

Truth or *dare**

ANDRÁS KORNAI

22.1 Background

Thirty years ago, September 1986 to be precise, when I arrived at Stanford, Dwight Bolinger threw a barbeque for the incoming linguistics class. I was absolutely shocked. Bolinger was a demigod¹ of linguistics, someone I knew only through his work and stories, both transmitted to me by Ferenc Kiefer, and now I get to talk to him in person, eat his pulled pork, and drink his beer? I was delighted to see that, retired or not, he was just as sharp in real life as his work. There were plenty of CSLI people and other folks loosely affiliated with the linguistics department around – this was the first time I met Lauri. When my student Gábor Recski came back from *Sem 2015, he told me a very similar story: in a sea of younger faces there was this older guy, asking very sharp questions about all aspects of the work, and he was stunned to find out this was Lauri Karttunen himself! Somehow, some of the interests and esthetics of the older generation will rub off on the students, and (largely unbeknownst to me) a great deal of what Lauri thought was interesting I also ended up being interested in.

Following on many years of research in this direction, recently Kart-

*Thanks to Tibor Beke (UMass Lowell), Tim Fernando (Trinity College Dublin), Ferenc Kiefer (HAS Research Institute of Linguistics), Márton Makrai (HAS RIL), Chris Piñón (Université Lille 3), Gábor Recski (HAS RIL) and Károly Varasdi (Henrich Heine University, Düsseldorf) for cogent criticism of the draft version.

¹http://zvon.org/comp/r/ref-Jargon_file.html#Terms~demigod

tunen (2014) offered a force dynamics²-like analysis of what he calls *two-way implicatives* following Karttunen (1971), verbs like English *manage* and Finnish *hennoa*, the implication of which Karttunen illustrates with *Hennoitko tappaa kissan?* ‘Did you overcome your pity to kill the cat?’, akin to ‘overcome your fear’ for *dare*. Here we will look at a finite-state account of these verbs, finally bringing together two of the main threads of Lauri’s research so far held together only by the esthetics. While hard to put in words exactly, clearly this esthetics has something to do with *mechanization*, going back at the very least to the the Calculus Ratiocinator of Leibniz, himself a tinkerer, building his Stepped Reckoner³ from cogwheels the same way as Lauri builds his morphophonology from FSTs. For Leibniz, the elementary pieces were Monads, for Karttunen they are finite mechanisms which must be, in a sense we will make more precise in Section 22.4, *memoryless*. The luxury of a Turing Machine *tape* is not given to us, everything we remember we must remember at great cost, building it into the state space of the machinery as we go along. Crucially, there are no variables, and there is no variable binding, so who does what to whom will need to be specified by a different mechanism. We will use *cases* for this purpose.

Frege (1892) already notes that the failure of the implicatures that such verbs carry will not render false the sentence that uses the verb, and most subsequent authorities, including Karttunen, share this view. What needs to be made explicit in this regard is that sentences whose meaning cannot be tied to truth conditions (other than truth conditions pertaining to the mind-state of the speaker and the hearer) actually demonstrate that truth-conditional semantics is a blunt instrument, incapable of assigning meaning to sentences. Indeed, our main thesis here will be that either we are interested in truth, or we are interested in what verbs like *dare* mean.

Our goal in this paper is to formalize a lexically driven analysis in terms of a mechanical, finite state calculus. Since such an analysis, practically the only one that makes sense from the perspective of the grammar, is compatible only with a weak form of commonsense reasoning, as opposed to a muscular Tarski-style theory of truth, we obtain the result claimed in the title.

²https://en.wikipedia.org/wiki/Force_dynamics

³https://en.wikipedia.org/wiki/Stepped_Reckoner

22.2 The elementary pieces

The good thing is that we already know what *dare* means: ‘to be brave enough to do something difficult or dangerous’ (Cambridge Dictionary of English); ‘to be brave enough to do something that is risky or that you are afraid to do’ (Longman); ‘to have enough courage or confidence to do something, to not be too afraid to do something’ (Merriam-Webster). This is a special case of the general analysis that Karttunen offers for the whole class of verbs: ‘overcoming an obstacle’ with the obstacle being *fear* for *dare*, *indifference* for *bother*, *empathy* for *hennoa* and so on.

