example it is obvious how to obtain *Ed arrived in the city* from the sentence *Ed arrived in the city by bus*. This inference corresponds simply to conjunction dropping in our logic. But of course there is much more to textual inference than simply dropping of conjuncts.

Inference in TIL is very rudimentary. We can 'drop clauses' like in most event semantics. From the sentence *Ed walked and Mary talked* we are able to infer both *Ed walked* and *Mary talked* by simply forgetting the respective clauses in the original representation.

We can do trivial inferences like identity and we can compose derivations:

$$\frac{}{s \to s} \qquad \frac{s \to t \qquad s \to r}{r \to t}$$

But note that the clauses we construct also satisfy the usual monotonicity patterns, both in positive and in negative form. Thus *Ed arrived in the city by bus* entails that *Ed arrived in the city*. But *Ed did not arrive in the city* entails that *Ed did not arrive in the city by bus*, while *Ed did not arrive in the city by bus* does **not** entail that *Ed did not arrive in the city*. Similarly a limited amount of 'going up and down the taxonomy' can be accounted for this way, using monotonicity markers. So *Ed arrived in the city* does entail that *A person arrived in the city*, since Ed is a person. Similarly *Ed arrived in Rome* should entail that *Ed arrived in a city*, as Rome is a city.

More importantly we concentrate efforts into using the context structure of our logic to provide inferences associated with kinds of verbs with implicative behavior. This work is discussed in [Nairn, Condoravdi, and Kartunnen., 2006]. Here we simply give an example of each one of the classes of "implication signatures" or implicative behavior described by Nairn, Condoravdi and Karttunen. There are six such classes, depending on whether positive environments are taken to positive or negative ones. Thus for example the verb "manage" takes positive predicates (e.g "Ed managed to close the door" → "Ed closed the door") to positive predicates and negative ones ("Ed didn't manage to close the door" → "Ed didn't close the door"). By contrast the verb "forget (to)" inverts the polarities: "Ed forgot to close the door" → "Ed didn't close the door" and "Ed didn't forget to close the door" → "Ed closed the door".

More complicated are the verbs that only show their implicative behavior either in positive or negative situations. For example we have positive implicatives like the verb "force (to)" takes positive polarities and produces positive polarities (e.g "Ed forced Mary to paint" → "Mary painted"), but if "Ed didn't force Mary to paint" we cannot tell whether Mary painted or not. While "refuse (to)" only works to produce negative polarity (e.g. 'Mary refused to sing" → "Mary did not sing"). There are also negative implicatives like "attempt (to)" (if you say "Ed didn't attempt to paint" this implies "Ed didn't paint", but nothing can be said about the positive version) and "hesitate (to)" which again only work for a negative polarity, but produces a positive one ("Ed didn't hesitate to leave" → "Ed left"), but if "Ed hesitated to leave" we cannot tell whether he left or not.

Finally we have factives and counterfactives, examples being "forget (that)" ("Ed forgot that Mary left" → "Mary left" and "Ed didn't forget that Mary left" → "Mary left" and "pretend that" ("Ed pretended that Mary left" → "Mary didn't leave" and "Ed didn't pretend that Mary left" → "Mary left"). And the neutral class, where we cannot say anything about the veridicality or otherwise of the complement ("Ed said/expected that Mary left"). Further work, to mark implicative behavior of verbs that do not take sentential complements has been also undertaken, but much more needs to be done.

After this 'impressionistic' discussion of the logical system TIL in an ideal world we would produce a proof system for it and prove some of its properties. But coming from these three distinct sources of inspiration (event semantics with description-logic like concepts plus McCarthy-like contexts), and having been built with efficiency of processing, readability of output and overriding respect for semanticist intuitions, this is not an easy task. This is not an easy task, because the sources of intuitions are very diverse and they do not fit together particularly well. One can have a logic of contexts as modalities over propositional logic easily, but contexts over *constructive* propositional or first-order logics are less well-understood [de Paiva, 1996]. Similarly I have not seen descriptions of constructive (or intuitionistic) event semantics. Even not considering the constructivity aspect, I have not seen a calculus with events and contexts in print. My expectation is that concepts can always be translated away to FOL easily, but the choice between the kinds of quantified calculi of contexts that one can have are not exactly trivial, even in the classical case. There are many design decisions to make as far quantified modal logics are concerned and the trade-offs between expressivity and tractability of implementation are hard to make (for a recent reasonable suggestion see Benzmüller's work [Benzmüller, 2013].

