

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Meaning in Context: Ontologically and linguistically motivated representations of objects and events

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1724262> since 2020-01-21T15:18:17Z

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Meaning in Context: Ontologically and Linguistically Motivated Representations of Objects and Events

Valerio Basile^a, Tommaso Caselli^b and Daniele P. Radicioni^{a,*}

^a *University of Turin, via Pessinetto 12, Turin, Italy*

E-mails: valerio.basile@unito.it, daniele.radicioni@unito.it

^b *CLCG, University of Groningen, Oude Kijk in't Jaatstraat, 26, 9712 EK, Groningen The Netherlands*

E-mail: t.caselli@rug.nl

Context is pervasive in all activities involving human beings as well as computer systems. It affects numerous aspects of our lives: how we understand the world, our communication, as well as the planning, carrying, and outcomes of activities. When focusing on understanding natural language, context plays a pivotal role that affects various facets ranging from the interpretation of speech signal to the identification of word meaning, composition of phrases and sentences as well as the intention of the discourse.

Context impacts the conceptualisation of human experience. Objects and events assume different cognitive salience according to their context of occurrence, thus determining access to partial relevant information rather than to all information. Understanding a piece of natural language requires a continuous interaction between abstract conceptualisations and their actual use or occurrence in a specific communicative situation. A possible example to clarify this distinction could be that of an orange being passed between two children, and that of the same orange peeled on a table: in the former context the roundness of the object prevails over other features, traits, or aspects that characterise the concept of “the fruit orange”. In the latter context, the edible features are those mostly relevant and conveyed. Similar mechanisms are at play with events. A well known example is that of “commercial events”. Different surface forms (i.e., verbs) can be used to express related, if not the same, knowledge about the the same scenario. For instance, the verbs *buy* and *sell* offer different interpretations of the same event, where none of them puts the focus on the object of the transaction, but rather on its participant. An extreme situation, in this case, is when different interpretations are imposed on the same event. For instance, a fatal shooting by the police can be interpreted as necessary self-defence (by the police officers) or as a murder (by the friends and family of the victim). Similar selectional mechanisms underlie figurative uses of word meanings, such as metonymy and metaphors among others, that intrinsically characterize the interface between knowledge and language — see the work by Gangemi et al. (2018) for a recent study on the ontological treatment of conceptual metaphors.

*Corresponding author. E-mail: daniele.radicioni@unito.it.

Embarking on an exhaustive review of the study of context in different disciplines such as Linguistics, Philosophy of Language, Cognitive Science and Artificial Intelligence, to mention a few, is out of the scope of this contribution. Classical model-theoretic semantics (such as proposed, e.g., by Frege, Russell, Carnap, and Tarsky, among others) was mainly concerned with the meaning of linguistic expressions, either underestimating or completely ignoring the role played by context or language in use: *“Indeed, meanings were abstracted from the linguistic items that have them, and (indicative) sentences were often equated with statements, which in turn were equated with propositions”* (Bach, 2004, p. 463). However, resistance to this vision has always been present. Malinowski (1922, 1935) was the first to develop a theory of context, that later influenced the work of Firth (1957), which is based on co-occurrence patterns in large language corpora (*“you shall know a word by the company it keeps”*). Morris (1938) followed the triadic semiotic relation initially proposed by Peirce, distinguishing between a semantic dimension (the study of the relations of signs to the objects to which they are applied), a pragmatic dimension (the study of the relations of signs to interpreters) and a syntactic dimension (the study of the relation of signs to one another). In a drastical and anti-theoretical vision, Wittgenstein (1953) defines meaning as *“use in the language”*. Situation Semantics (Barwise and Perry, 1983) proposes a theory of meaning and information content where context-dependence becomes an essential feature, rejecting possible world semantics assumptions and formalism. From a different perspective, Stalnaker (1998) identifies context as the body of information, or represented knowledge, that has been learned and is being presumed at a certain point by attending a certain discourse. This is reminiscent of notions of context developed in other formal semantics theories such as Discourse Representation Theory and File Change Semantics (Kamp and Reyle, 1993; Heim, 1983).

