ARTICLE IN PRESS

Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research xxx (2013) xxx–xxx

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Characterization of reasoning in terms of perceptual simulation

Hiroyuki Uchida *, Nicholas L. Cassimatis, J.R. Scally

Cognitive Science Department, Rensselaer Polytechnic Institute, United States

Received 4 February 2013; received in revised form 10 October 2013; accepted 28 October 2013

Abstract

Although characterizing reasoning and natural language semantics in traditional logic captures their complexity and productivity, accounting for the grounding of logical reasoning in perception raises several challenges. These include difficulties in explaining the integration of reasoning and perceptual processing, and in accounting for the evolution of human reasoning from sensorimotor origins. Central to these problems is the fact that traditional logic includes elements such as quantifiers and negation that do not obviously occur in perceptual representations. We propose a formal framework in terms of perceptual simulation that bridges this gap. We demonstrate that perceptual simulations have the power to explain crucial elements of logical human reasoning and also allow us to provide the first unified linguistic analysis of noun phrases, negative polarity items and branching quantifiers within a single cognitively motivated formal framework.

Keywords: Reasoning; Perceptual simulation; Quantifiers; NP semantics; NPI; Branching quantifiers

1. Introduction

Logical languages such as first-order predicate logic are commonly used to represent natural language semantics and to characterize human reasoning (cf. Gamut, 1991).[1] Such a logical approach reflects the compositionality and productivity of natural language semantics. Human reasoning also shares productivity and at least some deductive properties with logical inference.[2] On the other hand, a major problem with the use of logic for enabling natural language semantics and reasoning is the grounding of the semantic information in perception (cf. Harnad, 1990).

The problem is that if we assume that a first-order logic (FOL) or its equivalent enables human reasoning, then it means that the ingredients of reasoning involve abstract elements such as quantifiers, the negation operator '¬' and the disjunction connective '∨'. In contrast, it is not immediately obvious how such abstract elements could be included in perceptual representations such as visual images.[3] This makes it harder to account for how reasoning and perception integrate. As a result, cognitive theories

* Corresponding author. Address: Cognitive Science Department, Rensselaer Polytechnic Institute, Carnegie 108, 110, 8th Street, Troy, NY 12180, United States. Tel.: +1 (518) 423 4755; fax: +1 (518) 276 8268.
E-mail address: uchidh@rpi.edu (H. Uchida).

[1] We assume that humans reason with natural language semantic representations. Although this is not an absolute necessity for our account, most semantic theories (Carpenter, 1997; Gamut, 1991; Heim & Kratzer 1998) assume this. Also, since a crucial motivation for the formal language that we propose is the efficient data exchange between different computational methods that are used for formulating human-level intelligence, it makes little sense to multiply the number of knowledge representations discussed in this paper.

[2] Some researchers (Oaksford & Chater, 2006; Spivey, 2007) argue that human reasoning is not logical. However, even if human reasoning is not exactly like classical logic inference, a theory of reasoning somehow needs to explain those aspects of human reasoning that have motivated logical formalisms.

[3] In fact, it is not easy to represent quantifiers and propositional negations in terms of visual images. See Uchida, Cassimatis, and Scally (2012) for details.

1389-0417/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cogsys.2013.11.001

become fragmented and significant inefficiencies or inadequacies are introduced into computational systems attempting to combine reasoning and perception.

Admittedly, it is debatable whether human reasoning and perception should be integrated in the theory of human-level intelligence. However, as Barsalou (1999) pointed out and as our own earlier work (Uchida et al., 2012) also discussed, the view that reasoning occurs using representations similar to perceptual representations has several theoretical merits. It helps explain how cognition and perception are connected, how child cognition develops and how human cognition evolved from the cognition of animals with primarily perceptual and motor abilities. Further, an increasing amount of psychological and neurological data (Barsalou, Pecher, Zeelenberg, Simmons, & Hamann, 2005) is consistent with this theory.

In natural language interpretation, perceptual information such as visual images can continually interact with the semantics of natural language expressions.[4] Similarly, perceptual information can be integrated into general reasoning at any time (cf. Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995). For example, when a car driver perceives another car running too close to his car, he may drive away from that car, according to a general reasoning rule such as, 'If one's car is too close to a large object, one drives ones' car away from that object'. We may represent this rule by a first-order formula in (1).[5]

(1) $\forall x \forall y \forall z((\mathrm{Car}(x) \& \mathrm{Drive}(y,x) \& \mathrm{BigObject}(z) \& x \neq z \& \mathrm{Near}(x,z)) \rightarrow \mathrm{DriveAwayFrom}(y,x,z))$

However, in the above example, the perceptual system recognizes a particular car, say, John's car that John is driving, getting too close to another particular car, i.e., the car that John is seeing next to his car. Thus, if we account for human cognition by using first-order logic as in (1), there will be non-trivial inference steps in order to match the concrete perceived information with the rule stated by the abstract logical form in (1) so that John can conclude, 'John drives John's car away from the car running next to his car.'

To address these problems, we present an account of human reasoning in terms of simulations with perceptual mechanisms and propose a language that represents such simulations. Since the ingredients of this language have corresponding elements in perceptual representations such as visual images (see Section 2), our theory will be able to account for the integration between reasoning and perception as above, or the grounding of linguistic information in perception, in a more efficient manner.

In this paper, we compare our simulation language with first-order logic (FOL) before considering higher-order logics in future work. This is partly because, although the analysis of reasoning with first-order logic has the above mentioned grounding problem with regard to perception, an analysis using higher-order logic (cf. Gamut, 1991) is even worse in this regard.[6] Similarly, most established automated implementations of reasoning use first-order logic (cf. Fitting, 1996), which indicates that a theory of reasoning using higher-order logic still has a non-trivial hurdle to clear, since no proper scientific theory should require the supervision of the theorist to deal with new data. On the other hand, an analysis of natural language semantics and reasoning using higher-order logic (cf. Barwise & Cooper, 1981; Carpenter, 1997; Gamut, 1991) has several merits, such as the compositionality of interpretation and its ability to represent finer-grained entailment relations. In this regard, we hope that our future research can show that perceptual simulations can capture these merits of higher-order logic without sharing its demerits.[7]

As mentioned above, we characterize human reasoning in terms of simulations with perceptual mechanisms. Several cognitive scientists (Barsalou, 2009; Goldman, 2002; Gordon, 1995; Gordon & Cruz, 2002) also formulate human reasoning in terms of simulations. In addition to the efficient integration of perceptual information into reasoning, simulation theory naturally explains how the ability of reasoning (and the use of natural language associated with it) has evolved from the perceptual abilities shared by humans and their primate ancestors (Barsalou, 1999; Cassimatis, Murugesan, & Bignoli, 2009a). A theory that formulates reasoning and natural language interpretation in terms of a

[4] Humans often omit language expressions whose meanings are perceptually recoverable in the context, exemplified with fragmental utterances such as "Look" and "Which book?" – "The red one". There has been discussion whether such utterances are elliptical (Merchant, 2004) or truly sub-sentential (Stainton, 2004). On either view it is uncontroversial that inference is involved in fleshing out what was meant.

[5] For readability, we simplify logical formulas by omitting the details such as tense, modality and location. We use intuitive English expressions for predicates, such as 'BigObject' and 'DriveAwayFrom', ignoring the internal structures of the semantics of the corresponding English expressions for convenience.

[6] One can conclude this from the fact that the interpretation models for higher-order logics include functions that map sets of individuals to functions from sets of individuals to truth values (e.g., denotations of *every* and *some*), etc., whereas the interpretation models for first-order logic include only concrete individuals and the first-order sets that have either those individuals or ordered pairs of those individuals as members. In this regard, notice that each first-order interpretation model does not explicitly represent abstract notions such as quantifiers and negations.

[7] Sorted first-order logic can capture many elements of higher-order logic (cf. Fox & Lappin, 2004; Gamut, 1991). Thus, in order to formulate higher-order reasoning in our framework, we can aim to characterize the corresponding sorted first-order reasoning in terms of simulations. The so-called 'proportional' quantifiers, such as *most linguists*, are commonly taken to exceed the power of first-order logic. In this regard, we can introduce probabilistic elements to simulations, or posit certain internal structures for individual objects (Uchida & Cassimatis, 2010), extending the plural-individual structures in Link (1983) with our characterization of quantification in terms of simulations. We leave further details for another paper.

symbolic system that involves abstract elements (such as quantifiers) that are not grounded in perception would be less explanatory in this regard.

Although our research shares the basic motivations with other works on the simulation theories of human cognition that we mentioned above, this paper is less concerned about the psychological mechanisms of simulations. Instead, our focus is to show that simulations with only perceptual mechanisms, and therefore, without explicit representations of abstract elements such as quantifiers and the negative propositional operator, can nonetheless have the expressive power to capture the notions of quantification and negation. Note that we do not aim to show that concrete simulations as above do underlie all of human reasoning, only that they have the expressive power to deal with the crucial abstract elements in human reasoning. Also, this paper is not concerned about the exact computational implementation of the simulation language.

After explaining the basics of our simulation-based account of FOL inferences, we show that our characterization of reasoning in terms of simulations can straightforwardly account for several linguistic phenomena. First, since we formulate the semantics of both referential and quantificational noun phrases (NP) in terms of individual objects that appear in simulations, we can provide a uniform semantic analysis of the two types of NP. In contrast, as we discuss in Section 3 (see also Gamut, 1991), the semantic analysis of referential and quantificational noun phrases using FOL is problematic because of its lack of uniformity and difficulty to maintain semantic compositionality.[8]

Secondly, our simulation language can provide a uniform analysis of the distribution of negative polarity items (NPI), such as *any* in *Bob did not read any book*. Although a traditional logical analysis, such as Ladusaw (1979), can explain the distribution of NPI with the notion of downward entailing (DE) environments, there are some NPI data (cf. Kadmon & Landman, 1993) that are not easy to analyze by way of the DE environment. This logical analysis also fails to capture the commonality between the NPI *any* as above and the so-called free choice *any*, as in *Bob will read ANY book*, indicating a lexical ambiguity analysis of the linguistic form *any*, which is both theoretically and empirically undesirable (Horn, 1972; Kamp, 1973; Carlson, 1981; Kadmon & Landman, 1993). In contrast, we account for the NPI *any* and the free-choice *any* in fundamentally the same way, and thus, our proposal is at least compatible with the uniform lexical analysis of the NPI *any* and the free-choice *any*.

[8] Semantic theories using higher-order logic (cf. Barwise & Cooper, 1981; Carpenter, 1997) can provide a uniform and compositional account of NP semantics. As we indicated above, however, higher-order logic is not suitable for our main objective, i.e., maintaining the correspondence between reasoning and perception.

Finally, we show that our simulation language can provide a cognitively well-grounded account of branching quantifier propositions (Barwise, 1979; Henkin, 1959). Since branching quantifiers involve a non-linear scope order between quantifiers, it is difficult to deal with in a semantic analysis based on FOL, which linearizes all the quantifiers that appear in a proposition.

Crucially, we can deal with the above linguistic data without introducing any special mechanisms other than those that perceptual simulations already have. In fact, as we discuss in Section 3, the above three phenomena are expected from the way we continually introduce objects into perceptual simulations.

