Undefined 0 (2019) 1–0
IOS Press

Closing the Loop between Knowledge Patterns in Cognition and the Semantic Web

Aldo Gangemi *,
University of Bologna, ISTC-CNR, Rome, Italy
E-mail: aldo.gangemi@unibo.it

Abstract. We discuss currently open issues in the discovery and representation of knowledge patterns in computational processing of meaning, in order to improve interoperability and cognitive validity of web-based semantics. We present the current state of *knowledge patterns* (KP) in Knowledge Representation, the Semantic Web and Cognitive Sciences, focusing on an intensional abstraction of heterogeneous predicates as a formal foundation for KP.

Keywords: Relations, Roles, Frames, Knowledge Patterns, Cognitive Semantics, Semantic Interoperability

1. A Design Loop

After 20 years of Semantic Web, at least 60 years of attempts to build computational models of meaning, and 100 years from the publication of Ludwig Wittgenstein's *Tractatus Logico-Philosophicus* [65], let alone the previous footwork of philosophers, linguists, and logicians, the situation with publicly shared, rigorous representations of meaning is only partly satisfying. The deep learning turn in artificial intelligence is adding new means for inductive inference and pattern discovery, but not much to the general problem: what are the basic bricks of meaning, if any, and their viable computational representation? How to make them converge (or diverge) according to the needs for local efficacy and global interoperability?

In work presented in 2010 for the inaugural issue of this journal [26], those building blocks were identified in *Knowledge Patterns* (KP) [10][20], a semantic web generalisation of *frames* in cognitive science, linguistics, and sociology literature, which have played a substantial role in early knowledge representation. The proposed approach was to empirically collect and use KPs for design, re-engineering and interoperability across data, schemas, lexicons, and interaction.

*Corresponding author. E-mail: aldo.gangemi@unibo.it

While in 2010 the amount of known KPs was limited to certain well known ontology design patterns [20][50] and informal linguistic frames, with examples of how different data models and data structures could be made interoperable through them, from that time some advancements have been made, which are briefly summarized in Sect. 3.

It is now time to assess where we are, and to take another step towards an integration of scientific efforts from related disciplines ranging from cognitive neuroscience to knowledge representation.

2. KP as Relational Knowledge

As Dedre Gentner [30] stated in a crystalline way:

> the ability to perceive and use purely relational similarity is a major contributor –arguably the major contributor– to our species' remarkable mental powers.

Gentner's quotation gives us a starting point to propose a dual nature for knowledge patterns and relations: on one hand, they are intensional structures that represent certain invariant features of the world, making specific situations emerge out of the continuum of reality as perceived, memorised, and publicly recognised in human societies and individuals. On the other hand, they are relations with a precise extensional semantics. The extensionally ordered view of relations corresponds to the intensional view of knowledge patterns, and viceversa. Feature similarity, as processed by humans or machines in reasoning and learning, helps detecting a relation that has been already sensed in the past, as well as recognising a pattern that has been already stored in bodily, societal, or cultural memory.

The Semantic Web has started as a pragmatic way to use the Web as a platform to spread human semantics and human ability to process meaning. That platform was supposed to be decentralised, and to (unintentionally) realise the dream of a transparent negotiation of meaning, where entities have a public identity, with publicly known features that are encoded in public representations that can be dereferenced on the Web.

Eventually, the Semantic Web has created the conditions for web semantics to evolve: billions of multidomain Linked Data triples, the international acceptance of governmental linked open data, the F.A.I.R. data movement[1], and the crucial asset development for enterprise knowledge graphs, are all evidence for a paradigm shift. Yet, where real semantic interoperabil-

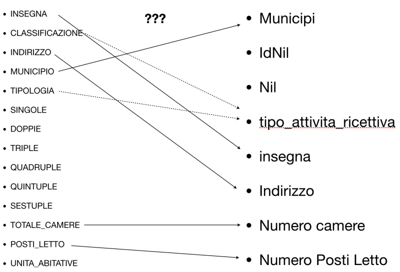

Fig. 1.: Incomplete mappings between accommodation schemas for Rome vs. Milan municipality data.

ity has succeeded, it has typically happened in a centralised way. Some examples are mentioned here:

- public administrations produce data with heterogeneous schemas, even for simple conceptualisations such as accommodations in Rome and Milan (Fig. 1): in most cases, it's centralised efforts to create shared schemas, and complex refactoring procedures after data ingestion, which alleviate the problem.[2] Exceptions such as naming hubs in sameas.cc, which are used by multiple distributed data providers, are not yet attacking the problem of relational meaning;
- web designers and content producers use their own tags, and only something like schema.org has enabled SEO and semantic search to take off;
- DBpedia has evolved a large schema for Wikipedia data that is partly dependent on Wikipedia Infoboxes, partly on collaborative design of classes and properties, however, data needs cleaning, and only a stronger semantics as shown in [47] is able to detect the most severe problems emerging from bulk reengineering practices. Active work is being done also on Wikidata [62], but catching regularities in its massively heterogeneous entities is still an elusive task;
- the decision making on sharing schemas is painful and subject to conflicts, let alone the cases when generic schemas, which are independent from an organisation's control, and do not necessarily cover the same semantics, are nonetheless assumed as standard. A proper practice may be instead to analyse the requirements extracted from scenarios or competency questions, as recommended by state-of-the-art agile methods such as eXtreme Design [51], and only later to align the resulting ontology to existing ones. The ArCo ontology network [9] demonstrates the advantages of this approach;
- a large amount of knowledge needs to be extracted from natural language, but the integration between natural language understanding, which is progressing towards shareable semantic representations such as AMR [4], and ontology design, is not yet widespread, despite the road has been opened by knowledge extraction methods [29] and massive integration of linguistic and factual resources [22].

Notwithstanding the long (15 years) activity of the Ontology Design Patterns community, with the substantial work collected in dedicated repositories,[3] or published (e.g. [35]), let alone the general agreement on reusing design patterns in ontology design for the Semantic Web and Conceptual Modelling [16], ontologies usually do not include their design practices (as possible e.g. with OPLA [36]), and modeling choices

[1] https://www.go-fair.org/fair-principles/

[2] Cf. the DAF platform supported by the OntoPiA ontology network in Italy `https://github.com/ontopia/ontopia`

[3] E.g. `http://www.ontologydesignpatterns.org`, `http://www.gong.manchester.ac.uk/odp/html/`, `https://github.com/INCATools/dead_simple_owl_design_patterns`

are scarcely documented, leading to difficulties in integrating schemas and their data. Even the FrameBase approach [54], which practically implements interoperability based on FrameNet frames as hubs for alternative schemas sharing a KP, does not yet look as a real game-changer, possibly due to its limited coverage.

Clearly, there is a difficulty in abstracting from local modeling choices, without a strong centralization, or a push towards the reuse of a quasi-standard. The reason probably lies in both the distance between domain expertise and ontology design practices, which require non-trivial logical competence, and in the existence of alternative terminologies, design solutions, and local alternatives, which make different but potentially overlapping schemas look farther than they actually are.

Semi-automated ways to design ontologies, to match them, or to inject interoperability, are still on the academic side of things, probably because of their limited friendliness, or conceptual coverage.

The suggestion here is to take the bull by the horns, which in this case means to accelerate the widespread collection of knowledge patterns *where they actually are*: existing ontologies, data models, large natural language corpora, linguistic resources, competency questions from formal and informal contexts, workflows, how-to repositories, commonsense knowledge bases; as well as *when they can be extracted at scale: automatic construction of knowledge bases, schema induction, language models*, etc. This is already happening (see Sect. 3), but at a pace and level of awareness that are too low to impact within a reasonable time. What is missing includes a good theory, and practical representation, collection, and reuse tools.

For example, much research for practical features and languages to facilitate ontology design, including e.g. OWL2 [40] *punning* (a.k.a. type reification [19]) and *keys* (a.k.a. *identification constraints* [8]), SPIN[4], ShEX[5], SHACL[6], OTTR[7], etc., mostly originates from the need to represent, evaluate, or respect implicit KPs, and to make existing ontologies and data satisfy KP cognitive requirements (Sect. 4).

The knowledge patterns emerging from this activity need to have both intensional and extensional representation. KPs should preserve their intensional nature, and avoid a strong commitment to specific logical primitives, in order to be robust against the evolution of knowledge representation, knowledge engineering, and alternative paradigms, including e.g. graph databases [53] and graph networks [7]).

However, a correspondence between KP intension and extension can be maintained by using a lightweight, pragmatic semantics, e.g. the Framester semantics, which is summarised in Sect. 6.

3. Where are we now?

What are the main research questions for KP research? A partial list is proposed here as a checklist for the next years.