Here and in what follows we will rely heavily on the ‘algebraic’ style of lexical semantics first described in Kornai (2010a) and Kornai (2012) and now spelled out in far greater detail in Kornai (in press) and in a series of more computationally oriented papers such as Kornai et al. (2015). Remarkably, no part of the apparatus employed in this paper was developed with implicative verbs in mind.

We will use hypergraph-style notation with two kinds of concepts, unaries, which we will write in `typewriter` font and binaries, written in SMALL CAPS. There are only three types of edges: 0 used both for attribution *John is brave* and for IS_A indiscriminately. In larger graphs, we will write dashed arrows, $- \rightarrow$ instead of $\overset{0}{\rightarrow}$. The two other arrow types are the familiar ‘1’ (ordinary \rightarrow in larger graphs) for subject, and ‘2’ (dotted arrow) for object. For ease of presentation, we introduce a special symbol @ that will be placed in the middle of arrows that should *in their entirety* be the terminal point of some other arrow, so that we can display the object of seeing in *video patrem venire* as




Figure 1. *Video patrem venire*

John dares VP is taken to mean *John does VP* in conjunction with *VP is risky*, and for the moment we leave open the issue whether it’s risky for John, or really risky for everybody. What we did here was to incorporate a hidden element, ‘object being risky’ for *dare*, ‘object being boring’ or ‘subject being indifferent’ for *bother*, and ‘subject being humane’ for *hennoa*. As it happens, the calculus we use is detached from part of speech, and we make no distinction between the adjective *risky*, the V’ *is risky*, and the noun *risk*, but this lack of morphosyntactic

typing will not be relied on heavily in what follows, except for making it easier to draw the graphs and talk about their nodes.

A significant advantage of relying on such hidden elements is that they all fit in the category of **difficulty**. As a result, they will participate in a larger frame involving **virtue** $\xleftarrow{1}$ **OVERCOME** $\xrightarrow{2}$ **difficulty**. The virtue, be it **bravery** as in the Cambridge and Longman dictionaries, or **courage/confidence**, as in Merriam-Webster, comes for free here, in that **bravery** $\xrightarrow{0}$ **virtue**, **diligence** $\xrightarrow{0}$ **virtue**, **decency** $\xrightarrow{0}$ **virtue**, etc. come from the lexical definition of these concepts, and no doubt **winter-hardiness**, required in the analysis of *tarjeta*, is a virtue, perhaps a monomorphemic one in Finnish, if Sapir-Whorf has any bite.⁴

Now it is plausible that **fear** $\xrightarrow{0}$ **difficulty** (or **obstacle**, as Karttunen (2014) has it) is somehow a piece of lexical knowledge, but this does not work for the other cases: it is rather unlikely that **indifference** or **empathy** are lexically specified for genus **difficulty/obstacle**. Such a conclusion, if not available lexically, must somehow be derived by a process of typecasting. The essence of the force-dynamic analysis is that *dare* and *hennoa* both IS_A **OVERCOME**. There is every reason to suppose that **OVERCOME** subcategorizes for a **power** subject and an **obstacle** object, and to make some action an instance of **OVERCOME**, its subject must be typecast to **power** and its object to **obstacle**. That **virtue** IS_A **power** is hard to deny, and by transitivity of IS_A we can treat all implicative subjects under this heading. With the implicative objects, this is less trivial: **empathy** gets to be an obstacle only because *hennoa* IS_A **OVERCOME**, and outside this frame we cannot draw the usual conclusions, e.g. that obstacles are bad things and therefore being humane is a bad thing.