Section 2 explains how simulations with only concrete, perceptual elements can still reflect abstract notions such as quantification and negation by way of implication. Section 3 provides independent linguistic motivations for using our simulation-based semantic representations, as we sketched above. Section 4 provides concluding remarks.

2. Reasoning with perceptual simulations

This section explains our formulation of reasoning in terms of simulations with perceptual elements. Section 2.1 provides the basics and Section 2.2 shows our treatment of more complex natural language expressions whose semantics are often taken to motivate the abstract operators and connectives in standard first-order logic.

2.1. Basics

As we explained in Section 1, an important motivation for accounting for human reasoning in terms of perceptual simulations is to ground reasoning in perception. Thus, we start our exposition of simulations with the ingredients of reasoning that have clearer reflections in perception. First, we argue that perceptual representations include a certain degree of internal structure. More specifically, we assume that perception can individuate objects and recognize certain kinds of properties and relations on objects. As support for this assumption, from the earliest ages (Wilcox, 1999) children are able to individuate objects from each other and from the background. Similarly, primates and mice can recognize individual objects and attribute certain perceptual properties and relations to those objects, such as color, smell and spatial relationship. Moreover, Hesslow (2002) discusses some neurological data that suggest that simulating perceptual experiences, possibly with hypothetical objects, is essentially the same as perceiving the analogous objects in reality, in terms of how the relevant areas in the brain are activated.

These suggest that including individual objects (which may be ‘hypothetical’, or have never been perceived before in reality) and first-order predicates on those objects as ingredients of reasoning does not go against the correspondence between reasoning and perception in terms of their basic expressive powers. Consider (2).

(2) a. Reasoning: b. Vision:

{Circle(c), Triangle(t), LeftOf(c,t)}

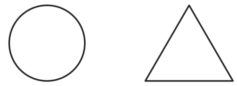

In (2a), which is meant to represent some possible ingredients of reasoning in an intuitive manner, the atomic formula 'Circle(c)' means that the property 'Circle' is attributed to the individuated object 'c.' For convenience, we use lower-case letters such as 'c' and 't' to identify the objects that appear in reasoning. In (2a), we have used the notations of first-order logic (FOL) for convenience, but this does not mean that our reasoning with simulation uses the rules of FOL. The notation in (2b), which is meant to capture some visual representations, is also chosen for convenience. That is, exactly how the visual mechanism encodes this imagery is not important in this paper. The essential point in (2b) is that vision recognizes two distinct objects, one of which is recognized to have the shape of a circle and the other is recognized to have the triangular shape. We assume that vision recognizes the spatial relation between these two objects that corresponds to the predicate 'LeftOf' in reasoning. Clearly, there is an isomorphism between (2a) and (2b) in terms of the individual objects and their properties and relations that we have just explained.

We are not arguing that all the properties and relations that humans use in reasoning are represented in each perceptual mechanism.[9] First of all, some properties or relations may be represented only in vision but not in the auditory system and vice versa. Also, exactly which visually relevant properties and relations are actually recognized by the human's visual mechanism and to what degree of granularity they are recognized would require a careful psychological investigation and thus, those are beyond the scope of this paper. Secondly, it is clear that humans can reason with non-perceptual concepts such as kin-ship relations and abstract mathematical concepts. We do not investigate how we can accommodate these abstract properties and relations in our formulation of reasoning in this paper, whose main goal is to maintain the basic correspondence between reasoning and perception. However, our general argument is that as long as the perceptual faculties are equipped with the basic structures to represent the attribution of at least some properties and relations to (possibly hypothetical) individual objects, using the same kind of structures in reasoning to attribute more abstract properties and relations to individuals (or hypothetically introducing new individuals and attributing properties to them) does not go against the fundamental correspondence between reasoning and perception.[10]

The commas inside the set in (2a) are interpreted as 'AND,' as standard in the set theory. The correspondence between (2a) and (2b) shows that the conjunction of multiple atomic predicate-argument statements in reasoning can be represented by one visual image that contains the elements that correspond to these statements.

(2) has shown that introducing the predicate-argument structures into reasoning does not go against the correspondence between reasoning and perception in terms of their basic expressive powers. In contrast, consider the semantics of a quantified first-order formula in (3b), which is often used to represent the semantics of the English sentence in (3a).

(3) a. Every boy runs.
  b. $\forall x(\mathrm{Boy}(x) \rightarrow \mathrm{Run}(x))$

At least at first sight, it is not clear how we can represent the information contained in (3b) in terms of a visual image such as the one in (2b).[11] A similar difficulty arises with negated propositions. For example, how can we represent the semantics of 'c is not a circle', which is commonly represented by a first-order formula, '¬Circle(c),' in terms of a visual image? Putting a triangle, for example, in the image is not satisfactory, since although we can tell the shape is not a circle, we can also tell the shape is not a square, or any other shape that is not a triangle.

Does the existence of quantifiers such as *every* and *some* and the negative operator *not* in natural language suggest that although maintaining the tight correspondence between reasoning and perception as above has strong merits, that position is not empirically sustainable? We argue that it is still possible to maintain the basic correspondence between reasoning and perception.

We deal with this challenge by characterizing reasoning in terms of simulations with perceptual elements. As we explained in Section 1, the simulation theory has strong motivations in cognitive science and some researchers have considered a symbolic reasoning mechanism based on this theory (cf. Barsalou, 1999). In this paper, our main focus is to show with linguistic examples that formalizing reasoning as simulations with perceptual representations allows us to maintain basic correspondence between the representations used in reasoning and perception.

[9] For arguments that human perception can recognize a limited kind of properties and relations, see Lakoff (1987) and Lakoff and Núñez (2001).

[10] Also, note that humans have difficulty grasping abstract notions, such as the vector spaces of n-dimensions with $n \geqslant 4$, which are harder to represent in perceptual imagery.

[11] The standard first-order interpretation models, which are made out of the domain of individuals and sets of those individuals and of the ordered n-tuples of those individuals, cannot directly represent the notion of quantification, either. Since the first-order interpretation models are formally well-behaved (cf. Hodges, 1993) while still representing the states of affairs in the world in a way that is adequate for natural language semantics, we argue that our characterization of reasoning in terms of perceptual simulations has the merit of including only those elements that have their denotations in first-order interpretation models.

Please cite this article in press as: Uchida, H., et al. Characterization of reasoning in terms of perceptual simulation. *Cognitive Systems Research* (2013), http://dx.doi.org/10.1016/j.cogsys.2013.11.001

Intuitively, perceptual simulations proceed just like simulations that physical scientists conduct by setting up idealized experimental situations in the physical world. For example, in order to prove the existence of gravity operating on objects with mass within a certain distance from the surface of the earth, physical researchers might actually toss objects with some mass into the air without any hindrance. The result of such a simulation, such as the falling of those objects, may support their physical theory. A perceptual simulation proceeds in a similar manner, only by way of perceptual representations, instead of being conducted in the physical world. For example, in order to simulate a coin-toss, one may mentally represent a coin tossed into the air in one's mind and this simulation may produce the result that the tossed coin falls to the ground in one's mind. For concreteness, we can say that a perceptual simulation takes in a concrete perceptual state of affairs as in (2) as input and produces a similar perceptual state of affairs as output. Remember that a perceptual representation such as (2b) is isomorphic to a set of concrete, atomic formulas such as the atomic formulas in (2a). Because of this, we can schematically represent a simulation process as in (4).

(4) A process of perceptual simulation:

{Coin(c), TossedIntoTheAir(c)} ⇒ [Simulation 1] ⇒ {Fall(c)}

Technically, a simulation takes in a set of atomic propositions,[12] $\{\Phi_1,\ldots,\Phi_n\}$, as input and produces a set of atomic propositions, $\{\Psi_1,\ldots,\Psi_k\}$, as output. Again, as we saw at (2) above, each proposition, Φ_i or Ψ_j, is a perceptually representable statement, maintaining the correspondence between general reasoning and perception at the level of the elements that are explicitly represented in simulations and the elements that are recognized in perception. Since the exact computational implementation of reasoning that we propose is beyond the scope of this paper, we are not concerned about exactly what algorithm allows us to produce particular outputs from the given inputs.

Some might wonder what independent motivations we have for including two stages of perceptual representations in the simulation process, namely, the input state of affairs and the output state of affairs. One justification for this is the so-called pattern-completion task conducted in human vision (cf. Barsalou, 2003). For example, humans can see a (mentally) complete visual image even though the actual physical image is made out of dots placed in a certain configuration, as in news-paper pictures.[13] More generally, the human's visual mechanism can take in an occluded image as input and somehow complete the missing elements. We take it as evidence that the perceptual mechanism can link one perceptual state of affairs to another by way of some sort of casual relation.[14] As we discussed above with regard to the perceptual mechanisms' recognition of predicate-argument structures, what is important is that the human's perceptual mechanisms are equipped with the basic ability to connect two perceptual states of affairs with the recognition of some sort of causal relation from the initial state to the following state. Once the perceptual mechanisms are equipped with that linking ability, making an enriched use of the same ability only in reasoning does not go against the fundamental correspondence between reasoning and perception, as we see with examples below.

Assuming that these assumptions about simulations are reasonably convincing, we explain how we can maintain the correspondence between the ingredients of reasoning and the perceptual representations in the presence of the notion of quantification in human reasoning. First, as a matter of notation, we henceforth represent (4) above as in (5).

(5) Simulation 1, based on: {Coin(c), TossedIntoTheAir(c)}
produces: {Fall(c)}

We call the input to the simulation, that is, the set above the horizontal bar in (5), the basis of the simulation and the output, the conclusion of the simulation. As we explained above, 'c' is used as the identifier of the object that appears in this simulation. We call it a term.

We start with our treatment of a universally quantified proposition, e.g., 'Every planet rotates.' Imagine that we tell somebody, say, Meg, to do a mental simulation with an object that she only recognizes as a planet. This object is new to Meg and she knows nothing else about this object, not even its name. Now assume that Meg runs a simulation as instructed and then it produces the result that this object rotates. We can symbolically represent this simulation as in (6).

(6) Simulation 2, based on: {New(p), Planet(p)}/
produces: {Rotate(p)}

The formula 'New(p)' in (6) is a little different from other formulas. While normal formulas indicate that the corresponding states of affairs hold in (Meg's conception of) the external world, the formula 'New(p)' indicates where in Meg's reasoning the object 'p' is introduced for the first time. In (6), it indicates that 'p' is introduced newly in the assumption of this simulation. We could omit this

[12] By 'atomic propositions', we mean propositions that can be represented by simple predicate-argument structures, $\text{Pred}(a_1,\ldots,a_n)$, without propositional connectives or operators.

[13] We think that the presence of perceptual illusions (Luckiesh, 1995) also indicates that perception can connect two states of affairs by way of a causal link.

[14] Also, see Garbarini and Adenzato (2004) for a mechanism of 'as-if' neural simulation, which involves the 'assumption-conclusion' structure that is analogous to the 'input-output' structure in (4). See also Hesslow (2002) for some neurological evidence that humans can simulate perceptual experiences in a hypothetical setting that may involve an assumption-conclusion structure.

'syntactic-sugar' formula while keeping track of where in Meg's reasoning each object is introduced in some other way, but for presentational convenience, we keep it in the symbolic representation.

Our complete theory of cognition assumes that mental simulations may exhibit varying degrees of uncertainty. The main reason for this assumption is that human cognition can leave some room for uncertainty, especially when general reasoning is concerned. Having said that, in this paper we only discuss simulations that produce their outputs with full certainty.