1. what KPs are known?
2. is KP coverage enough to approximate human knowledge patterns?
3. how to extend, evolve, learn, or discover KPs?
4. how to enrich automated reasoning with an intensional characterisation of KPs?
5. how to use intensional KPs to foster interoperability independently from the local representation of an ontology or conceptual model? In other words, how to employ KPs in ontology reengineering and ontology matching?
6. what is the intensional difference between frames, roles, and selectional constraints or types?
7. how to formalize KP compositionality?
8. how to study higher levels of semantics, such as modalities, opinion, emotions, metaphors, narratives, and other macrostructures?

This is definitely an ambitious research programme, which has been partly carried out in the last 10 years.

Concerning known KPs and their coverage, some progress has been made, for example the Framester [22] knowledge graph is able to represent any linguistic or ontology predicate as a KP, and to reconcile it to a foundational layer initially provided by FrameNet [55] frames, and extended by aligning and incorporating predicates from multimodal linguistic, data, and multimodal resources. Hundreds of thousands of KPs have been automatically extracted from existing repositories [43],[22]. Many more KPs can be extracted from existing data [49], or informal graphs such as Wikipedia links [44]. Now state-of-the-art formal knowledge extraction tools such as FRED [29] (see also Sect. 6) are able to extract KP-based knowledge graphs from text, and aligning them to Framester-based predicates.

The ability to make use of KPs for interoperability has been proved e.g. by FrameBase [54]. Another ex-

[4] https://www.w3.org/Submission/spin-overview/
[5] https://www.w3.org/2013/ShEx/Primer
[6] https://www.w3.org/TR/shacl/
[7] http://ottr.xyz/

periment has been described in [1] about reconciling different but related knowledge graphs (extracted from text), by exploiting KP embeddings and combinatorial optimisation.

An example of using KPs from the DOLCE foundational ontology [52] to clean up a knowledge graph is described in [47]. Examples of using KPs for representing higher levels of meaning are described in [28] about using a two-tier semantics for extracting knowledge graphs from text, and over-describing them with opinion KPs that improve the state of the art in aspect-based sentiment analysis. Another recent example [23] is about representing conceptual metaphors as KP mappings, and attempting to use the resulting knowledge base for both detection and generation of metaphors.

Tooling is also important for supporting the adoption of KPs. Recommendations such as ShEX and SHACL, jointly with languages such as SPIN and OTTR, seem valuable commodities to that purpose, even though they do not directly indicate KPs as use cases. Named graphs are also an underexploited possibility for adding a "graph of KP-based graphs" to complex ontologies and data.

Whatever the language or recommendation, there seems to be emerging a need for "packaging" subsets of axioms from an ontology, which make sense as a unit for designing, unit testing, rapid prototyping, customising inferential procedures. A reasonable hypothesis is that such packaging is ultimately motivated by the intuitive necessity to make ontologies better correspond to the cognitive principles used to organise knowledge, a.k.a. knowledge patterns (Sect. 4).

Finally, some recent work is being carried out outside of the semantic web community, but it has a lot of relevance for KP research in discoverying patterns that can be easily converted into KPs. Some examples includes end-to-end neural frame detection systems (e.g. [59], and automatically constructed commonsense repositories from large scale data (e.g. Atomic [56]).

In the next sections we revisit the cognitive foundations of KP (Sect. 4), and the deeper problem of intensional compositionality (Sect. 5). Concerning KP semantics, in Sect. 6 we summarize a long-standing investigation into the nuances of intensional KP representation. How full-fledged reasoning with KP compositionality might impact existing automated reasoning techniques? Could we reduce the computational complexity of knowledge graphs and their matching by counting on the schematic nature of KPs, and automated translation into existing logical languages?

4. KP and Cognition

The term *knowledge pattern* was firstly introduced by de Beaugrande [12]:[8]

> the availability of global patterns of knowledge cuts down on non-determinacy enough to offset idiosyncratic bottom-up input that might otherwise be confusing.

However, the idea of recurrent, invariant units of knowledge was already present in philosophy, psychology and sociology as *schemata*, at least since [48]:

> La logique égocentrique est plus intuitive, plus «syncrétique», que déductive ... Elle emploie des schémas personnels d'analogie, souvenirs du raisonnement antérieur, qui dirigent le raisonnement ultérieur sans que cette influence soit explicite.[9]

Notably, in the same period (1970-1980) more notions were being introduced to characterize cognitive structures that were supposed to bridge research in linguistics, artificial intelligence, knowledge representation, etc. These include *Frames* in linguistics [17], later defined in FrameNet[10] as:

> a schematic representation of a situation involving various participants, props [inanimate entities, *ed.*] and other conceptual roles, each of which is a frame element

and in artificial intelligence [39], defined as:

> a remembered framework to be adapted to fit reality by changing details as necessary ... a frame is a data-structure for representing a stereotyped situation.