Before turning to the details of the finite state analysis, let us consider how an ordinary sentence, such as *John dared to criticize the mayor* will be analyzed. The matrix verb is *dare* and it is John who does the daring, so we have **John** $\xleftarrow{1}$ **DARE**, and it is also John who does the criticizing, so we have **John** $\xleftarrow{1}$ **CRITICIZE**. How the subject equi is effected during the parsing process is something we leave to the phenogrammar, the point here is that few grammarians (including LFGers who would use an XCOMP here) would seriously doubt that the subject of both verbs is the same **John**. The object of the criticism is no doubt the **mayor**, and the object of daring is the entire *criticizing*

⁴Inquiries to Finnish-speaking friends produced *pakkasenkestävä*, but with a cautionary note that this is probably more appropriate for plants than for animals or humans, and *kylmänkestävä* with perhaps a bit less strict subcategorization.

the mayor.




Figure 2. *John dared to criticize the mayor*

As we said, *dare* IS_A *overcome*, and whoever *dares* actually HAS power, or at least he thinks he does, to overcome risk. In fact, we can rely on a dictionary definition of *bravery*, *courage* as ‘power to overcome risk’ or ‘power to overcome fright caused by risk’, or even ‘power to overcome one’s own fright caused by risk’. Neither the final cause nor the precise application site of the power ends up being very relevant for the task at hand, which is to explain certain implications, whose non-fulfillment makes sentences using *dare* infelicitous, but not outright false. With *criticizing the mayor* it is rather clear that mayors are powerful people, and criticizing the powerful is dangerous. But when we say

- (1) #*John dared to chew gum*

we need to abductively infer some theory that makes chewing gum dangerous. Perhaps John had throat surgery, and the wounds have not quite healed. Perhaps he is in the presence of some superior who considers this disrespectful. Perhaps he was told the gum could be laced with poison. There are many theories that would make the use of *dare* felicitous, and we need not choose among them. But we do need to draw the implication from *dare* to risk or danger. As with *obstacle* versus *difficulty* above, we need not be very precise which of these terms is operative here. What matters is the semantic concept, not the (English) printname we assign to it.

We will posit $\text{DARE} \xrightarrow{2} \text{danger}$ as part of the lexical entry of *dare*, i.e. the selectional restriction that its object is dangerous. Similarly, the object of *deign* is low status (or perhaps the subject is high status), that of *remember* is hard to memorize, and so on. *Manage* has an object that is simply difficult, manner unspecified. Since it is the relationship of the subject to the object that is getting characterized by the implicative verb, we often have a choice between alternative framings: e.g. with *deign* we may be describing (i) the subject as high status; (ii) the object as low status; or (iii) the subject as higher status than the object. These alternatives are logically equivalent, since by default things are neither high nor low status. Yet it is quite conceivable that different speakers have different lexical entries for *deign*, and different

lexicalization options may play differently with negation. Translational near-equivalents in different languages may also differ only in the choice between (i)-(iii).

22.3 The mechanism

Again we begin with the easy pieces. The theory we rely on takes unaries simply to be bundles of essential properties, so for example the *fox* is **four-legged**, **animal**, **hairy**, **red**, **clever**. Since we make no distinction between attribution (*the fox is hairy*) and IS_A (*the fox is an animal*), we need not worry about the type signature of selectional restriction: by demanding DARE $\xrightarrow{2}$ danger we have accomplished DARE $\xrightarrow{2} X, X \xrightarrow{0}$ danger without lifting a finger. In general (not just for the sake of making this particular case come out right) IS_A is a derived notion, simply there to achieve economy in the system by means of lazy default inheritance. If x IS_A y , then the small set of essential properties associated to x is a superset of those associated to y , and elementary pieces of link-tracing logic, such as x IS_A $y \wedge y$ IS_A $z \Rightarrow x$ IS_A z , or x IS_A $y \wedge y$ HAS $z \Rightarrow x$ HAS z follow without any stipulation.