A reviewer suggests Distributed Description Logic [Borgida and Serafini., 2003] and its development [Serafini and Homola, 2012], a specific form of distributed version of description logics, that uses the mechanism of *bridge rules* as a model of contexts.This is an interesting system, but thoroughly classical in its way of dealing with negation, disjunction and implication. This reviewer also suggests (alternatively) a classical version of contexts over description logics[Klarman and Gutierrez-Basulto, 2010] which I was not aware of, but looks a well-developed, classical version of [de Paiva and Alechina, 2011], in that we have a collection of modal contexts over a description logic basis. But for TIL we would prefer to have modal, constructive contexts over a basis that looked more like a quantifier-free constructive logic of events. This makes for the readability of the representations, a crucial factor when it is not yet clear that we can cover all the kind of language phenomena that we want to deal with.

Hence we instead take the ”translation route” and aim to show that some of the fragments of Natural Logic considered in the literature can be translated into the system TIL.

Mapping to Syllogistic Systems

The program of Natural Logic can be characterized as done in [Pratt-Hartmann, 2004] as follows

> By a *fragment* of a natural language we mean a subset of that language equipped with semantics that translate its sentences into some formal system such as first-order logic. The familiar concepts of *satisfiability* and *entailment* can be defined for any such fragment in a natural way. The question therefore arises for any given fragment of a natural language, as to the computational complexity of determining satisfiability and entailment within that fragment.

The characterization above presupposes notions of satisfiability and entailment defined in terms of mathematical models. But surely once sentences of a fragment of natural language are translated into a formal system we can also define for this fragment notions of entailment and derivability based only on rules of derivation, the proof-theoretic way. This should work whether one is concerned with computational complexity or not. The label of *proof-theoretic semantics* has been coined for a general version of this program. What we are after is a proof-theoretic semantics of TIL that would help us measure not only the computational complexity of determining entailment within the system, but also its expressivity and coverage when compared to other computational semantics formalisms.

Following on the footsteps of Moss and Ian Pratt-Hartmann we can start by looking at how TIL deals with syllogisms. As far as semantics is concerned, TIL is very much like Moss' systems.

For example: **R3** *All boys are mammals.*
Conceptual Structure:
role(cardinality-restriction,boy-2,all(pl))
role(cardinality-restriction,mammal-4,pl)
role(copula-pred,be-3,mammal-4)
role(copula-subj,be-3,boy-2)
subconcept(boy-2,[List1])
subconcept(mammal-4,[List2])
Contextual Structure:
context(t)
instantiable(be-3,t)
instantiable(boy-2,t)
instantiable(mammal-4,t)
top-context(t)

Concepts in TIL are interpreted as subsets of a collection of sets, so

subconcept(boy-2,[List])

means that the boys referred to by the skolem boy-2 are a subset of all the boys in the universe in the senses of the word 'boy' recognized in our ontology. Similarly the eating concept in our sentence is a subset of the eating events in the universe.

subconcept(eat-3,[List])

This is a bit non-intuitive for proper names, but we still use subsets and not individuals. This semantics agrees with Moss' semantics for the system of syllogistic logic in [Pratt-Hartmann, 2004].

Syllogistic Logic of *All*

This is the simplest fragment considered by Moss. Its syntax consists of a collection of unary atoms (for nouns). The sentences are the expressions of the form *All p are q*, only.