We also show only 'factual' simulations in this paper, that is, those simulations that are consistent with the knowledge of the actual world in the mind of the person who conducts the simulations (who we call the **agent** henceforth). More specifically, everything that the agent knows is true in the real world is also true in each factual simulation that the agent conducts. For example, in each of Meg's factual simulations, planets have exactly those properties that she knows they have in the real world. We leave the discussion of 'counterfactual' simulations that involve something that contradicts what the agent knows or believes is the case in the real world for a different paper.

Now, suppose that Simulation 2 in (6) is fully certain and factual. This means that Meg concludes that the state of affairs represented in the conclusion necessarily follows from the state of affairs represented in the basis and that this simulation is consistent with Meg's knowledge of the states of affairs in the real world. All the ingredients that appear in Simulation 2 are concrete, that is, the object 'p' corresponds to the object that Meg recognizes as a planet, and all the properties are properties of such individual objects. Now, consider what we can tell from Simulation 2. Again, Meg initially recognizes the object 'p' only as a planet. If Meg can conclude with full certainty that this object rotates, then that output of the simulation must hinge only on 'p' being a planet, together with the characteristics of planets in Meg's knowledge of the real world. In other words, from the fact that Meg's simulation in (6) has produced the output '{Rotate(p)}' with full certainty on the sole basis of the object 'p' being a planet as in the real world and nothing else, it can be inferred that Meg believes that any object that is a planet will produce the same result, as long as planets are as they are in the real world. That is, Meg believes that every planet rotates (in the real world).

Since this universally-quantified statement can be read off Meg's simulation result in (6) above, we can take Simulation 2 as an unambiguous indicator of the truth of this universally-quantified proposition. In other words, in our semantic theory, we can use the symbolic representation of Simulation 2 in (6) to represent the proposition, 'Every planet rotates', instead of the first-order formula '$\forall x(\text{Planet}(x) \rightarrow \text{Rotate}(x))$'.

We move onto our treatment of existential quantifiers. Suppose that Meg simulates with a new individual, 'b,' which she only knows as a boy. Suppose that Meg's simulation concludes with certainty that 'b' likes another new individual, 'g,' which Meg recognizes only as a girl. Again, this is a factual simulation and hence, boys and girls are as they are in the real world. We can represent this simulation

(7) Simulation 3, based on: {New(b), Boy(b)}/
produces: {New(g), Girl(g), Like(b,g)}

as in (7).

The individual object 'b' is assumed to be new to Meg and by way of the reasoning as above, an object that is newly introduced in the basis of a simulation implies universal quantification. That is, from the result of Simulation 3, which we assume that Meg confirms with full certainty, we can deterministically recover the implication that any individual that can replace 'b' in the basis (i.e., any individual that is a boy) produces the same result in the output, as long as things are as they are in the real world. Unlike Simulation 2, however, the output of Simulation 3 includes another object, that is, 'g.' As we suggested above, we assume that this object 'g' is new to Meg as well. From Meg's viewpoint, we can say that Meg has introduced a new object in the output of Simulation 3. Meg has introduced this object 'g' in response to the object 'b' in the basis. As a result, when we read the universal implication as above with regard to 'b', and replace this object with another boy, say, 'b′', then g in the output is replaced with a possibly different object, say, 'g′', corresponding to 'b′'. Descriptively, the implication that we can read off Simulation 3 is that for each object that is a boy, it introduces a (possibly) different object, which is a girl that the boy likes. That is, 'Every boy likes a (different) girl.' Again, as we can deterministically recover this quantified proposition from Simulation 3, we let the symbolic representation of Simulation 3 represent this proposition in the formal implementation. In other words, our simulation language uses Simulation 3, instead of '$\forall x(\text{Boy}(x) \rightarrow \exists y(\text{Girl}(y)\&\text{Like}(x,y)))$,' to represent the indefinite narrow-scope interpretation of 'Every boy likes a girl.'

Note that the introduction of a new object in the output of a simulation has the existential implication on condition that the basis holds. Introducing a new object in Meg's mind normally implies its existence in Meg's mind. However, positing an object in the basis of a simulation has a hypothetical status, in the same way that positing an object in the premise of a conditional statement has a hypothetical status, as we can see in the statement, "If we posit a new boy, 'b', in the basis of a simulation, that simulation produces an output which includes another new individual, 'g', who is a girl that 'b' likes." As we can see in this statement, the existence of 'g' is not entailed by Simulation 3 either, since its existence depends on the existence of 'b' in the basis. However, with Simulation 3, each time the existence of some boy is confirmed, for example, by a perceptual mechanism that has recognized a boy in the real world, the reasoning mechanism can conclude that there exists a girl for that boy. In this explanation, the new object 'g', which is introduced in the output of Simulation 3, does not imply universal quantification. This is because the conclusion of a simulation on its own does not trigger any further simulation based on that. Simulation 3 implies that any individual that can replace 'b' will produce a girl 'g' which 'b' likes in the output, but replacing 'g' with another girl individual that 'b' likes does not lead to any further output.

In the formal implementation of simulations, the computational system will compute the universal replacement possibility as above with regard to any new object that appears in the basis of a simulation. Similarly, for any new object that appears in the conclusion of a simulation, the existential implication is calculated on condition that the basis assumption holds. In this paper, we are proposing this reasoning mechanism as the underlying theory behind the cognitive ability of reasoning. In other words, we are not arguing that every time a person performs a quantificational reasoning, she literally goes through the simulation process as above. On the other hand, the above explanation assumes that humans are capable of simulating in the above manner as a matter of ability and it follows from there that humans' reasoning can involve some objects that have not been perceived by a perceptual mechanism. For example, for Meg to run Simulation 2 as above successfully, the new object does not need to appear in a perceptual representation. In fact, if 'p' corresponded to some object that Meg's perception has already recognized in the real world, we could not read off the universal implication as above, and hence Simulation 2 would have a different implication, as we discuss later.

We do not think that this reasoning mechanism's ability to posit an object that has not yet been recognized by a perceptual system poses a fatal problem for our theory of reasoning. Our main goal is to achieve the basic correspondence between reasoning and perception in terms of the foundational representational expressive powers of these mechanisms, not in terms of which (representable) information these systems actually represent at the level of performance. Thus, we do not require each item that appears in reasoning in performance to appear in a perception mechanism. Such a requirement is cognitively undesirable. For example, the general reasoning mechanism may store some universally applicable rule in the form of Simulation 3 above. Suppose that a perceptual mechanism such as vision recognizes some boy, say, Bob. This implies that there exists a girl that Bob likes. Given this computation in the reasoning mechanism, the perceptual mechanism might look for an actual girl that confirms this reasoning result. If the agent fails to perceive such confirmation for a long time, the agent may become uncertain about Simulation 3. Note that in this intuitively plausible spontaneous interaction between the reasoning mechanism and the visual mechanism, it is not the case that these two mechanisms represent exactly the same individuals at every instant.[15]

We have dealt with the narrow scope reading of *Every boy likes a girl*, but this sentence also has the wide-scope reading of the indefinite, 'There is a girl that every boy likes.' How can we represent this proposition in our simulation language? Consider (8).

(8) Simulation 4, based on: {}/produces: {New(g), Girl(g)}
Simulation 5, based on: {New(b), Boy(b)}/produces: {Like(b,g)}

In (8), we use the same terms, 'b', 'g', that we used in (7) but this is only for convenience and we assume these objects are newly introduced in (8). First, although simulations normally have a non-empty basis, we assume that humans can also simulate without any hypothetical assumptions. This does not mean that simulations can produce something out of absolute nothing. As we have indicated above, the factual knowledge that the person who runs the simulation has about the real world is assumed to be compatible with these simulations. Now, (8) contains two simulations. Consider Simulation 4 first, which we assume takes precedence relative to Simulation 5. As in the explanation above, the empty basis of Simulation 4 means that the person who runs this simulation produces this output without assumptions, given her general knowledge about the states of affairs in the real world. The individual 'g' is newly introduced in the output of this simulation and as we explained above, this implies existential quantification. However, since there is nothing in the basis of Simulation 4, 'g' does not depend on any other object, such as 'b' in Simulation 5. Turning our eyes to Simulation 5, the newly introduced boy individual, 'b', leads to the above mentioned universal implication. On the other hand, the girl individual 'g' in the output of Simulation 5 is not new. Given the above mentioned precedence relation between Simulation 4 and Simulation 5, 'g' in Simulation 5 refers back to 'g' in the output of Simulation 4. Thus, 'g' does not covary with 'b', nor with any other individual. In this way, from (8), we can deterministically recover the proposition, 'There is a girl that every boy likes,' and in the symbolic system that we develop, we let (8) represent this proposition, instead of the first-order formula, $\exists y \forall x(\text{Boy}(x) \rightarrow (\text{Girl}(y)\ \&\ \text{Like}(x,y)))$.

Next, we show our treatment of referential expressions and negation, both of which appear in (9).

(9) Jack does not like Meg.

As we indicated above, our theory of reasoning has developed out of the usefulness of perceptual simulations. However, we are not arguing that the reasoning mecha-

[15] We also assume that the objects that appear in the perceptual representations might not exist in the external world. The existence of perceptual illusions suggests that the perceptual mechanisms can represent non-existing objects.

nism operates only by way of simulations. Naturally, we assume that this mechanism is also equipped with the basic function of representing a factual state of affairs observed in the real world. For example, suppose that Meg above knows a particular individual, say, Aaron Rogers, and also knows that this individual plays football. Then, our reasoning mechanism can store this information as in (10).

(10) In Reality: {NameOf(ar, AaronRogers), Football(fb), Play(ar,fb)}

(10) means that Meg's knowledgebase contains the factual information about the world that the individual, 'ar', whose name is Aaron Rogers, is engaged in the activity, 'fb', that has the property of American football, or more concisely, 'Aaron Rogers plays American football.' Here, we treated the proper name, 'Aaron Rogers,' as a property of individual objects, based on several motivations, such as the presence of more than one individual in the world who has this name and also the presence of some language, such as, Italian, which adds determiners to proper names to form argument NPs, suggesting that proper names have the same status as common nouns such as *man*, *teacher*, etc., which are normally taken to denote properties. Also, we treated activities as objects and used the term 'fb', to which we attributed the property of American football. We believe that treating events and activities as objects have many empirical merits in natural language data analysis. However, these are not essential for the main merits of our theory of reasoning. The only essential point in (10) is that recognizing 'ar' and 'fb' in the real world has the effect of fixing the referents as the objects identified in the real world. Note that the basic property of the terms in (10) out of context is exactly the same as the terms in the other examples above, such as 's', 'b', and 'g,' that is, all these terms are identifiers of particular objects. However, including these objects in the knowledgebase of the confirmed states of the real world leads to a different implication from the objects that are introduced newly in simulations.

With our treatment of proper names as above, we now provide our treatment of 'Jack does not like Meg' in (9) above. Consider (11).