We make purposely very little distinction between an individual fox, the species *Vulpes vulpes*, the set of foxes in the world, or the class of potential foxes in all possible worlds. Even though such differences can be sufficient for distinguishing different senses of the same word, (compare *Jack is blond* ‘has blond hair’ to *Jack is a blond* ‘fulfills the criteria for the stereotype’), this is a peculiarity of English syntax, and as such it has no place in the semantics.

What implicative verbs bring to the table (working memory) are little lexically prespecified hypergraphs that demand abductive inferencing over and above the normal inferencing process, which we take to be (hyper)graph unification. We assume, as is standard, that verbs can subcategorize for their arguments: for example *elapse* demands a time interval subject. If we read that *A sekki elapsed* we know that this must refer to some time period even if we do not know the details of the sekki⁵ system. This is part and parcel of knowing what *elapse* means: people who cannot make this inference are not in full possession of the lexical entry. In the representation of *elapse* there is thus a direct prespecification $\text{time_period} \xleftarrow{1} \text{elapse}$, which contrasts with the prespecification inherited from a nominative linker NOM that subjects are agentive (active, causing, volitional).

There is little reason to suppose that time intervals are inherently

⁵<http://eco.mtk.nao.ac.jp/koyomi/faq/24sekki.html.en>

active, that they causally contribute to their own elapsing. To the contrary, *time_interval* IS_A *abstract_object* and would therefore inherit features of the latter such as LACK $\xrightarrow{2}$ *volition* that would directly contradict agentivity. But once *sekki* appears in the subject slot of *elapse*, it IS_A *time_interval* and the inheritance of non-volitionality from *abstract_object* is blocked, since the specific, lexically prespecified case will block the general inheritance mechanism as a matter of course. Similarly, in the representation of *dare*, the object of daring IS_A *risk*, and this will block the general assessment that e.g. chewing gum is not normally considered a risk.

Another aspect of the analysis that is relatively easy in the framework presented here is adding *overcoming* to the small set of preexisting force dynamic primitives *letting*, *hindering*, and *helping*. Since this would take us far afield from the issue at hand, we only sketch the basic mechanism: verbs can, and often do refer to the state *before* and *after* the action they depict: for example *go* simply means the conjunction **before** AT **source**, **after** AT **goal**; while telic verbs carry the specification **after done**. With these pieces at hand, we can dispense with the force-dynamic diagrams entirely in favor of analytic statements such as *overcome* ‘the agent, initially weaker than the object, is subsequently stronger’. Since **before** and **after** are independently justified by other verbal phenomena, as is the generic comparative binary relation ER (abstracted from *greater*, *bigger*, *more*), we can say **before** (**force** $\xleftarrow{2}$ OVERCOME ER **force** $\xleftarrow{1}$ OVERCOME), **after** (**force** $\xleftarrow{1}$ OVERCOME ER **force** $\xleftarrow{2}$ OVERCOME). As usual, it matters but little whether we call the basis of the comparison *force*, *power*, *might*, *heft*, *momentum*, or something else, there is a single concept here, and we chose to call it **force** mainly to make clear our indebtedness to Talmy (1988) and Jackendoff (1990).

In a standard system of logic such as first order predicate calculus (FOPC) we would need at least two variables *x* and *y*, a one-place predicate *force* and seven two-place predicates *SubjectOf*, *ObjectOf*, *Before*, *After*, *IsA*, *Has*, and *>*, to express approximately the same meaning in a conjunctive formula:

- (2) *Before*((*x IsA force* & *SubjectOf(x,overcome)*) & *y IsA force* & *ObjectOf(y,overcome)* & *y > x*),*overcome*) & *After*((*x IsA force* & *SubjectOf(x,overcome)*) & *y IsA force* & *ObjectOf(y,overcome)* & *x > y*),*overcome*)

even with the dubious expedient of reusing *x* and *y* in the *Before* and *After* subformulas. To simplify matters, notice that the variable *x* is

just a paraphrase ‘the force of the subject of overcome’ and similarly y is ‘the force of the object of overcome’ so what we need to handle first are binary relations that have a first and a second argument.