Macrostructures [61], defined as:

> higher-level semantic or conceptual structures that organise the 'local' microstructures of discourse, interaction, and their cognitive processing.

Scripts [58], defined as:

[8] A close usage of the term can be found earlier in a "creative engineering" book [2]: "*knowledge pattern ... by this is meant the knowledge and experience applicable to the technique of synthesis ... There are three important parts to the knowledge pattern as regards creative work, (1) scientific knowledge, (2) design curiosity, and (3) the ability to generalize experience.*"

[9] "Egocentric logic is more intuitive, more "syncretic", than deductive ... It uses personal patterns of analogy, memories of previous reasoning, which direct the subsequent reasoning without this influence being explicit."

[10] `https://framenet.icsi.berkeley.edu/fndrupal/glossary`

a structured representation describing a stereotyped sequence of events in a particular context.

As de Beaugrande noticed about those different notions, "*These large-scale knowledge configurations supply top-down input for a wide range of communicative and interactive tasks.*". In fact, there seems to be a common intuition concerning *invariances* shared by multiple situations, typically featuring an internal order, and being applied to multiple reasoning and interaction activities.

Since something can be invariant only if it remains unchanged under transformations[11] that span through time, space, observers, physical conditions, constituency, measurement, procedural constraints, etc.,[12] and since knowledge patterns are representations of situations, they reflect that those situations remain unchanged under some transformation of features that are not relevant for the pattern to be applicable (they "offset idiosyncratic bottom-up input").

For example, a red ball might still be a red ball after being deflated, but a red ball to play volley cannot. Throwing paper waste on the street may be the same action on any street, but in a country the same action can be tolerated, in another not. A slap is a slap, but it could be voluntary or not, an aggression or a joke, according to the intention of who's slapping, or to the observer's perspective.

Knowledge patterns contain invariant features that make them appropriate as abstract data structures to be remembered/stored, and, as Minsky [39] noticed about frames, they can be *adapted to fit reality by changing details as necessary*. Minsky's intuition can be used to propose KP dynamics as striking a balance between invariances (converging to universal patterns) and localities (tending to pattern divergence, adaptation or blending). Cognition works with patterns, but updates them to local observations, which are unique, because of the richness and compositional interference of actual (multi-modal) perceptions. This tension is reflected also in Barsalou's simulation theory [5], which can be summarised as the defense that concepts are *grounded* by multi-modally-informed, situated simulations of the external world. Barsalou [6] also proposes that concepts can be shared thanks to a huge *coordination activity* aimed at establishing a common ground for mutual understanding.

The balance suggested by Minsky, and re-proposed by Barsalou as social coordination, has some analogy in inconclusive results of neurological experiments that aimed at finding evidence for, or against, contextual dependency of core cognitive processes. For example, a recent fMRI metastudy by David Wisniewski [64] starts from the following dilemma:

> Some suggested that intentions representations in the fronto-parietal cortex change flexibly when external demands change (context-dependent coding). Others suggested that these representations are encoded in an abstract format that is not affected by changes in external demands (context-invariant coding)

It then revisits the literature on goal-oriented action and context, and finds that the stability/flexibility dynamics (which corresponds to the common-ground/adaptation dynamics of Minsky's and Barsalou's) is a motivation for inconclusiveness:

> results to date are mixed, showing context-dependence in some, but context-invariance in other cases ... depending on characteristics of intentions as well as environment, intentions can either be encoded in a context-dependent or a context-invariant format ... to achieve both stability and flexibility of behavior under constantly changing external demands

Two questions emerge then for a computational treatment of knowledge patterns: what features characterize a pattern? how to be tolerant to pattern adaptation?

On one hand, since patterns have inherent invariances, they are useful to make predictions, to create expectations, to quickly judge something, to catch opportunities (affordances), to avoid obstacles, to diagnose a medical condition, to hypothesize a natural law, to establish a social norm, to maintain a physical, social, or individual equilibrium, etc. This massive importance make them key to interoperability across multiple representations.

On the other hand, in many contexts a pattern can be used analogically, approximately, partially, while still retaining some of its explanatory power. In other words, patterns retain their usefulness even when they do not fully correspond to a situation.