(11) In Reality, {NameOf(jk, Jack), NameOf(mg, Meg)}
Simulation 6, based on: {Like(jk,mg)}/produces: {⊥}

Again, the inclusion of the terms, 'jk' and 'mg', in the knowledgebase of the states of affairs in the real world has the effect of fixing the referents of these terms to the corresponding individuals observed in the real world. In Simulation 6, '⊥' means that this simulation fails. Intuitively, this means the following. Suppose that somebody, say, Meg, runs a simulation about the two individuals that she knows in the real world, say, Jack White and Meg White of the White Stripes, with the basis that Jack likes Meg. Suppose that Meg concludes that this simulation leads to contradiction (i.e., '⊥' in Simulation 6 indicates contradiction), given her knowledge about the states of affairs in the real world. This simulation result implies that the basis of this simulation is not compatible with what Meg believes about the states of the real world.[16] In this way, the above simulation result amounts to negating the basis of this simulation. That is, from this simulation result, we can infer Meg's belief: 'Jack White does not like Meg White.' Since we can deterministically infer this proposition from (11), again, in our theory of reasoning, we can let (11) represent the proposition, 'Jack White does not like Meg White.'

Before we end this subsection, we briefly discuss a 'more concrete' case of universal quantification. For discourse like "I went to the party yesterday. Everyone was drunk", our treatment of universal quantification at (6) above might seem less natural.[17] However, in our account, the universal quantification in the suggested example is treated in the same way we accounted for "Every planet rotates" at (6); the difference is how the domain restriction for the universal quantification applies. Intuitively, in (6) above, the term 'p' can potentially be any planet in the real world. In contrast, for the discourse example above, the term that is introduced newly in the basis of the corresponding simulation ranges only over the (human) individuals that the speaker met at the party that he went to yesterday (or further pragmatic domain restriction may apply depending on the context of the utterance). Although it might seem unnatural to posit a new object when it can only range over the individuals that the agent personally observed yesterday, recall that we are not assuming that the agent literally runs such a mental simulation with universal implication each time he or she envisages the corresponding universally quantified proposition. The point is that the simulation language has the basic expressive power to represent such a proposition.

That being said, the precise formulation of the domain restriction suggested above requires some further work, both in terms of its empirical adequacy and formal details. For example, we would need to consider the division of labor between the explicit domain restriction that is done in terms of the additional formulas put in the basis of the simulation in question and the implicit domain restriction that applies to the semantic models in which such a simulation language is interpreted. If we restrict the domain explicitly, then the formal language to achieve that would exceed the first-order expressive power. For instance, we might first introduce a term

[16] Recall that in this paper, we discuss only factual simulations. For our treatment of counterfactual simulations, see Uchida, Cassimatis, and Scally (2013).

[17] We thank an anonymous reviewer for suggesting this example.

that denotes the plural individual (cf. Link, 1983) that the speaker saw in the party yesterday[18] and then enable the language to express the restriction that a new term that we will put in the basis of the simulation for the above discourse example ranges only over the atomic components of that plural individual. Also, notice the importance of the past tense in the discourse example above. In fact, if we turn our example at (6) into the past tense, that is, *Every planet rotated*, the default interpretation seems to be "Every planet that the speaker observed rotated", which would require a domain restriction similar to the one for the above discourse example. We leave further details of our precise analysis for another paper.

As we stated before, we only discuss factual simulations in this paper in which, for example, planets are exactly as they are in the real world. Thus, as far as the proposition that we discussed at (6) is concerned, the expressive power of our formal language is basically the same as FOL.[19] However, the potential expressive power of the suggested language is stronger. For example, if we remove the restriction of factual simulations as above and assume that no facts in the real world need to hold in the simulation in question, it can represent a modal proposition such as "Every planet necessarily rotates". For further details of our treatment of modal reasoning in terms of simulations, see Uchida et al. (2013).

2.2. More complex examples

This subsection shows how our characterization of reasoning in terms of perceptual simulations can handle some other natural language expressions whose semantics are often taken to motivate the use of a standard logical language.

We start with a conditional sentence. Consider (12).

(12) If Avril sings, Jack is happy.

Our treatment of (12) is in (13).

(13) In Reality: {NameOf(av,Avril), NameOf(jk,Jack)}
Simulation 7, based on: {Sing(av)}/produces: {Happy(jk)}

Again, the terms 'av' and 'jk' in Simulation 7 refer back to the corresponding individuals in the real world, so neither of them leads to a quantificational implication.

(14) summarizes the basic simulation schema that we have explained so far, where $\Delta_1 \sim \Delta_3$ are sets of atomic (i.e., 'operator-free') propositions.

(14) In Reality, Δ_1
S_i, based on: Δ_2/produces: Δ_3

As we have explained above, what the agent regards as facts in the real world remain true in any factual simulation that he conducts. Thus, assuming that S_i is a factual simulation, (14) means that in the agent's characterization of reality, if the assumptions in Δ_2 are added to the facts in Δ_1, then it follows that Δ_3.

We have also assumed that the relation between the basis (i.e., assumption) and the conclusion of a simulation is the same as logical entailment relation, when the simulation is a fully certain one. Thus, (14) means that when Δ_1 holds, Δ_2 entails Δ_3. This characterizes the so-called monotonicity of classical logic, including FOL. This means that if FOL can infer a set of propositions Δ_j from a set of assumptions Δ_i (which can be empty), it follows that FOL can also infer Δ_j from Δ_i plus any other assumptions (cf. Gamut, 1991). For example, suppose that the agent concludes with full certainty that given the facts in reality, there exists an American who speaks Russian. In the schema in (14), this will be (15).

(15) In Reality, Δ_1
S_i, based on: {}/
produces: {New(a), American(a), Speak(a, Russian)}

Now, given the monotonicity discussed above, if the agent can conclude with full certainty without any assumption that there exists an American who speaks Russian (as far as the facts in Δ_1 are maintained), then from there, it follows that any simulation that contains any assumption maintains that conclusion, (again, as far as the facts in Δ_1 remain true in that simulation).

Our assumption that the relation between the basis and the conclusion of a simulation is the same as material implication (i.e., logical entailment) in classical propositional logic is partly due to our current purpose, that is, we are aiming to show that simulations can characterize the semantics

[18] Although the formal details of plural individuals are not important in this paper, each plural individual is made out of multiple atomic individuals. When we attribute a property to a plural individual, the property is attributed either in a distributive manner to each atomic component of the plural individual, e.g., for the sentence *Five students passed the final exam*, or collectively to the plural individual, e.g., for *Five students met up at Penn Station*, depending on the property in question and possibly on the context in which the sentence is uttered. For further details, see Link (1983). Unlike us, Link uses explicit quantifiers to quantify over the components of plural individuals.

[19] If we consider scope dependencies between multiple quantifies, the expressive power of our formal language is a little different from that of FOL. We discuss a relevant example later in Section 3.

[20] For example, by truthfully asserting "I will (definitely) meet you tomorrow", the speaker is not committing himself to the truth of a conditional statement; "If I die tonight, I will meet you tomorrow" (cf. Gauker, 2005). Allott and Uchida (2009) has shown that this and other typical counter-examples against the material implication account of natural language conditionals are misguided, but it is at least debatable if the semantics of the form $(P \rightarrow Q)$ adequately captures the semantics of the natural language construction; "If P, then Q" (since in classical logic, whenever Q is true, $(P \rightarrow Q)$ is also true).

[21] For example, we might assume that the relation between the basis and the conclusion of a simulation corresponds to some sort of causal relation, which can still be certain or uncertain to various degrees. See Cassimatis et al. (2009a).

of classical FOL. Thus, if material implication in classical logic turns out to be inadequate for human reasoning,[20] we may consider a different formulation of simulations.[21] However, the monotonicity of classical logic as above enhances its predictive power so that it can cover a potentially infinite number of inference patterns in a sound and complete manner (cf. Uchida, 2013). For example, although we assumed at (15) above that the conclusion of S_i there will also hold in any other factual simulation that the agent conducts with the same reality set, this does not mean that the agent needs to be aware of each conclusion proposition in every such factual simulation. In application, we will only show the relevant conclusions for each inference and maintaining the monotonicity of the underlying inference system makes it easier in application to omit the conclusions that are irrelevant for the simulation in question (since the derivation of the relevant conclusions that are explicitly shown is not affected irrespective of what other formulas are omitted either in the basis set or in the conclusion set). Also, in our account at (14), modifying what facts in Δ_1 remain true in the simulation in question (or, possibly, spontaneously modifying the agent's *characterization* of reality depending on the purpose of the simulation) will affect which conclusions are derived (see Uchida et al., 2013 for further details). Thus, for the purpose of this paper, we maintain the monotonicity of simulations as above.

Next, consider the disjunctive statement in (16).

(16) Avril sings or Jack sings.

We let (17) represent the proposition expressed with (16). We abbreviate 'Simulation' as 'S' henceforth.

(17) In Reality: {NameOf(av,Avril), NameOf(jk,Jack)}
S8, based on: {Sing(av)}/produces: {⊥}
S9, based on: {Sing(jk)}/produces: {⊥}
S10, based on: {Hold(S8), Hold(S9)}/produces: {⊥}

S8, which is based on the assumption that Avril sings, fails. This result implies the proposition, 'Avril does not sing,' as we explained above. Similarly, S9 represents the proposition, 'Jack does not sing.' S10 is based on the assumption that both S8 and S9 are as above, that is, both simulations output '⊥', meaning that they both fail. Descriptively, the basis of S10 assumes that both Avril does not sing and Jack does not sing. Now, S10 based on this assumption also produces '⊥', which means that this simulation fails, which implies that it is not the case that both Avril does not sing and Jack does not sing. By way of De Morgan's law (that is, 'Not(Not(Avril sings) and Not(Jack sings))' $\Longleftrightarrow$ 'Avril sings or Jack sings'), this is equivalent to the proposition, 'Avril sings or Jack sings'.

There are two things that could potentially be problematic with our treatment of disjunctions in (17). The apparent complexity of the representation is not a serious problem. It is well-known that humans have more difficulty processing disjunctive statements, say, in comparison to conjunctive statements. The first problem is the use of the atom 'Hold(S8)' and 'Hold(S9)' in the basis of S10. This means that the reasoning system can meta-represent the other two simulations, that is S8 and S9, and attribute to them some property, that is, the property denoted by 'Hold.' Since it is hard to argue that a perceptual mechanism, such as vision, can meta-represent some process that it goes through, such as the pattern-completion process that we mentioned above, and then attribute some property to it, our treatment of disjunction at (17) seems more problematic than any other that we have provided so far.

The second problem may actually turn out to correctly reflect human cognition. Our treatment of disjunctive statements in (17) means that human reasoning cannot directly represent the semantics of disjunction. We dealt with the disjunctive proposition only by way of an equivalent proposition that contains only a conjunction and several negations. At the moment, we are not sure if this captures human cognitive abilities correctly.

We do not address these two potential problems here but we leave a few comments. First, the above treatment of disjunction is provisional and there is a way of dealing with disjunction in a more direct manner, though at the moment, we prefer the above formulation since in this way we can completely cover all the inferences supported in standard first-order logic. Note that in standard propositional logic, whose properties standard first-order logic inherits, all the formulas with the connectives, '¬, &, ∨, ¬,' can be represented (by way of logical equivalence) with formulas with only '¬, →.' Since our simulation language represents the semantics of '→' in terms of the 'basis-conclusion' structure and can also implicate negation, the analysis along the line of (17) allows us to cover all the propositions representable in first-order logic.