As readers of Kornai (2010a) will know, we will use *machines* in the sense of Eilenberg (1974) to do this work. A binary machine will have three *partitions*. The 0th or common partition will store all properties that hold for the entirety of the machine, e.g. that its printname is OVERCOME, or that various things may hold BEFORE or AFTER the overcoming. As with unaries, we do not make strict distinctions between a machine, its printname, or a pointer to the machine. The distinguished 1st partition will store properties true of the subject (1st argument), and the undistinguished 2nd partition to store properties that hold of the object (2nd argument).

This skeletal apparatus is already sufficient for providing the meaning representation of the bound morpheme *-er*: this will be a machine that has printname ER, and the only property it enjoys is that the relation $>$ holds between whatever is stored on its distinguished (1st) and undistinguished (2nd) partition. Largely similar is the machine for the possessive relation: we give it printname HAS (but it could be called OWN or POSS just as well) and this has implications for the subject, seen as the controller, and the active maintainer of the possession relation; and also for the object, seen as the controlled element, the one that cannot end the relation by itself.

We have already discussed unaries, seen as Eilenberg machines with two partitions, the 0th holding the printname, and the 1st (distinguished) partition holding (pointers to) the characteristic properties of the word being modeled (recall fox above). Note that IS_A is not a machine: we say x IS_A y iff the properties associated to x are a superset of those associated to y . Here x and y are variables in the metalanguage, not in the object language. There are no ternary or higher arity predicates – this is discussed in Kornai (2012).

We use *cases* such as NOM and ACC to effect the linking of subjects to the first, and objects to the 2nd partition at least in nominative-accusative languages: for ergative-absolutive this would be different. At this level, the mechanism is very similar to the linking theories of Kiparsky (1987) and Bierwisch (1988), see Chapter 5.5.1 of Butt (2006). The main technical innovation is the elimination of variable binding by means of machine substitution, taking advantage of the extra modeling capability offered by the relational monoid X that is at the base of Eilenberg machines. Members of X are relations (subsets of X^2), and if X has only two elements u, v we can use right multiplication with the one-member relation (v, v) to model binding at u , and right multipli-

cation with (u, u) to model binding at v . This way we can guarantee, without recourse to variables, that the formula (or hypergraph) associated to *Brutus killed Caesar* differs exactly in the expected way from the one associated to *Caesar killed Brutus*.

Putting this all together, we start from highly skeletal lexical entries such as *dare* = **brave**, **do**, and obtain, by finite state means, implications such as *John chewed gum* and *chewing gum was risky (for John)*. Again, it should be noted that the lexical entry for *dare* was not created post festum, just to support this analysis, but existed in the 4lang conceptual dictionary for some years now, see github.⁶ We owe a recent bugfix, shifting the verbal base from *try* to *do*, to the heightened scrutiny the entry received during the writing of this paper, but this does not affect the main point, which is to get from the subject's bravery (as posited in the actual entry) to the object's riskiness (the notion we relied on above).

This inferencing is accomplished by the same process, substitution *salva veritate* of the definition of *brave* in the subject position, so that we obtain 'subject HAS courage'. Now what about *courage*? Again substituting the definition 'will ER fear' we obtain 'subject has greater willpower than fear' and it is by yet another substitution, that **fear** is not just any old mental state, but one that **danger** CAUSE that we finally derive the conclusion that the object of *dare* is indeed dangerous. As in many parts of the lexicon, there can be serious disagreement over how much of this is precomputed and already stored in the lexicon, see e.g. Pinker and Prince (1988), and defending one analysis over the other would take as far afield from the central theme of this paper, which is the finite state mechanism, the weak (proto)logic calculus one can use in deriving the meaning of the whole from the meanings of the parts.