The semantics is given by the following definition: A model $\mathcal{M}$ is a collection of sets M, and for each noun p we have an interpretation $[\![p]\!] \subseteq M$.

$$\mathcal{M} \models \textit{All p are q} \qquad \text{iff} \qquad [\![p]\!] \subseteq [\![q]\!]$$

The proof system for this fragment is given simply by the two rules:

$$\frac{}{\textit{All p are p}} \qquad\qquad \frac{\textit{All p are n} \qquad \textit{All n are q}}{\textit{All p are q}}$$

The system TIL satisfies trivially the rules for the *All*-fragment. The first 'axiom' corresponds to examples such as *All boys are boys*, which is an odd sentence, but not logically problematic.

Semantically the second rule is just transitivity of subset containment. For TIL the transitive inference is simply climbing up the concept hierarchy.

$$\frac{\textit{All boys are mammals} \qquad \textit{All mammals are animals}}{\textit{All boys are animals}}$$

Syllogistic Logic of *Some*

The sentences are only the expressions of the form *Some p are q*. The semantics is given by a definition similar to the above: A model $\mathcal{M}$ is a collection of sets M, and for each noun p, q we have an interpretation $[\![p]\!], [\![q]\!] \subseteq M$, such that:

$$\mathcal{M} \models \textit{Some p are q} \qquad \text{iff} \qquad [\![p]\!] \cap [\![q]\!] \neq \emptyset$$

The proof system has three rules:

$$\frac{\textit{Some p are q}}{\textit{Some q are p}} \qquad \frac{\textit{Some p are q}}{\textit{Some p are p}} \qquad \frac{\textit{All q are n} \quad \textit{Some p are q}}{\textit{Some p are n}}$$

The first rule is confusing, if one is thinking of traditional quantifiers and of *Some* as the existential quantifier, as existentials are not symmetric. But if

$$\frac{\textit{Some doctors are women}}{\textit{Some women are doctors}}$$

is a valid inference. The second rule seems a grounding rule and the third rule

$$\frac{\textit{All doctors are rats} \quad \textit{Some women are doctors}}{\textit{Some women are rats}}$$

is at least parallel to the rules for modalities in constructive logic.

The system TIL satisfies all the rules for the *Some*-fragment. The semantics is the same as Moss', non-empty intersection of subsets.

The inference relation between 'all' and 'some' can be outsourced in TIL. Universal quantifiers are considered a cardinality restriction: as we saw in the example *All boys are mammals*, we treat *all* as a kind of measure of the subset of the the boys that are mammals, and this measure is the whole subset under discussion. This allows us to have some *arithmetic of cardinality* that is convenient for quantities (*three boys slept* entails that *two boys slept*) and also, as a side-effect, allows us to leave to the designer of the system the option of deciding whether to have existential import or not, i.e. whether to read "for all" meaning only over non-empty domains or not.

The Bridge system has a 'poor man's' inference system called Entailment and Contradiction Detection (ECD), which has not been properly described in the literature (only the patent is available) but where some of the choices above are employed. ECD has a table of relationships between the 'cardinality-restrictions' that are postulated. This is useful if you want to keep your logic options open, but the trade-offs with more conventional systems have not been investigated.

Syllogistic with noun-level negation?

One can add complemented atoms $\overline{p}$ on top of the language of *All* and *Some*, with interpretation via set complement: $[\![\overline{p}]\!] = \mathrm{M} \setminus [\![p]\!]$.

But if one has a system like

$$\left.\begin{array}{l} \mathcal{S} \quad \left\{ \begin{array}{l} \textit{All } p \textit{ are } q \\ \textit{Some } p \textit{ are } q \\ \textit{All } p \textit{ are } \overline{q} \equiv \textit{No } p \textit{ are } q \\ \textit{Some } p \textit{ are } \overline{q} \equiv \textit{Some } p \textit{ aren't } q \end{array} \right. \\ \\ \qquad \textit{Some non-}p \textit{ are non-}q \end{array} \right\} \quad \mathcal{S}^{\dagger}$$

things can get strange, as explained by Moss-Hartmann in [Pratt-Hartmann and Moss, 2009]. Instead of adding negation as they do (as it is a classical version of negation) we next consider only *positive* versions of their system.