Secondly, with regard to the meta-representational status of the labels 'S8' and 'S9' in (17), we might speculate that this ability of meta-representing some perceptual processes and attributing properties to them might provide the missing link between the high intelligence of modern humans and the lower cognitive abilities of our evolutionary ancestors. Obviously, the validity of this argument will depend on future researches in cognition.

At the moment, we assume that the proposed reasoning mechanism can provide identifiers for the simulations that it conducts and then attribute the property, 'Hold', to those individuated simulations. In contrast, we find no strong reason to assume that the perceptual mechanisms are equipped with this sort of reflective ability. However, we think that this divergence is less problematic than the divergence that would be introduced with the use of abstract operators such as quantifiers in first-order logic. After all, the perceptual mechanisms can individuate certain processes that take place naturally in the real world. Thus, providing identifiers for some operations (i.e., simulations) that can be taken as processes does not modify the basics too much.

We end this section by showing that we can deal with a basic scope ambiguity between a universal quantifier and a negation. Consider (18).

(18) a. Not every student is noisy. (i.e., $\neg\forall x(\text{Student}(x) \rightarrow \text{Noisy}(x)))$
b. No student is noisy. (i.e., $\forall x(\text{Student}(x) \rightarrow \neg\text{Noisy}(x)))$

We analyze (18a) with (19a) and (18b) with (19b).

(19) a. S11, based on: {Student(st), New(st)}/produces: {Noisy(st)}
S12, based on: {Hold(S11)}/produces: {⊥}
b. S13, based on: {Student(st′), Noisy(st′), New(st′)}/
produces: {⊥}

We assume that ‘st’ and ‘st′’ are introduced newly in S11 and in S13 respectively, which lead to the universal implication as we explained above. S11 implies, ‘Every student is noisy’. Since S12 based on the validity of S11 fails, S12 implies the negation of ‘Every student is noisy’, which is (18a). Thus, in our theory of reasoning, we let (19a) represent (18a). S12 meta-represents another simulation, that is, ‘S11,’ and assumes that it holds. To avoid such reflective element, we could alternatively postulate a complex basis, putting the whole of S11 as the basis of S12. But there are pros and cons between these two formulations and since we use the identifying terms for simulations in our treatment of disjunctions anyway, we maintain the formulation in (19a) in this paper. As for our treatment of (18b), S13 in (19b) corresponds to the first order formula, ‘$\neg\exists x(\text{Student}(x)\&\text{Noisy}(x))$’, which is logically equivalent to ‘$\forall x(\text{Student}(x) \rightarrow \neg\text{Noisy}(x))$’ in (18b).

Before we move onto the next section, we briefly discuss other formalisms of reasoning and natural language semantics to which our formalism is similar in some sense. First, our formalism is similar to Discourse Representation Theory (DRT), Kamp and Uwe (1993), and File Change Semantics (FCS), Heim (1983), in that the formal reasoning mechanism can continually introduce new objects and that the interpretations of these objects are influenced where in the formal semantic structures these objects are introduced. However, from a formal viewpoint, discourse referents in DRT and their equivalents in FCS are variables. This is essential since these theories use explicit quantifiers that bind these variables. It is true that DRT shows quantifiers in discourse representation structures mostly when they are dealing with complex sentences including plurals, as we discuss briefly below. However, some researchers (cf. Abusch, 1994) use explicit quantifiers to deal with the partial vs. exhaustive interpretations of pronouns bound by the singular indefinite in donkey anaphora sentences and standard DRT updates the value assignments to discourse-referent variables (cf. Kamp, van Genabith, & Uwe, 2005), which again indicates that discourse referents are essentially variables from a formal viewpoint.[22] In contrast, the terms in our simulation language are basically the same as constant terms in first-order logic.[23] This is because we read the abstract notion of quantification indirectly from the result of simulations (which in themselves include only concrete elements as we explained above). Whether it is better to explicitly represent all the contributing factors that validate inference depends on the main focus of the theory. But since our main purpose is to ground reasoning in perception, our formal representation omits the abstract information that we can read off the concrete simulations.

The above-mentioned difference between the object terms in our simulation language and the discourse-referent variables in DRT is partly due to the different main purposes. That is, the main reason why we have been developing our formal language is that it is useful for our hybrid human-level intelligence architecture (cf. Cassimatis et al., 2010), in which we can alternate between different computational methods depending on the purpose, such as switching between deductive inference using the simulation language and visual computation using graph structures (cf. Riesen & Bunke, 2012) or topological spaces (cf. Duan, 2003). Using explicit operators and variables in the deductive language poses difficulty for such a spontaneous switch between alternative representation structures. In contrast, DRT puts more emphasis on fine-grained descriptions of

[22] Discourse referents have been formulated as variables partly because DRT has developed as one of the theories of dynamic semantics (cf. Chierchia, 1995), which regards the semantics of natural language expressions in terms of their ‘context change potentials’ and formalizes such semantics in terms of their continual updating of value assignments to variables (cf. Peregrin & von Heusinger, 2004). We agree that DRT could have developed in a way that is closer to our simulation language, but we also speculate that that has not happened mainly because their original motivation and their main focus so far have been different from ours, as we come back to shortly.

[23] Technically, if we assume that the sequent: ‘…, *Planet*(*earth*) ⇒ *Rotate*(*earth*)’ is provable in standard classic logic, where ‘…’ represents some premises that are not explicitly shown, and if we also assume that the constant term ‘*earth*’ does not appear anywhere else in this sequent other than the two explicit occurrences above, then it follows that the hidden premises ‘…’ include either ‘$\forall x(Planet(x) \rightarrow Rotate(x))$’ or a formula that entails this universally quantified formula. This shows that, given our implicit approach to inference as we explained in the main text, the term ‘p’ that we used in (6) above can have basically the same property as the constant term ‘*earth*’ in the above logical sequent. Since our current formulation continually introduces new objects (i.e., new terms), in the most straightforward formal implementation of this account (cf. Cassimatis, Murugesan, & Bignoli, 2009b), we will assume *open world assumption*, OWA (Fagin, Kolaitis, Miller, & Popa, 2005), in which the uniqueness of the names (i.e., terms) is not necessarily maintained across different stages of reasoning (i.e., two distinct names may turn out to denote the same individual as the reasoning proceeds further), but there is also a way of avoiding OWA in our formulation of simulations if it turns out to be undesirable. Such technical details however are beyond the scope of this paper.

natural language semantic data. For that purpose, using variables that can be bound by various kinds of quantifiers and operators is convenient. For example, consider (20).

(20) a. *All* parents who have a single $child_1$ spoil him_1
b. *Most* parents who have a single $child_1$ spoil him_1
c. *Some* parents who have a single $child_1$ spoil him_1

In (20), how many of the single children are spoiled depends on which quantifier is in the main-clause subject NP of the sentence. In some other data (cf. Kratzer, 2005), such a quantificational element is contextually provided. This sort of data has motivated a semantic theory (cf. Heim, 1982; Kratzer, 2005) in which the indefinite lexically introduces a free variable with its nominal restriction, such as '{x: Single-child(x)}', and then the variable x can be bound by a various kind of quantifier that is provided either in the sentence or in the pragmatic context. Since discourse referents in DRT are formally variables, DRT can deal with (20a–c) in essentially the same manner, introducing various kinds of quantifiers in their discourse representation structures (cf. Kamp et al., 2005). In a sense, whether we can maintain the uniqueness of our simulation language will then depend on whether we can handle such complex data with our operator-free and variable-free simulation language. We leave this for future research, but we can deal with the sort of examples in (20) by introducing the kind of plural objects that we discussed briefly in Section 2.1 or probabilistic elements into simulations (cf. Domingos & Lowd, 2009).

Secondly, as an attempt to ground reasoning in perception or in some analogous form of denotational representations, we would have to mention Jackendoff's works (1990; 2002) and Situation Theory (Barwise & Perry, 1983). However, since their semantic representations are developed primarily by way of the analysis of natural language semantic data, they tend to introduce whatever complexity reasoning and natural language semantics require into their perceptual or denotation structures. In contrast, the main point of our approach is to assume whatever perceptual representations that are best-motivated in the theory of perception such as vision, and then propose a reasoning mechanism that can be well-grounded in such independently motivated perceptual representations. That is why our approach does not include quantifiers or other abstract notions in our formal representations.

3. Linguistic implications

In this section, we show that our simulation-based formulation of reasoning can provide a uniform and cognitively well-founded account of several difficult natural language semantic problems.

Firstly, our simulation-based characterization of reasoning deals with the semantics of both quantificational and referential NPs in a uniform manner. As we discuss below, this formulation enhances the syntax-semantics correspondence and thus, simplifies the pairing of phonological strings with their semantics. Consider the English sentences in (21), together with the first-order formulas in the parentheses representing their semantics.

(21) a. $[_{NP}$ Every student] is noisy. ($\forall x$(Student(x) $\rightarrow$ Noisy(x)))
b. $[_{NP}$ Some student] is noisy. ($\exists x$(Student(x) & Noisy(x)))
c. $[_{NP}$ David] is noisy. (Noisy(david))

Traditional syntactic analyses suggest that all of the three expressions inside the square brackets in (21) occupy the same syntactic position, that is, the external argument (or subject) position of the predicate expression *is noisy*.[24] However, this predicate-argument relation that we can recognize in the natural language structure is maintained only in the first-order formula in (21c). In the quantified first-order formulas in (21a) and (21b), the bits that correspond to the NPs in the natural language syntax are spread all over the formulas and if these actually represent the semantics of the corresponding natural language sentences, it is not clear how we can provide the compositional derivation of such semantics starting with the lexical semantics.

The compositionality problem can be resolved by using a higher-order logic to represent the semantics, such as 'Every(Student)(Noisy)' for (21a) and 'Some(Student)(Noisy)' for (21b), where the higher-order predicates 'Every' and 'Some' denote two-place relations between the set of students denoted by 'Student' and the set of noisy individuals denoted by 'Noisy', so that 'Every(Student)(Noisy)' is true if and only if the set of students is a subset of the set of noisy individuals and 'Some(Student)(Noisy)' is true if and only if the intersection of the set of students and the set of noisy individuals is not empty (cf. Barwise & Cooper, 1981; Gamut, 1991). However, using this higher-order logic analysis alone does not resolve the lack-of-uniformity problem between quantificational and referential NPs. Also, as we suggested in Section 1, it is very hard to make higher-order logic computable in automated reasoning and the inherent undecidability of higher-order logic poses further theoretical challenge.

[24] Chomskyan syntacticians may postulate an implicit movement operation called Quantifier Raising, QR (cf. May, 1977). However, the main motivations for QR are semantic, lacking in clear evidence in the overt phonological expressions (cf. Uchida, 2008). Also, in order to maintain the uniformity for (21a∼c), we would need to quantifier-raise the referential NP in (21c) to an operator position, which is lacking in convincing independent motivations and also makes the syntactic pairing of phonological strings with their semantics unnecessarily complex (the same point can be posed for Montague's analysis, 1973, of referential NPs that type-lifts both the quantificational and referential NPs in the semantics). Also, note that using QR on its own does not solve the compositionality problem associated with the first-order formulas in (21a-b).

In contrast, consider our treatment of (21a–c) in (22a–c).