22.4 Memory

Starting perhaps with Yngve (1961), linguists have long wrestled with assessing the impact on sentence processing of the limitations of short-term or working memory (Miller, 1956). The bulk of this work concerns syntax and takes it for granted that the central issue is dealing with the linear succession of words. Island parsing techniques, based on the idea that a full parse may be built from well-understood subgrammars, came two decades later (Carroll, 1983), and it was only under the impact of Ken Church's famous declaration, *parsers don't work*, that interest in partial parses, such as offered by light parsing (Abney, 1991), was

⁶<https://github.com/kornai/4lang>

beginning to be seen as legitimate.

If our focus is with semantics, the defining data structure of sequential processing, the tapes common to FSA, FSTs, and Turing machines, appear neither relevant nor particularly useful. Clearly, humans have huge long-term memory, but there is no reason whatsoever to suppose that this memory is sequentially organized outside of procedural/episodic memory. In particular, the bulk of linguistic information is stored in the lexicon, a device that is best thought of as random access. The classic model of computation best suited to random access of this sort is the Kolmogorov \mathbb{B} -complex (**KBC**), originating with Kolmogorov (1953) – for a more accessible introduction see Ch. 1 of Uspensky and Semenov (1993). There are more modern concepts, in particular the Storage Modification Machine of Schönhage (1980), the Pointer Machine of Shvachko (1991), and the Random Access Computer of Angluin and Valiant (1979) – for a good discussion, see Gurevich (1988).

We are in no position to make a compelling case for one over the other, but the key issue, as emphasized by Gurevich, is that all these models are “more appropriate for lower time complexities like real time or linear time” than the standard Turing Machine. There is in fact an important line of research, starting with Quillian (1967) and today best exemplified by the Abstract Meaning Representation (AMR) theory Banarescu et al. (2013), that shares much of the tools, concepts, and formal underpinnings of the theory described above. All these theories, called ‘algebraic conceptual representation’ (ACR) in Kornai and Kracht (2015), take the lexical entries, and the knowledge representations, to be (hyper)graphs, just as the **KBC** family of models. While AMR operates with hypergraph rewriting, and the theory presented here uses Eilenberg machines to build larger graphs from the elementary (lexical) pieces, **KBC** directly transduces one graph to the next one.

Whichever model we choose (and there is quite a rich variety to choose from, for example Kornai (1987) uses equational systems as defined in Curry and Feys (1958)), the key, most compelling notion these theories have since Quillian is *spreading activation*. This is island parsing writ large, beginning with nouns, named entities, and NPs, detection of case marking, assembly of clausal structure, and verbal slot filling. At every stage, morphemes, words, or larger lexical entries are active, and by spreading activation so are their links. A structure is detected whenever two such spreading waves of activation meet. Pragmatics, in the sense relevant to our understanding of *dare* and other implicatives, is simply an effort to find paths where none initially exist. There is clearly no link, at least initially, from *chewing gum* to *danger*.

But the post-verbal position really compels a reading whereby *chewing gum* is the object of *dare*, so we make the link, and now *chewing gum* is risky. Speakers are of course very aware that such sense-making activity is under way. They use it to eliminate redundancy, and abuse it to set up semantic traps like *when did you stop beating your wife*.

22.5 Conclusions

The lower predicate calculus (FOPC) is a wonder to behold: it is highly expressive, has good model theory, excellent proof theory, and a completeness theorem cementing the link between these two. It enjoys the compactness and Löwenheim-Skolem properties – in fact, according to Lindström’s Theorem, it is the strongest among those systems of logic that have these two properties. Though Montague (1973) had good reasons to reach for higher order calculi, there have been continued efforts to bring natural language semantics down to first order, see in particular Blackburn and Bos (2005). But these efforts do not go far enough to curb the logic’s hunger for resources: the simple task of checking that each quantifier bounds some variable is conjectured (Marsh and Partee, 1984) to be properly context sensitive, outside the power of indexed grammars, let alone the mildly context sensitive class.