A different take on context is somehow implied in Frame Semantics (Fillmore, 1976, 1982). Case frames are presented as elements that characterise scenes, or abstract situations, that necessarily involve some sort of meaning representations structured in accordance with some motivating context: *“words, and other linguistic forms and categories are seen as indexing semantics or other cognitive categories which are recognized as participating in other larger conceptual structures of some sort, all of this made intelligible by knowing something about the kinds of settings or contexts in which a community has found a need to make such categories available to its participants”* (Fillmore, 1982, p. 119).

In the domain of ontology and ontology modelling, the notion of context is not less complex nor debated. Ontologies are agreements on conceptualisations (Uschold and Gruninger, 1996) and their designing processes require that the intended models are made explicit (Guarino, 1997), so as *“to restrict the number of possible interpretations only to the ones specified by the basic ontological categories used to describe the domain”* (Richards and Simoff, 2001, p. 122).

A number of contextual ontology models have been proposed, mostly intended for software development in the fields of pervasive and ubiquitous computing applications such as, e.g., the work by Gu et al. (2004); Baldauf et al. (2007). Context is intended as a device to grasp surrounding information to enable systems and devices to effectively provide appropriate information or services to the users, by adapting their behaviours according to the current situation. Of course, many definitions for context have been provided, most of which are listed and surveyed in a work by Zainol and Nakata (2010). More recently, on a different line of research, context has been investigated by starting from the analysis of relationships and events (Guarino and Guizzardi, 2015, 2016). In this setting, events emerge from *scenes* as the result of a cognitive process that focuses on relationships; within this perspective, the scene is the representational mechanism hosting the relationships in their dynamics. The scene is herein the ontological counterpart of our notion of context.

In parallel with the work of ontologists, the Computational Linguistic community organically developed its own modeling for context, as a stepping stone for Natural Language Understanding and Generation. One of the major paradigmatic shifts in NLP is the use of so called *distributional* representations. This paradigm is based on a re-discovery of Firth's notion of meaning (Firth, 1957), in an attempt to address some of the limitations of knowledge-based approaches, in particular the lack of dynamicity, or context. Distributional semantics research produced a variety of models: the earlier algorithms, such as Latent Semantic Analysis (Landauer and Dumais, 1997) and Latent Dirichlet Allocation (Blei et al., 2003) were based on manipulations of co-occurrence matrices and exhibited varying degrees of interpretability. Later models are instead based on the prediction of the contexts of words, typically employing neural networks (Mikolov et al., 2013).

It is a matter of fact that the context delivered in current mainstream Natural Language Processing works is as shallow as the distributional interpretation of context.

Distributed representations are mostly intended to model similarity and analogy, such as in the work by Drozd et al. (2016). Such representations have been proven to be beneficial in lots of downstream tasks thanks to their geometric properties, such as investigating how the sense of a term varies over time (Kulkarni et al., 2015; Hu et al., 2019). However, distributional semantics tools can be hardly tailored to deal with issues such as predication, compositionality, lexical inferences, quantification and anaphora. Provided that the extent to which semantic properties can be reduced to combinatorial relations is in general controversial, some of the main issues have been individuated in a seminal Dagstuhl seminar (Kamp et al., 2014), and include: *i*) Polysemy and Vagueness, including type coercion, metonymy, metaphor and figurative language; *ii*) Inference and Reasoning, including deduction, induction, and abduction, defeasibility and default logics; *iii*) Compositionality, including function application, and semantic roles; *iv*) Modality and Negation, including Deontic logic, Epistemic logic and reasoning about knowledge and beliefs. While attempts have been made to reconcile the compositional and distributional perspectives (Grefenstette, 2013; Van de Cruys et al., 2013), and to understand their correlation extrinsically (Baan et al., 2019; Korrel et al., 2019; Chrupała and Alishahi, 2019), many of the aforementioned problems remain open — see for instance the work by Glockner et al. (2018), highlighting how modern neural systems fail to generalize over very simple lexical-based inference tasks.

To summarise, ontological analysis is a good candidate to provide formal tools to investigate context, possibly in its interplay with events and relations, complementing the family of approaches underlying distributional semantics. Among the aims of this special issue there is the intent to redraw the boundaries of these perspectives on context, and their interactions.