(22) a. S14, based on: {New(st), Student(st)}/produces: {Noisy(st)}
b. S15, based on: {}/produces: {New(st′), Student(st′), Noisy(st′)}
c. In Reality: {NameOf(dv, David)}
S16, based on: {}/produces: {Noisy(dv)}

Note that our theory formalizes the semantics of all the argument NPs in (21) using the same theoretical item, that is, the argument terms: 'st', 'st′' and 'dv', in (22).[25] The difference in their interpretations arises because of the different structural contexts where these terms are introduced into reasoning. In (22a), 'st' is introduced newly in the basis of a simulation and hence leads to the universal implication. In (22b), 'st′' is introduced newly in the conclusion of a simulation and hence leads to the existential implication. Finally, in (22c), 'dv' is in the knowledgebase of the real world, fixing the reference of 'dv' to the corresponding individual in the real world. Because our account deals with the semantics of both quantificational NPs and referential NPs by way of concrete terms in this way, syntactic arguments in natural language structures are also the arguments of the corresponding predicates in our semantic representations. This makes the relation between the natural language syntax and their semantics simpler, as well as providing an additional explanation for why quantificational NPs occupy argument positions in natural language structures, that is, they are arguments in the syntax since they are arguments in the semantics as well.

We are not arguing that all the NPs that are in the same syntactic position should receive the same semantic interpretation in one's theory. The issue here is more a matter of simpler syntax-semantic relation, such as the simple instruction to interpret each argument NPs in the surface syntax as an object term in the simulation language, together with further parametric specification such as interpreting a universal quantifier NP as a new term in the basis of the simulation in question while interpreting a proper name NP as an object term in the reality set, etc. Of course, in order to prove that the natural language interpretation into our simulation language is formally simpler than the natural language interpretation into standard FOL, we will need to provide the precise syntax and semantics of our simulation language.[26] However, given the semi-formal nature of this paper, we leave that for a different paper.

Another linguistic merit of our simulation theory of reasoning is our treatment of negative polarity items. Pre-theoretically, it is clear that negative polarity items (NPI) such as *any* occurs only in a particular structural environment. However, providing a uniform theoretical definition of this structural environment is not easy. Consider (23).

(23) a. Jack does not like *any* girl
b. If Meg reads *any* book, I am surprised
c. Jack likes every student who reads *any* logic book
d. *Jack reads *any* book
e. *Jack likes a student who reads *any* logic book

The article *any* in (23a–c) can be interpreted as NPI, whereas *any* in (23d) and (23e) can only be interpreted as the so-called 'free choice' *any* (which has different interpretation[27] and requires obligatory phonological focus). It is not clear whether the natural language syntactic configurations that license *any* in (23a–c) form a natural class. It is not very explanatory to simply enumerate each syntactic configuration that licenses an NPI, such as 1) in the scope of a negation as in (23a), 2) inside the *if*-clause as in (23b) and 3) inside the nominal restriction of a universally quantified NP as in (23c) (but not in the nominal restriction of an indefinite). Thus, some semanticists (Ladusaw, 1979) assume that NPIs require a particular semantic environment, that is, the so-called downward entailing environment. They use a certain logical entailment test that arguably helps identify the downward entailing environment.

We do not discuss whether the notion of downward-entailing environment can provide a theoretically uniform account of NPI distributions in this paper. The general introduction to the topic in Kearns (2000), Heim and Kratzer (1998) and critical reviews such as Heim (1984) suggest that it is at least not a trivial task. As one example, however, which Kadmon and Landman (1993) discussed in details, some propositional attitude adjectives license NPI, such as *I'm surprised that Bob made any progress*, but at least apparently, the embedded proposition 'Bob made any progress' is not in the DE environment in the standard definition as in Ladusaw (1979). Also, Ladusaw's analysis does not provide a clear common ground between the NPI *any* and the free choice *any*, as in (23d–e), in which *any* has some sort of universal interpretation similar to *every*. This indicates a lexical ambiguity for the natural language expression *any*, but as Kadmon & Landman suggested, further data analysis goes against such a lexical ambiguity hypothesis, and generally speaking, when the linguistic form (i.e., *any*) is the same, if there is a way of inferentially deriving the difference from the shared basic meaning lexically encoded by that form, that analysis is preferable to the analysis that assumes that the linguistic form is mapped into two, unrelated interpretations.

In contrast, our simulation-based characterization of reasoning can straightforwardly provide a uniform account of NPI distribution in a cognitively well-grounded manner.

[25] DRT can achieve a similar kind of uniform NP semantics, which we discussed at the end of Section 2.

[26] It will also be worth considering whether such a formal simplicity is important in the natural language semantic theory.

[27] *I do not eat anything* with its NPI interpretation means that the speaker eats nothing, while *I do not eat ANYTHING* with its free choice interpretation means that it is not the case that the speaker eats everything.

We also show later that we can deal with the free choice *any* in essentially the same manner, with a minimal parametric difference about where in simulations the individual objects corresponding to the NPI and free-choice *any* are introduced.

First, we represent the semantics of (23a), (23c) and (23e) using our simulation language.

(24) a. In Reality: {NameOf(jk, Jack)}
S20, based on: {New(gl), Girl(gl), Like(jk,gl)}/
produces: {⊥}
b. In Reality: {NameOf(jk, Jack)}
S21, based on:
{New(st), New(lb), Student(st), LogicBook(lb), Read(st,lb)}
produces: {Like(jk,st)}
c. In Reality: {NameOf(jk,Jack)}
*S22, based on: {}
produces:{New(st), New(lb), Student(st), LogicBook(lb), Read(st,lb), Like(jk,st)}

Note that (24a) correctly represents the NPI reading of (23a) and (24b) represents the NPI interpretation of (23c). S20 in (24a) means that if we simulate on the basis that Jack as we know in the real world likes a newly introduced individual, 'gl', which is a girl, then that simulation fails. This implies that for any girl, 'gl', if we assume that Jack likes 'gl', then that simulation fails. This in turn implies the correct NPI interpretation of *Jack does not like any girl*, that is, 'Jack likes no girl.' Similarly, since both 'st' and 'lb' are introduced newly in the basis of S21 in (24b), S21 implies that for any student 'st' and for any logic book 'lb' if 'st' reads 'lb', then Jack as we know in the real world likes 'st' (note that it is enough for each student to read at least one logic book to be liked by Jack in this reading). Again, this captures the correct NPI interpretation of (23c). Next, consider why NPIs are licensed in (23a) and in (23c) according to (24a) and (24b). Note that the terms that correspond to the NPs whose determiners are the NPI *any*, that is, 'gl' for *any girl* and 'lb' for *any logic book*, both appear in the basis of a simulation. Though we do not provide the exact semantic composition rule from the natural language expressions, the rough licensing rule of the NPI *any* would be like the following: 'The term that corresponds to the argument NP whose determiner is the NPI *any* **needs to** appear in an existing basis.' Here 'existing basis' means that the basis must occur in a simulation calculated based on the natural language structure independent of the presence of the NPIs in question. For example, the basis-conclusion structure of S20 in (24a) is constructed because of the negation *not* in *Jack does not like any girl*, and the basis-conclusion structure of S21 in (24b) is constructed because of the semantics of the universal determiner *every* in *every student who*.... Similarly, the *if*-conditional in (23b) generates a similar basis-conclusion simulation structure in our formulation and the term for the NPI argument *any book* appears in the basis of that simulation.

In contrast, the reason why (23e) does not have the NPI interpretation of *any logic book* is that the term that corresponds to the NP whose determiner is the NPI *any* does not appear in the basis of a simulation structure that is constructed based on the natural language structure of (23e), as shown in (24c), which again is constructed independent of the presence of the NPI in (23e). Note that the term in question, that is, 'lb', which is a logic book (i.e., for *any logic book* in (23e)), appears in the conclusion of S22. In order for the NPI to be licensed, 'lg' would need to appear in the basis of this simulation. If we ignore the NPI licensing, (24c) is a well-formed expression in our simulation language. It means that Jack (as we know in the real world) likes a student who reads a logic book.

The above discussion indicates that our simulation language can provide a uniform licensing condition for the NPI *any*. That is, all the terms that are associated with the NPI *any* must occur in the basis of a simulation (which is calculated from the syntactic structures of the natural language sentences independent of the NPI). Investigating whether we can cover all the natural language data that include the NPI *any* is beyond the scope of this paper, but it is clear that we can cover all the constructions in (23) in a uniform manner.

Next, we briefly sketch our treatment of the free choice *any*. Consider (23d) above. Although that English sentence cannot have the NPI interpretation, it can have the free choice *any* interpretation, which is equivalent to 'Jack reads every book.' We analyze the free choice interpretation of *Jack reads ANY logic book* in terms of the simulation structure in (25).

(25) In Reality: {NameOf(jk, Jack)}
S23, based on: {New(lb′), LogicBook(lb′)}/
produces: {Read(jk, lb′)}

Note that the term 'lb′' in S23 does appear newly in the basis of this simulation. The difference between the NPI *any* as in *Jack does not read any logic book* and the free choice *any* in *Jack reads ANY logic book* in our account is that although both the NPI *any* and the free choice *any* require the associated terms to occur in the basis of a simulation, only the free choice *any* can construct its own basis-conclusion structure in our semantic representation (just like *every* in *Jack reads every book* can construct an analogous basis-conclusion as in S23). In this sense, our account provides a certain degree of uniformity between the NPI *any* and the free choice *any*.

In the above formulation, we specified the parametric difference between the NPI *any* and the free-choice *any* as part of their lexical specifications. However, in order to maintain the uniform lexical encoding between the NPI *any* and the free-choice *any* in a stricter manner, we could derive the difference inferentially. For example, we can assume that the lexical item *any* (whether it is NPI or free-choice) encodes that each expression in the form of *any* plus common noun, such as *any girl* in (23a), requires the object term corresponding to this NP to appear in the basis of a simulation. With (23a), the object term corresponding to *any girl*, that is, 'gl', is already in the basis of a simulation in (24a), according to the translation of the entire English sentence in (23a) into the simulation in (24a) and for reasons of parsimony, we do not need to do anything special for satisfying the lexical requirement of *any* as above. This leads to the NPI reading. Now, for *any logic book* in *Jack reads ANY logic book* above, the lexical requirement of *any* is the same, but simply translating the sentential structure *Jack reads NP* (ignoring the internal structure of *NP* = *any logic book*) into a simulation does not put the object term for *any logic book* in the basis of a simulation.[28] Hence, as some sort of last resort mechanism, we can construct a new basis-conclusion structure just for satisfying the lexical requirement of *any* and put the object term for *any logic book* in that basis, as in S23 in (25). This last-resort measure leads to the free choice interpretation. The nearly obligatory phonological focus assignment on the free choice *any* (as in *ANY logic book*) might be a reflection of the last-resort nature of the basis construction just for satisfying the requirement of the lexical item *any*, and a similar explanation can be provided for why the free choice interpretation of *any* is generally marked, often requiring further pragmatic/contextual support.

In the above explanation, we adopted the viewpoint of natural language interpretation for convenience, but we can translate simulations into natural language expressions in a similar manner.

We leave for future research further details about the lexical semantics of *any*, as well as the difference between the free choice *any* and the universal quantifier *every*, but we briefly discuss one complication. Consider (26).