Given that the mainstream theory of linguistic quantification still leaves a lot to be desired (Kornai, 2010b), a less ambitious form of logic that reduces quantification to genericity (use of defaults), and sacrifices scope effects may actually be a better fit with the semantics of natural language, and this is what we propose here. The system presented here is not even fully Boolean: conjunctions are everywhere, but disjunction is hard, and negation practically nonexistent, except for overt marking of failure of defaults, for example in taking **blind** to mean LACK $\xrightarrow{2}$ **sight**, or even a unary analysis with **sight** $\xrightarrow{0}$ **lack**. This actually goes a long way towards explaining classic pragmatic puzzles, why *#blind stone* is infelicitous, and how can a *cup* be felicitously defined as a ‘small round container for liquids, with or without a handle’, given that everything is with or without a handle. At the same time it suggests caution with regards to relying on negation as *the* touchstone for implicature, especially as there is a whole lot of inferencing going on in positive contexts already, and the kind of logical semantics that takes Booleans for granted has no traction whatsoever over the central body of primary linguistic data.

Broadly speaking, all ‘algebraic’ approaches (among which we count not just classic AI models and AMR, but also Pāṇini) are like Gen-

erative Semantics⁷ in that they proceed from meaning representation to surface form directly, without any reliance on LF, and view interpretation as the inverse task, analysis by synthesis. Such systems have no room for quantifier scoping,⁸ and have been roundly rejected by AI greats like McCarthy (2005): “people who put knowledge into computers need mathematical logic, including quantifiers, as much as engineers need calculus”. However, putting knowledge into computers is a goal very different from the goals of linguistics, in that people can acquire a great deal of knowledge for which there is no language-level support, starting with elementary arithmetic. To sustain this kind of scientific knowledge needs a logic that is considerably stronger than the proto-logic we rely on here.

The problem with these stronger logics is not so much that they are undecidable (they are), or hard to compute with (all systems in common use carry SAT⁹ and other NP-complete problems), for clever limitations placed on the mechanism can avoid these pitfalls. The problem is that they are *unlearnable* in a strong sense. What we need is not just a model, but a path from primary data to this particular model within a reasonable hypothesis space. The general problem of learning FSA is already incredibly hard (Angluin, 1987), but somehow we all learn what *dare* means, so this must be a rather simple information object, selected from a very narrow hypothesis space. Our candidate information object is a graph with a few links to nodes already in the system. Here we laid out a plausible path toward a learnable theory, but there is a clear esthetics-driven choice here: one must decide whether to search for *truth*, in the narrow sense of catering to the needs of knowledge engineers, or try to build explanatory models that can learn *dare*.

References

- Abney, Steven. 1991. Parsing by chunks. In R. Berwick, S. Abney, and C. Tenny, eds., *Principle-based parsing*, pages 257–278. Kluwer Academic Publishers.
- Angluin, Dana. 1987. Learning regular sets from queries and counterexamples 75:87–106.
- Angluin, Dana and Leslie Valiant. 1979. Fast probabilistic algorithms for Hamiltonian circuits and matchings 18:155–193.
- Banarescu, Laura, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and

⁷https://en.wikipedia.org/wiki/Generative_semantics

⁸A clever escape route may be left open by micro-timing the spreading, but this is a rabbit hole we do not want to go down.