In 1993 the workshop “Using Knowledge in its Context” co-located with the International Conference on Artificial Intelligence (IJCAI) kick-started an interdisciplinary effort to investigate ways of modelling context and its effects in an attempt to bring together approaches, frameworks and theories scattered across disciplines, promoting cross-fertilization and interactions (Brezillon and Abu-Hakima, 1995). This initiative has then evolved into a series of conferences (“Conference on Modeling and Using Context”, CONTEXT, started in 1997 and still running) and into a new interdisciplinary journal, “Modeling and Using Context”. In 2017 the workshop “Contextual Representation of Events and Objects in Language” (CREOL) has been organized, aimed at promoting the investigation of context from the perspective of Natural Language Processing (NLP) through the use of ontologies, i.e., repositories of explicit knowledge about the world, and to highlight limits of current approaches in NLP. The contributions and discussions around CREOL, whose second edition was held in 2019, also in the frame of the Joint Ontol-

ogy Workshops¹, prepared the ground for this special issue, entitled “Meaning in Context: ontologically and linguistically motivated representations of objects and events”. Applied Ontology qualifies as the most appropriate venue to foster this discussion.

Four contributed articles were selected for inclusion in this volume. Besides technical quality, impact of the ideas, and presentation, these works were judged particularly fit to the themes of the special issue. All of them approach the idea of context from a particular perspective, either focusing on events or objects and other entities, thus providing complementary views on this complex notion.

In “What to consider about events: A survey on the ontology of occurrents”, Rodrigues and Abel (2019) present an in-depth survey of 11 ontologies available in literature, including some of the most popular upper ontologies (e.g., BFO, UFO, DOLCE, YAMATO, SUMO, GFO, among others). The contribution develops illustrating similarities and differences across them by focusing on four key aspects, or facets, of occurrents: their definition and the relationship with respect to participation, mereology, and causation. Interestingly, the latter three facets involve some degree of acknowledgement of context. Such facets are reviewed and critically assessed following three paths: uncontroversial aspects (i.e., aspects that are shared by all or the majority of the surveyed ontologies); complementary aspects (i.e., aspects that although not present in most ontologies do not present particular issues to be integrated or reconciled with commitments of other ontologies); and conflicting aspects (aspects that cannot be reconciled with commitments of some other ontology). It is well accepted across different ontological frameworks that occurrents are disjoint entities when compared to continuants, that they differentiate between representations of change and maintenance of states, and that have continuant participants. Participation qualifies as a central aspect of occurrents. There is a general agreement on the fact that participants may play different roles in the context of the occurrent in which they participate. At the same time, disputes (or contradictions) are still pending concerning the types (categories) of participants that can participate in an occurrent, or on the requirements of explicitly having participants for the happening of occurrents. It also appears undisputed that occurrents may have temporal parts. However, according to the views adopted in modelling a specific ontology, differences emerge on aspects such as the existence of instantaneous temporal parts. The last facet, causation, is also problematic: some ontologies do not explicitly include it (e.g. SEM, ESO and VEL) or use it to describe different aspects of the ontology itself (e.g. YAMATO). When explicitly modeled, there is agreement on a general notion of causation, while incompatible aspects concern the extension of causation to continuants, or the possibility of having bidirectional or simultaneous causations. Complementary aspects present interesting connections across frameworks that are not easy to identify. They may allow, however, extensions as well as enable interoperability of ontologies. Awareness of commonalities and (compatible and not compatible) differences across ontologies should force modelers, and also users, to take into account the consequences they may face and present when committing to specific aspects for occurrents.

In “Embodied Contextualization: Towards a Multistratal Ontological Treatment” Bateman et al. (2019) provide an investigation on the role of simulation in the contextualization of ontological knowledge. The authors propose a framework aimed at bridging the gap between natural language semantics and upper level ontologies, often highly underspecified with respect to contextual knowledge. The framework is instantiated to support reasoning for domestic service robots. Agents run simulations testing aspects such as affordances, in order to select the appropriate semantics to interpret orders in a given (but often implicit) context. This process allows the agent to assess the outcome of underspecified actions such as “bring plates”, in different situational contexts, such as stacking a set of plates, appropriate in the

¹<https://creol2019.di.unito.it/>

context of putting them away, vs. distributing them, appropriate when setting the table. The framework is supported by an ontological backbone comprising multiple layers (hence “multistratal”) for different levels of abstraction, and holds the promise for a deep and transparent, and thus explainable, connection between natural language and robotic action.