(26) a. If Meg reads *any* book, I am surprised. (NPI)
b. If Meg reads *ANY* book, I am surprised. (Free choice??)
c. If Meg reads every book, I am surprised. (Universal quantifier)

Notice that it is difficult to get the free choice interpretation of *ANY* in (26b). Even if we put a phonological focus on *ANY*, the most salient interpretation still seems to be truth-conditionally the same as the NPI interpretation (in which the speaker is surprised if Meg reads a book, whatever book it is). The subtle difference between (26a) and (26b) is that (26b) emphasizes the lack of any exception for this conditional statement (i.e., an additional emphasis on 'whatever book it is' above). The 'last resort' construction of a new basis that leads to the free-choice interpretation, which we speculated in the previous paragraph, might explain why it is exceptionally difficult to get the free choice interpretation with (26b). That is, the simulation structure for (26b) already contains a basis that is constructed by the conditional form *If…, (then)…*, in which we can put the object term for the *any* NP. Consider (27).

(27) In reality, {NameOf(meg, Meg), The-Agent(a)}
S24, based on: {New(b), Book(b), Read(meg,b)}/
Produces: {Surprised(a)}

We regard the speaker of the utterances in (26) as the agent in the simulation in (27). Again, since the conditional sentence constructs a basis-conclusion structure independent of *any*, we can put the new object for the *any* NP in that basis, which leads to the NPI interpretation.

Following our 'inferential' account of free-choice *any*, we can speculate that the last-resort construction of a new basis-conclusion structure can only apply when there is a very strong pragmatic or contextual justification. As for our data-judgment, we could probably force the free choice interpretation with (26b) by providing such an adequate pragmatic context, but the free choice reading is easier to get with *If Meg eats ANY food*, *I will be worried* (*since some food can be bad for Meg's health*), which has the same *if*-conditional structure. We assume here that this free choice interpretation is truth-conditionally the same as the semantics of *If Meg eats every food, I will be worried,* which can be characterized in terms of (28).

(28) In reality, {NameOf(meg, Meg), The-Agent(a)}
S25, based on: {New(f), Food(f)}/produces {Eat(meg,f)}
S26, based on: {Hold(S25)}/produces {Worried(a)}

Note that S25 has the basis-conclusion structure that is independently constructed because of the *any* NP. Again, comparing (26b) with the above alternative conditional sentence, we can tell that when *ANY* appears in the scope of the conditional *if*, the availability or saliency of the free choice interpretation seems to depend on both semantic and pragmatic factors. Thus, a more precise formulation of our account will require further linguistic data analysis in future research. In this paper, we only suggest that the 'last-resort' account of the free-choice interpretation of *any* might be explaining why getting the free-choice interpretation is extra-difficult when the NP in question appears in a traditional 'downward-entailing' environment, as in (26b).

The above discussion does not mean that we always account for the free choice *ANY* and the universal determiner *every* in terms of the same simulation. Explaining precisely how we distinguish them is beyond the scope of

[28] For example, the sentence *Jack reads Brighton Rock* will be translated into our formal representation: "In Reality: {NameOf(jk, Jack}, NameOf(br, Brighton Rock)}, S23': {}/produces:{Read(jk,br)}." Since the term 'br' for the object NP *Brighton Rock* appears in the conclusion of S23' according to this translation, the object NP *any book* in (23d) is also placed in the conclusion of a simulation of the same form, which violates the requirement of the NPI in our account. According to the inferential derivation of the free choice interpretation that we provide shortly, we can save this situation by forcing the construction of a new basis-conclusion structure as the last-resort measure and then putting the term for the *any*-NP in the basis, but this last-resort measure leads to the free-choice interpretation. See the following main text.

this paper, but we briefly mention three points. First, as we discussed with the discourse example, "I went to the party yesterday. Everybody was drunk", in Section 2, the universal NP tends to be associated with a particular group of individuals that are salient in the context. In contrast, the free choice *any* tends to expand the domain as wide as possible (cf. Kadmon & Landman, 1993) and it is not easy to use it to describe a particular state of affairs that one has observed himself. For example, note that using the free choice *ANY* instead of *every* in the above discourse example makes it sound really odd.[29] We can accommodate this difference between *every* and *ANY* using the domain restriction that we suggested in Section 2.

Secondly, the free choice *ANY* is often associated with some sort of modal interpretation. For example, consider the sentence *ANY logician can prove that*. Once we extend our simulation language to accommodate modal reasoning (cf. Uchida et al., 2013), we can differentiate *every* and *ANY* in terms of their (possibly) different interactions with modal elements.

Thirdly, in this paper, we only focus on the truth-conditional semantics, since our main goal is to show that our operator-free and variable-free language can represent the semantics of FOL. However, it is well-known that natural language expressions may carry non-truth conditional interpretations. For example, Powell (2003) differentiates the semantics of the definite article *the* and the demonstratives *this* and *that* in terms of their procedural instructions about which individual objects (or which 'individual concepts', in his term) are accessed. The different procedural instructions may lead to the same truth-conditional content, but that does not mean that the lexical encodings of *the*, *this* and *that* are the same.

We also omit our detailed treatment of the above-mentioned *I'm surprised that Bob made any progress*, which is at least apparently problematic for the traditional downward entailing analysis in Ladusaw (1979). Our account will be similar to Kadmon & Landman's (1993) in that we associate some sort of negative implication with the proposition in question. For example, the above sentence indicates that the speaker had expected that Bob would not make any progress. This implicit negative connotation is then formalized by way of simulations in the form of S20 in (24a) and the object term for *any progress* can be placed in the basis of this simulation structure.

Finally, perceptual simulations can provide a cognitively well-grounded account of the so-called branching-quantifier interpretations (Barwise, 1979; Henkin, 1959), which standard first-order logic cannot represent. Consider (29).

(29) a. A friend of every villager and a relative of every townsman hate each other.

b. $\forall x \forall y \exists u \exists v((\mathrm{Villager}(x) \& \mathrm{Townsman}(y)) \rightarrow (\mathrm{FriendOf}(u,x) \& \mathrm{RelativeOf}(v,y) \& \mathrm{Hate}(u,v) \& \mathrm{Hate}(v,u)))$

The problem of the linear first-order formula in (29b) for representing the semantics of the English sentence in (29a) is that in this logical form, both the existential quantifiers take scope narrower than both the universal quantifiers. As a result, this logical form means that the friend u can covary with each townsman y as well as with each villager x. Similarly, the form means that the relative v can covary with each villager x as well as with each townsman y. Intuitively, in (29a), the friend only covaries with each villager and the relative covaries only with each townsman. However, this 'partial' covariance is impossible to represent using the traditional first-order formula, since the four quantifiers as in (29b) need to be linearly ordered in this logic. In order to represent the above reading, we need to use the so-called branching quantifier, as in (30), cf. Hintikka and Sandu (1997).

(30) $$\begin{pmatrix} \forall x \exists u \\ \forall y \exists v \end{pmatrix} ((V(x) \wedge T(y)) \rightarrow (Fof(u,x) \wedge Rof(v,y) \wedge H(u.v) \wedge H(v,u))$$

Now, unlike the traditional linear first-order logic, our simulation language can use (31) to represent the same covariation that (30) expresses.

(31) S27, based on: {New(v), Villager(v)}/produces: {FriendOf(fr,v)}
S28, based on: {New(tw), Townsman(tw)} produces: {RelativeOf(rel,tw)}
S29, based on: {Villager(v),Townsman(tw)} produces: {Hate(fr,rel), Hate(rel,fr)}

S27 and S28 in (31) are straightforward, both of which precede S29. We can calculate these precedence relations by way of the subset relations between the basis sets, that is, {Villager(v)}⊂{Villager(v),Townsman(tw)} between S27 and S29 and {Townsman(tw)}⊂{Villager(v),Townsman(tw)} between S28 and S29, ignoring the 'syntactic-sugar' formulas 'New(v)' and 'New(tw)', which we have included for presentational convenience, that is, for showing where the terms are newly introduced explicitly.[30] Since the individual 'v' is introduced newly in the basis of S27, given the above precedence relation, S27 has the universal implication, 'Each villager has a (possibly) different friend.'

[29] Probably, we can still use the free choice *any* in this discourse with an expression that explicitly pins the domain to a particular group of people, such as *ANYbody there was drunk*, although even this does not save the oddity too much. This oddity might also be because of the modal element normally associated with *any*, which we discuss next.

[30] When we omit these syntactic sugar formulas, the precedence relation between simulations will become crucial for calculating where each term is newly introduced, but we omit a more precise formulation here.

In the same way, the individual 'tw' is introduced newly in the basis of S28, and thus, S28 has the universal implication, 'Every townsman has a different relative.' Now, the terms 'v' and 'tw' in S29 are not new and hence refer back to 'v' and 'tw' in S27 and S28. The same applies to the terms 'fr' and 'rel.' Note that 'fr' is introduced newly in the conclusion of S27, and hence the friend 'fr' covaries only with 'v', that is, each villager. Similarly, 'rel' is introduced for the first time in the conclusion of S28, whose basis contains only 'tw' as new individual. Thus, 'rel' cavaries only with 'tw', that is, each townsman, as desired. Neither 'fr' nor 'rel' picks up the force of universal quantification since neither is in a basis of a simulation. The above precedence relations between S27 and S29 and between S28 and S29 are important, since if 'v', 'tw', 'fr', 'rel' were taken to be introduced newly in S29, it would lead to a completely different implication.

The logic with branching quantifiers used in (30) has the same expressive power as second-order monadic logic (cf. Väänänen, 2007). This indicates that our simulation language also is as expressive as a second-order monadic logic. But we leave this formal investigation for another paper.

As we discussed in Section 1, we have accounted for the three kinds of linguistic data using only the basic mechanisms of perceptual simulations. In this framework, it makes sense that both quantificational and referential NPs occupy the same structural positions in overt natural language strings, since both of them refer to the objects that appear in simulations. Also, given the basis-conclusion structure that we assume in simulations, it is natural to expect that there exists a natural language expression that requires an object to be placed in the basis of a simulation. Finally, we have explained above how successive constructions of multiple simulations (some of which may refer back to the previous simulations) can naturally explain the existence of partial scope relations as in branching quantifier propositions. These considerations suggest that the data that we have discussed in this section can be taken as a support for our characterization of reasoning (and therefore natural language semantics) by way of perceptual simulations.

4. Conclusion

We have provided an account of reasoning that maintains the basic correspondence between the elements of reasoning and of perception. Because of the great superficial differences between the elements of logical reasoning and perceptual processes, it was not at first sight apparent how our simulation theory of reasoning can deal with abstract non-perceptual notions involved in human reasoning, such as quantification and propositional negation. To address this challenge, we have formulated our theory of reasoning in terms of simulations with perceptual representations. We have shown that we can infer the notions of quantification and negation as implications of particular simulation results. We have also shown that our simulation-based formulation of reasoning can deal with all the other propositional operators and connectives in traditional first-order logic. This means that our theory can potentially account for many aspects of natural language semantics and reasoning more straightforwardly than a theory using standard first-order logic.

We have also demonstrated that our theory also has independent merits in linguistic data analysis. We have shown our uniform semantic treatment of both quantificational and referential NPs. We have also shown that our simulation-based account can provide a uniform licensing configuration for the negative polarity item (NPI) *any*. Finally, we have shown that our simulation language can represent the so-called branching quantifier interpretations, which a linear first-order logic cannot represent. We have accounted for a branching-quantifier proposition in terms of two mutually-independent simulations which are both incorporated into a later simulation. We believe that our cognitively well-founded account is explanatory with regard to why human reasoning involves partial scope specification that is involved in branching-quantifier propositions.