Positive sillogistic $\mathcal{S}^p$ + names

The language $\mathcal{S}^p$ is the language of *All* and *Some*, with no negation, where the superscript p is used for for positive syllogisms.

$$\frac{}{\textit{All } p \textit{ are } p} \qquad \frac{\textit{Some } p \textit{ are } q}{\textit{Some } p \textit{ are } p} \qquad \frac{\textit{Some } p \textit{ are } q}{\textit{Some } q \textit{ are } p}$$

$$\frac{\textit{All } p \textit{ are } n \quad \textit{All } n \textit{ are } q}{\textit{All } p \textit{ are } q} \qquad \frac{\textit{All } n \textit{ are } p \quad \textit{Some } n \textit{ are } q}{\textit{Some } p \textit{ are } q}$$

$$\frac{}{J \textit{ is } J} \qquad \frac{J \textit{ is } M \quad M \textit{ is } F}{J \textit{ is } F} \qquad \frac{J \textit{ is a } p \quad J \textit{ is a } q}{\textit{Some } p \textit{ are } q}$$

$$\frac{\textit{All } p \textit{ are } q \quad J \textit{ is a } p}{J \textit{ is a } q} \qquad \frac{M \textit{ is a } p \quad J \textit{ is } M}{J \textit{ is a } p}$$

The system TIL satisfies the rules in $\mathcal{S}^p$ + names. The names component of the syllogistic rules (below the line) correspond to sentences like:

1. *Jon is Jon*;
2. *Jon is Mary* and *Mary is Fred* entails *Jon is Fred*;
3. *Jon is a man* and *Jon is a doctor* entails *Some men are doctors*;
4. *All men are mammals* and *Jon is a man* entails *Jon is a mammal*;
5. *Mary is a cat* and *Jon is Mary* entails *Jon is a cat.*

The system TIL was not conceived to interpret syllogisms. Mapping copula constructions is but a small part of it. Event semantics in general is about transitive, intransitive, ditransitive, etc. verbs and their modifiers. The system TIL should be able to deal with $\mathcal{S}^p$ and its relational cousin $\mathcal{R}^p$ (where relations interpret transitive verbs) and much more. How do we measure this much more? And how do we prove this conjecture? In principle noun negation should not present problems, as it can be outsourced to the ontology (a non-student is a person who is not a student) but sentential negation, as we saw, is dealt in TIL via *negative contexts*.

Contexts and their possible-worlds intuitions send us looking for modal and hybrid logics, which carve out decidable tracts of first-order logic, but do so via translations. In previous work we suggested that McCarthy's contexts should be interpreted as modalities in constructive propositional logic [Mendler and de Paiva, 2005], using multi-modalities in the system we dubbed CK_n. While CK_n does model notions of context inspired by propositional attitudes and possibly some of the implicative behaviors of complemented verbs, it does so at the wrong level of granularity for the TIL system, as modalities in CK_n are applied to full sentences, while we want to apply them to fragments of sentences.

Further work

The work on TIL and on proof theoretical semantics for Natural Logic is just starting. To the extent that Natural Language semantics must be experimental in its foundations, the project requires implementations and analysis of real life sentences to justify its claims of suitability for computational linguists. But as an application of logic, especially proof theory, to a new area and its ability to generate new problems and solutions in logic itself, this approach has already borne some fruits. A thorough mapping of the relationships between basic constructive modal logics has begun, the existence of (versions of) constructive hybrid logics has been established, as has the existence of constructive description logics. and even of contextual constructive description logics Complexity results for these new logics are being investigated.

Summarizing the work on TIL so far we can say that the logical system comes from the confluence of linguistic intuitions from event semantics, description-logic-like systems and McCarthy's logics of context. Quantification is mostly done via the new notion of instantiable (or uninstantiable) concepts in given contexts. This mechanism allows us to deal with non-existing entities, as well as with intensional predicates in what we claim is an inference-friendly way. We can see that translations of the fragments originally considered by Moss are soundly interpreted in TIL, but the fact that our system starts from a constructive propositional basis causes some problems. Much remains to be done to provide an actual proof system for TIL and to substantiate the claims of its suitability for computational semantics. But in particular the work on sub-sentential semantics started by Francez and others seems to hint at constructivity as an essential feature of these kinds of systems for Natural Logic.