⁹https://en.wikipedia.org/wiki/Boolean_satisfiability_problem

- Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186. Association for Computational Linguistics.
- Bierwisch, Manfred. 1988. *Thematic grids as the interface between syntax and semantics*. Lectures.
- Blackburn, Patrick and Johan Bos. 2005. *Representation and Inference for Natural Language. A First Course in Computational Semantics*. CSLI.
- Butt, Miriam. 2006. *Theories of Case*. Cambridge University Press.
- Carroll, John A. 1983. *An island parsing interpreter for the full augmented transition network formalism*. ACL Proceedings, First European Conference.
- Curry, Haskell B. and Robert Feys. 1958. *Combinatory Logic I*. North-Holland.
- Eilenberg, Samuel. 1974. *Automata, Languages, and Machines*, vol. A. Academic Press.
- Frege, Gottlob. 1892. On sense and reference. In A. Martinich, ed., *The Philosophy of Language*, pages 36–56. Oxford University Press (4th ed, 2000).
- Gurevich, Yuri. 1988. On Kolmogorov Machines and related issues 35:71–82.
- Jackendoff, Ray S. 1990. *Semantic Structures*. MIT Press.
- Karttunen, Lauri. 1971. Implicative verbs 47(2):340–358.
- Karttunen, Lauri. 2014. Three ways of not being lucky.
- Kiparsky, Paul. 1987. *Morphosyntax*. ms.
- Kolmogorov, Andrei N. 1953. O ponyatii algoritma 8(4):175–176.
- Kornai, András. 1987. Finite state semantics. In U. Klenk, P. Scherber, and M. Thaller, eds., *Computerlinguistik und philologische Datenverarbeitung*, pages 59–70. Georg Olms.
- Kornai, András. 2010a. The algebra of lexical semantics. In C. Ebert, G. Jäger, and J. Michaelis, eds., *Proceedings of the 11th Mathematics of Language Workshop*, LNAI 6149, pages 174–199. Springer.
- Kornai, András. 2010b. The treatment of ordinary quantification in English proper 54(4):150–162.
- Kornai, András. 2012. Eliminating ditransitives. In P. de Groote and M.-J. Nederhof, eds., *Revised and Selected Papers from the 15th and 16th Formal Grammar Conferences*, LNCS 7395, pages 243–261. Springer.
- Kornai, András. in press. *Semantics*. Springer Verlag.
- Kornai, András, Judit Ács, Márton Makrai, Dávid Márk Nemeskey, Katalin Pajkossy, and Gábor Recski. 2015. Competence in lexical semantics. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics (*SEM 2015)*, pages 165–175. Association for Computational Linguistics.

- Kornai, András and Marcus Kracht. 2015. Lexical semantics and model theory: Together at last? In *Proceedings of the 14th Meeting on the Mathematics of Language (MoL 14)*, pages 51–61. Association for Computational Linguistics.
- Marsh, William and Barbara Partee. 1984. How non-context-free is variable binding? In M. Cobler, S. MacKaye, and M. Wescoat, eds., *Proceedings of the West Coast Conference on Formal Linguistics III*, pages 179–190.
- McCarthy, John J. 2005. Human-level ai is harder than it seemed in 1955.
- Miller, George A. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information 63:81–97.
- Montague, Richard. 1973. The proper treatment of quantification in ordinary English. In R. Thomason, ed., *Formal Philosophy*, pages 247–270. Yale University Press.
- Pinker, Steven and Alan S. Prince. 1988. On language and connectionism: Analysis of a parallel distributed processing model of language acquisition 28:73–194.
- Quillian, M. Ross. 1967. Semantic memory. In Minsky, ed., *Semantic information processing*, pages 227–270. MIT Press.
- Schönhage, A. 1980. Storage modification machine 9(3):490–508.
- Shvachko, Konstantin V. 1991. Different modifications of pointer machines and their computational power. In *Proc. 16th Mathematical Foundations of Computer Science*, pages 426–441. Springer.
- Talmy, L. 1988. Force dynamics in language and cognition 12(1):49–100.
- Uspensky, Vladimir and Alexei Semenov. 1993. *Algorithms: Main Ideas and Applications*. Springer.
- Yngve, Victor H. 1961. The depth hypothesis. In R. Jakobson, ed., *Structure of Language and its Mathematical Aspects*, pages 130–138. American Mathematical Society.