This paper has ignored uncertainty and counterfactual reasoning. However, it is possible in future work that we can straightforwardly deal with the uncertainty of human reasoning by adding 'weights' to simulations (cf. Domingos & Lowd, 2009). For example, we can state that the output of a simulation follows from the basis with less than 100% certainty. Since our simulations reflect the logical 'hypothesis-conclusion' relation in terms of some kind of 'causal' relation from the input to the output, our treatment of counterfactual reasoning will also be straightforward. The only remaining issue with counterfactual simulations will be how to calculate which facts in reality are 'inherited by' (i.e., maintained in) counterfactual simulations and how to represent such inheritance or lack of inheritance in our simulation language. We leave the technical details for future research (see Uchida et al., 2013 for basic formulations).

Acknowledgements

The authors would like to thank Nicholas Allott, Paul Bello, Selmer Bringsjord, Robyn Carston, Robert Truswell, Catherine Wearing and the members of the Human-Level Intelligence Laboratory at RPI for discussions on this work and for comments on earlier drafts of this paper. The authors are also grateful to the editor and the anonymous reviewers whose comments on the earlier draft have helped improve the content and the presentation of the paper. This work was supported in part by grants from the Office of Naval Research (N000140910094), the Air Force Office of Scientific Research (FA9550-10-1-0389), and MURI award (N000140911029).

References

Abusch, D. (1994). The scope of indefinites. *Natural Language Semantics, 2*(2), 83–135.

Allott, N., & Uchida, H. (2009). Natural language indicative conditionals are classical. *UCL Working Papers in Linguistics, 21*, 1–17.

Barsalou, L. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577–609.
Barsalou, L. (2003). Situated simulation in the human conceptual system. *Language and Cognitive Processes, 18*, 513–562.
Barsalou, L. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London: Biological Sciences, 364*, 1281–1289.
Barsalou, L., Pecher, D., Zeelenberg, R., Simmons, W. K., & Hamann, S. B. (2005). Multi-modal simulation in conceptual processing. In W. Goldstone, B. Love, A. Markman, & P. Wolff (Eds.), *Categorization inside and outside the lab: Essays in honor of Douglas L. Martin* (pp. 249–270). Washington: American Psychological Association.
Barwise, J. (1979). On branching quantifiers in English. *Journal of Philosophical Logic, 8*, 47–80.
Barwise, J., & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy, 4*, 159–219.
Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge, Mass: MIT Press.
Carlson, G. (1981). Distribution of free-choice *any*. In R. Hendrick, C. Masek, & M. Miller (Eds.), *Papers from the seventeenth regional meeting of the Chicago Linguistic Society*. CLS: University of Chicago.
Carpenter, B. (1997). *Type-logical semantics*. Cambridge, Mass.: MIT Press.
Cassimatis, N. L., Murugesan, A., & Bignoli, P. G. (2009a). Reasoning as simulation. *Cognitive Processing, 10*(4), 343–353.
Cassimatis, N. L., Murugesan, A., & Bignoli, P. G. (2009b). Inference with relational theories over infinite domains. In H. Chad Lane & H. W. Guesgen (Eds.), *Proceedings of the twenty-second international FLAIRS conference 2009* (pp. 21–26). Menlo Park, California: AAAI Press.
Cassimatis, N. L., Bignoli, P. G., Bugajska, M. D., Dugas, S., Kurup, U., Murugesan, A., et al. (2010). An architecture for adaptive algorithmic hybrids. *IEEE Transactions on Systems, Man, and Cybernetics, Part B, 40*(3), 903–914.
Chierchia, G. (1995). *Dynamics of meaning: Anaphora, presupposition and the theories of grammar*. Chicago: The University of Chicago Press.
Domingos, P., & Lowd, D. (2009). *Markov logic: An interface layer for artificial intelligence*. San Rafael, CA: Morgan & Claypool.
Duan, Y. (2003). *Topology adaptive deformable models for visual computation*. PhD thesis, State University of New York, Stony Brook, NY.
Fagin, R., Kolaitis, P. G., Miller, R. J., & Popa, L. (2005). Data exchange: Semantics and query answering. *Theoretical Computer Science, 336*(1), 89–124.
Fitting, M. (1996). *First-order logic and automated theorem proving* (2nd ed.). Berlin: Springer-Verlag.
Fox, C., & Lappin, S. (2004). An expressive first-order logic with flexible typing for natural language semantics. *Logic Journal of the International Group for Pure and Applied Logic, 12*, 135–168.
Gamut, L. T. F. (1991). *Logic, language and meaning* (Vols. 1–2). Chicago: University of Chicago Press.
Garbarini, F., & Adenzato, M. (2004). At the root of embodied cognition: Cognitive science meets neurophysiology. *Brain and Cognition, 56*, 100–106.
Gauker, C. (2005). *Conditionals in context*. Cambridge, Mass.: The MIT Press.
Goldman, A. (2002). Simulation theory and mental concepts. In J. Dokic & J. Proust (Eds.), *Simulation and knowledge of action* (pp. 1–20). Oxford: Blackwell.
Gordon, R. (1995). The simulation theory: Objections and misconceptions. In M. Davis & T. Stone (Eds.), *Folk psychology* (pp. 100–122). Oxford: Blackwell.
Gordon, R., & Cruz, J. (2002). Simulation theory. In L. Nadel (Ed.), *Encyclopedia of cognitive science*. Macmillan.
Harnad, S. (1990). The symbol grounding problem. *Physica D, 42*, 335–346.
Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. PhD thesis, University of Massachusetts.
Heim, I. (1983). File change semantics and the familiarity theory of definiteness. In R. Bäuerle, C. Schwarze, & A. von Stechow (Eds.), *Meaning, use and the interpretation of language* (pp. 164–190). Walter de Gruyter & Co.
Heim, I. (1984). A note on negative polarity and downward entailingness. In C. Jones & P. Sells (Eds.), *Proceedings of NELS 14, GLSA*. Amherst: Linguistics Dept., University of Massachusetts.
Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
Henkin, L. (1959). Some remarks on infinitely long formulas. In *Infinite methods. Proceedings of the symposium on foundations of mathematics* (pp. 167–183). Warsaw: Panstwowe Wydawnictwo Naukowe and Pergamon Press.
Hesslow, G. (2002). Conscious thought as simulation of behavior and perception. *Trends in Cognitive Science, 6*(6), 242–247.
Hintikka, J., & Sandu, G. (1997). Game-theoretic semantics. In J. van Benthem & A. ter Meulen (Eds.). *Handbook of logic and language* (Vol. 1, pp. 361–410). Cambridge, Mass.: MIT Press.
Hodges, W. (1993). *Model theory*. Cambridge: Cambridge University Press.
Horn, L. (1972) *On the semantic properties of logical operators in English.* Ph. D. Thesis, UCLA, reproduced by the IULC, 1976.
Jackendoff, R. (1990). *Semantic structures*. Cambrige, Mass.: MIT Press.
Jackendoff, R. (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford: Oxford University Press.
Kadmon, N., & Landman, F. (1993). Any. *Linguistics and Philosophy, 16*, 353–422.
Kamp, H. (1973). Free choice permission. In *Proceedings of the Aristotelian Society*, London.
Kamp, H., & Uwe, R. (1993). *From discourse to logic: Introduction to model-theoretic semantics of natural language, formal logic and discourse representation theory*. Dordrecht: Kluwer Academic Publishers.
Kamp, H., van Genabith, J., & Uwe, R. (2005). Discourse representation theory. In D. Gabbay & F. Guenthner (Eds.). *Handbook of philosophical logic* (Vol. 15, pp. 125–394). Dordrecht: Springer.
Kearns, K. (2000). *Semantics*. New York: Macmillan.
Kratzer, A. (2005). Indefinites and the operators they depend on: From Japanese to Salish. In G. N. Carlson & F. J. Pelletier (Eds.), *Reference and quantification: The Partee effect* (pp. 113–142). Stanford: CSLI Publications.
Ladusaw, W. A. (1979). *Polarity sensitivity as inherent scope relations*. PhD thesis, University of Texas, Austin, reproduced by the IULC, 1980.
Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: The University of Chicago Press.
Lakoff, G., & Núñez, R. (2001). *Where mathematics comes from: How the embodied mind brings mathematics into being*. New York: Basic Books.
Link, G. (1983). The logical analysis of plurals and mass terms. In R. Bäuerle, C. Schwarze, & A. von Stechow (Eds.), *Meaning, use and the interpretation of language* (pp. 303–323). Berlin: Walter de Gruyter & Co.
Luckiesh, M. (1995). *Visual illusions*. New York: Dover Publications Inc.
May, R. (1977). *The grammar of quantification*. PhD thesis, MIT, Cambridge, Mass.
Merchant, J. (2004). Fragments and ellipsis. *Linguistics and Philosophy, 27*(6), 661–738.
Montague, R. (1973). The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, & P. Suppes (Eds.), *Approaches to natural language: Proceedings of the 1970 Stanford workshop of grammar and semantics* (pp. 221–242). Dordrecht: Reidel.
Oaksford, M., & Chater, N. (2006). *Bayesian Rationality*. Oxford: Oxford University Press.
Peregrin, J., & von Heusinger, K. (2004). Dynamic semantics with choice functions. In H. Kamp & B. Partee (Eds.), *Context dependence in the analysis of linguistic meaning* (pp. 255–274). Amsterdam: Elsevier.
Powell, G. (2003). *Language, thought and reference.* PhD thesis, University College London.

Riesen, K., & Bunke, H. (2012). Classification and clustering of vector space embedded graphs. In C. H. Chen (Ed.), *Emerging topics in computer vision and its applications* (pp. 49–70). New Jersey: World Scientific.

Spivey, M. J. (2007). *The continuity of mind*. New York: Oxford University Press.

Stainton, R. (2004). The pragmatics of non-sentences. In L. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 266–287). Oxford: Blackwell.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science, 268*(5217), 1632–1634.

Uchida, H. (2008). *Scope alternation and type logical grammar.* PhD thesis, University College London.

Uchida, H. (2013). Logic in pragmatics. *Lingua, 133*, 336–359.

Uchida, H., & Cassimatis, N. L. (2010). Singular NP lattice without explicit quantification. In *Eighteenth annual meeting of the European Society for philosophy and psychology, ESPP2010* (pp. 187–189). Germany: Bochum & Essen.

Uchida, H., Cassimatis, N. L., & Scally, J. R. (2012). Perceptual simulations can be as expressive as first-order logic. *Cognitive Processing, 13*(4), 361–369.

Uchida, H., Cassimatis, N. L., & Scally, J. R. (2013). Modal reasoning as simulation. In *IACAP2013 (the annual meeting of the International Association of Computing and Philosophy)*. University of Maryland at College Park, USA.

Väänänen, J. (2007). *Dependence logic*. Cambridge: Cambridge University Press.

Wilcox, T. (1999). Object individuation: Infants' use of shape, size, pattern, and color. *Cognition, 72*, 125–166.

Please cite this article in press as: Uchida, H., et al. Characterization of reasoning in terms of perceptual simulation. *Cognitive Systems Research* (2013), http://dx.doi.org/10.1016/j.cogsys.2013.11.001