References

Benzmüller, C. 2013. A top-down approach to combining logics. In *5th International Conference on Agents and Artificial Intelligence (ICAART 2013), Barcelona, Spain*.

Bobrow, D.; de Paiva, V.; Condoravdi, C.; Crouch, R.; Karttunen, L.; King, T.; Nairn, R.; and Zaenen, A. 2005. A basic logic for textual inference. *In Procs. of the AAAI Workshop on Inference for Textual Question Answering, Pittsburgh PA* 27.

Borgida, A., and Serafini., L. 2003. Distributed description logics: Assimilating information from peer sources. *Journal of Data Semantics*.

Condoravdi, C.; Crouch, D.; Everett, J.; de Paiva, V.; Stolle, R.; Bobrow, D.; and van den Berg, M. 2001. Preventing existence. In Welty, C., and Smith, B., eds., *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, 17–19.

Crouch, D., and King, T. 2005. Unifying lexical resources. In *Proceedings of the Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes*.

de Paiva, V., and Alechina, N. 2011. Contextual constructive description logics. In *Proceedings of ARCOE2011*.

De Paiva, V.; Bobrow, D.; Condoravdi, C.; Crouch, R.; Karttunen, L.; King, T.; Nairn, R.; and Zaenen, A. 2007. Textual inference logic: Take two. *Proceedings of the Workshop on Contexts and Ontologies, Representation and Reasoning* 27.

de Paiva, V. 1996. Natural deduction and context as (constructive) modality. In *Modeling and Using Context, Proceedings of the 4th International and Interdisciplinary Conference CONTEXT 2003, Stanford, CA, USA*, Springer Lecture Notes in Artificial Intelligence, vol. 2680, June 2003.

Klarman, S., and Gutierrez-Basulto, V. 2010. ALC–ALC: A context description logic. In *Proceedings of JELIA-10*.

Lasersohn, P. N. 2006. Event-based semantics. *Encyclopedia of Language and Linguistics, 2nd edition, ed. by K. Brown* 4(1):316–320.

Lenat, D.; R.V.Guha; K.Pittman; D.Pratt; and Shepherd, M. 1990. Cyc: Towards programs with common sense. In *Communications of the ACM, Vol.33, No.8*.

Maxwell, J., and Kaplan, R. 1996. An efficient parser for lfg. In *Proceedings of the First LFG Conference*, CSLI Publications.

McCarthy, J. 1993. Notes on formalizing context. In *Proc. of the 13th Joint Conference on Artificial Intelligence (IJCAI-93)*, 555–560.

Mendler, M., and de Paiva, V. 2005. Constructive ck for contexts. *Proceedings of the Worskhop on Context Representation and Reasoning, Paris, France*.

Moss, L. 2010. Logics for two fragments beyond the syllogistic boundary. *Fields of Logic and Computation: Essays Dedicated to Yuri Gurevich on the Occasion of His 70th Birthday* 6300(1):538–563.

Moss, L. 2012. Inference in a boolean fragment. *Theories of Everything: In Honor of Ed Keenan (eds. Thomas Graf, Denis Paperno, Anna Szabolcsi, and Jos Tellings). UCLA Working Papers in Linguistics 17* 261–273.

Nairn, R.; Condoravdi, C.; and Kartunnen., L. 2006. Computing relative polarity for textual inference. *Proceedings of ICoS-5 (Inference in Computational Semantics). April 20-21, 2006. Buxton, UK.*

Niles, I., and Pease, A. 2001. Towards a standard upper ontology. In Welty, C., and Smith, B., eds., *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, 17–19.

Pratt-Hartmann, I., and Moss, L. 2009. Logics for the relational syllogistic. *Review of Symbolic Logic, Vol. 2, No. 4* 647–683.

Pratt-Hartmann, I. 2004. Fragments of language. *Journal of Logic, Language and Information, 13(2)* 207–223.

Serafini, L., and Homola, M. 2012. Contextualized knowledge repositories for the semantic web. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*.