In “Context For Language Understanding by Intelligent Agents” McShane and Nirenburg (2019) illustrate a complete architecture aimed at providing language-endowed intelligent agents (LEIAs) with many kinds of context. This work mainly proposes a conceptual framework rather than an implemented system. In this perspective, artificial agents need two sorts of context. The first one, *horizontal* context, is intended to handle the processing of perceptual stimuli over time, in diachronic fashion. The second one, *vertical* context, is aimed at recalling knowledge elements along with their associated processing routines in a multistage process of analysis. Seven layers are proposed to deal with both fundamental and rather subtle linguistic phenomena. These include: Preprocessing and Syntactic Parsing, Syntactic Integration and Recovery, Basic Semantic Analysis, Reference: Initial Analysis, Extended Semantic Analysis, Situational Reasoning, and Using Large Corpora as Extended Context. These layers consider many relevant linguistic phenomena (such as, e.g., dealing with scalars, gapping, VP ellipsis), and the authors introduce a relevant principle: sometimes complex semantic analysis can be performed during basic semantic analysis (such as conventional metaphor, which is lexically recorded and stored in some knowledge base), whilst in other cases (such as, e.g., metonymies or non-conventional metaphors featured by copular construction), it requires multi-stage efforts. Similar multi-round processing efforts seem to be required to deal with coreferential phenomena, arranged into three main cases: textual coreference, coreference with elements that are not mentioned in the text, and cases where the referring expression is new in the discourse. With respect to the degree of ontological precision underlying the overall framework, no specific ontology is described, but the work rather postulates the availability of a large knowledge base (endowed with about ten thousand concepts and a lexicon with thirty thousand senses) implementing a slot-filling formalism.

In “Implicit Entity Linking in Tweets: an Ad-hoc Retrieval Approach” Hosseini et al. (2019) explore the Natural Language Processing task of identifying entities in text that are only indirectly mentioned, e.g., in the utterance “Angelina Jolie and her husband”, denoting Brad Pitt. In the proposed approach, explicitly mentioned entities are automatically linked to a knowledge base, and they provide the context for the disambiguation of the implicit entity. Therefore, in this work, the context of an entity is considered as fixed, and formalized as a function of a knowledge graph. Additionally, temporal context plays a role when the data are augmented in order to create training material for a supervised classifier, along with ontologically-driven targeted search. The different contexts of an entity thus emerge as a function of how the entity is searched (e.g., Angelina Jolie as an actress vs. as a philanthropist) and when.

We hope that this special issue will contribute to stimulate research from different communities in exploring the connections between context and different approaches to representations of meaning with an open-minded approach.

References

- Baan, J., Leible, J., Nikolaus, M., Rau, D., Ulmer, D., Baumgärtner, T., Hupkes, D., and Bruni, E. (2019). On the realization of compositionality in neural networks. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 127–137, Florence, Italy. Association for Computational Linguistics.
- Bach, K. (2004). Pragmatics and the philosophy of language. In L., H. and G., W., editors, *The Handbook of pragmatics*, chapter 21, pages 463–487. Balckwell.

- Baldauf, M., Dustdar, S., and Rosenberg, F. (2007). A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing*, 2(4):263–277.
- Barwise, J. and Perry, J. (1983). *Situations and attitudes*. Cambridge, MA and London: MIT Press.
- Bateman, J. A., Pomarlan, M., and Kazhoyan, G. (2019). Embodied contextualization: Towards a multistratal ontological treatment. *Applied Ontology*, 14(4).
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.
- Brezillon, P. and Abu-Hakima, S. (1995). Using knowledge in its context: Report on the ijcai-93 workshop. *AI magazine*, 16(1):87–87.
- Chrupała, G. and Alishahi, A. (2019). Correlating neural and symbolic representations of language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2952–2962, Florence, Italy. Association for Computational Linguistics.
- Drozdz, A., Gladkova, A., and Matsuoka, S. (2016). Word embeddings, analogies, and machine learning: Beyond king - man + woman = queen. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3519–3530, Osaka, Japan. The COLING 2016 Organizing Committee.
- Fillmore, C. J. (1976). Frame semantics and the nature of language. In *Annals of the New York Academy of Sciences: Conference on the origin and development of language and speech*, volume 280, pages 20–32.
- Fillmore, C. J. (1982). Linguistics in the morning calm. *Linguistics Society of Korea. Frame Semantics. Seou: Hanshin*.
- Firth, J. (1957). *Papers in linguistics*. London. Oxford UP.
- Gangemi, A., Alam, M., and Presutti, V. (2018). Amnestic forgery: An ontology of conceptual metaphors. In *Formal Ontology in Information Systems - Proceedings of the 10th International Conference, FOIS 2018, Cape Town, South Africa, 19-21 September 2018*, pages 159–172.
- Glockner, M., Shwartz, V., and Goldberg, Y. (2018). Breaking nli systems with sentences that require simple lexical inferences. In *The 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, Melbourne, Australia.
- Grefenstette, E. (2013). Towards a formal distributional semantics: Simulating logical calculi with tensors. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 1–10, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Gu, T., Wang, X. H., Pung, H. K., and Zhang, D. Q. (2004). An ontology-based context model in intelligent environments. In *Proceedings of communication networks and distributed systems modeling and simulation conference*, volume 2004, pages 270–275. San Diego, CA, USA.
- Guarino, N. (1997). Understanding, building and using ontologies. *International Journal of Human-Computer Studies*, 46(2-3):293–310.
- Guarino, N. and Guizzardi, G. (2015). “we need to discuss the relationship”: Revisiting relationships as modeling constructs. In *International Conference on Advanced Information Systems Engineering*, pages 279–294. Springer.
- Guarino, N. and Guizzardi, G. (2016). Relationships and events: towards a general theory of reification and truthmaking. In *Conference of the Italian Association for Artificial Intelligence*, pages 237–249. Springer.
- Heim, I. (1983). File change semantics and the familiarity theory of definiteness. *Semantics Critical Concepts in Linguistics*, pages 108–135.
- Hosseini, H., Nguyen, T. T., Wu, J., and Bagheri, E. (2019). Implicit entity linking in tweets: An ad-hoc retrieval approach. *Applied Ontology*, 14(4).
- Hu, R., Li, S., and Liang, S. (2019). Diachronic sense modeling with deep contextualized word embeddings: An ecological view. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 3899–3908.
- Kamp, H., Lenci, A., and Pustejovsky, J. (2014). Computational models of language meaning in context (dagstuhl seminar 13462). In *Dagstuhl Reports*, volume 3. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Kamp, H. and Reyle, U. (1993). *From logic to discourse*. Dordrecht: Kluwer.
- Korrel, K., Hupkes, D., Dankers, V., and Bruni, E. (2019). Transcoding compositionally: Using attention to find more generalizable solutions. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 1–11, Florence, Italy. Association for Computational Linguistics.
- Kulkarni, V., Al-Rfou, R., Perozzi, B., and Skiena, S. (2015). Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web*, pages 625–635. International World Wide Web Conferences Steering Committee.
- Landauer, T. and Dumais, S. (1997). A solution to plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2):211–240.
- Malinowski, B. (1922). Argonauts of the western pacific: an account of native enterprise and adventure in the archipelagoes of melanesian new guinea.
- Malinowski, B. (1935). *Coral Gardens and Their Magic: A Study of the Methods of Tilling the Soil and of Agricultural Rites in the Trobriand Islands. II, the Language of Magic and Gardening*. RoutledgeFalmer.
- McShane, M. and Nirenburg, S. (2019). Context for language understanding by intelligent agents. *Applied Ontology*, 14(4).

- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Morris, C. W. (1938). Foundations of the theory of signs. In *International encyclopedia of unified science*, pages 1–59. Chicago University Press.
- Richards, D. and Simoff, S. J. (2001). Design ontology in context—a situated cognition approach to conceptual modelling. *Artificial Intelligence in Engineering*, 15(2):121–136.
- Rodrigues, F. H. and Abel, M. (2019). What to consider about events: A survey on the ontology of occurrents. *Applied Ontology*, 14(4).
- Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language, and Information*, 7(1):3–19.
- Uschold, M. and Gruninger, M. (1996). Ontologies: Principles, methods and applications. *The knowledge engineering review*, 11(2):93–136.
- Van de Cruys, T., Poibeau, T., and Korhonen, A. (2013). A tensor-based factorization model of semantic compositionality. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1142–1151.
- Wittgenstein, L. (1953). Philosophical investigations. London, Basic Blackw.
- Zainol, Z. and Nakata, K. (2010). Generic context ontology modelling: A review and framework. In *2nd International Conference on Computer Technology and Development*, pages 126–130. IEEE.