An extreme case happens when a KP is used to denote the special or unique quality of a situation, e.g. when a politician has a lot in common with a sportsman, or a gangster (cf. the cases described in an exploration of knowledge patterns emerging out of Wikipedia links [45]), or when one recognizes the

[11]Cf. Paul Dirac [13]: *"The important things in the world appear as invariants ... of ... transformations"*.

[12]Cf. [42] for a detailed study on invariance and objectivity, and [31] for Gibson's psychological theory of how invariances in stimulus-energy pair permanent ("projectable") properties in the environment ("affordances").

unique way of nodding by a friend (uniqueness prizes "idiosyncratic bottom-up input").

In addition, due to their tolerance to modification, knowledge patterns have a dynamics: they are *adaptable* (e.g. when applying a `Too Much` frame to situations as different as food consumption, sunlight, or amusement), and can be *learned* or *discovered* by a human or a machine from a collection of examples.

Adaptability results in compositional problems that are easily interpreted by humans, but remain opaque even to sophisticated logical methods. For example, when we apply a `Too Much` frame to food consumption (e.g. *too much sugar*) vs. amusement (e.g. *too much fun* situations, knowledge patterns show a peculiar compositionality, which we discuss in Sect. 5).

5. KP Compositionality

Knowledge Patterns, variously called schemas, frames, scripts, scenes, modeling components, data modeling patterns, etc., have been proposed as the core building blocks in ontology design [26], providing cognitive relevance, explicit situation boundaries, independence from a particular formalism, under the assumption of direct associations to modeling requirements. For example, in the classical blocks world example of AI, a generic $\texttt{Over}(o_1, o_2)$ frame involving a vertical spatial relation between any two physical objects satisfies a modeling requirement that only takes into account the relative position of the objects.[13] However, if the requirements include the knowledge whether the two objects touch each other or not, a richer $\texttt{On/Above}(o_1, o_2, c)$ frame that requires a role for the contact situation will be needed. The richer frame is actually the composition of the $\texttt{Over}(o_1, o_2)$ and $\texttt{Contact}(o_1, o_2, c)$ frames.

A KP can be represented in a specific logical language, but it should also preserve an intensional representation that is invariant across logical languages. In their original presentation of KPs in knowledge representation, Peter Clark and colleagues [10] indicated category theory as the most adequate abstraction for KP representation. More recently, Oliver Kutz and colleagues [14] chose a close approach for the representation of conceptual blending. Our intention here has been to start from a more traditional mathematical framework, close to existing KR languages: a two-tier intensional/extensional logic, which can use the same basic semantic web languages or knowledge graphs in use today (see Sect. 6).

For example, a `Playing Music` KP (represented here as a first-order predicate) $\texttt{PM}(p, i, c, t, tim, loc)$, with role projections (Sect. 6) such as `player`, `instrument`, `composition`, `tempo`, `time`, `location`, etc. Semantic types (cf. Sect. 6) [63] could be added to those roles, e.g. a player should be a `person`, an instrument should be `tempered`, a composition should be in `written form`, a tempo should be in a certain `range`, etc. However, specific applications of `Playing Music` might force roles to accept an untempered musical instrument, an AI playing a part, a section that is not written, but improvised, etc.

`Playing Music` could also be used to refer to a metonymically related situation, e.g. when one plays music *on* an audio system: in this case interpretation needs to reconstruct a composition of default music playing, its recording, and its reproduction. We may want to treat this as two separate `Playing Music-1` and `Playing Music-2` frames, but `Playing Music-2` is the result of a composition depending on `Playing Music-1`.

Literature on compositionality is huge (cf. [34] for a recent palette of positions), but the basic argument is about the asymmetry between symbolic and semantic compositionality:[14] *is the meaning of a structure entirely determined by the meaning of its constituents*? There are multiple reasons why the answer is "not always". We can consider several classes of asymmetry between purely symbolic and semantic compositions, here examplified with natural language cases:

- anaphoric composition: *They got* ***married. She*** *is beautiful*;
- modal composition: *23-year-old man dies after* ***fake doctor*** *administered unidentified treatment via injection*;
- hidden relations: *this* ***plaid jacket with hood*** *is made of cotton*;

[13]A reviewer wondered if relations like `Over` are "just" predicates, while KPs should be "configurations of interrelated predicates". As argued later in Sect. 6, a KP is an intensional (reified) view of a multigrade predicate, i.e. a predicate with arbitrary arity. Hence, all predicates can be either elementary or composed KPs. For example, `Over` may be used with more arguments for time, spatial context, amount of touching between the objects, etc.

[14]Following common practice in knowledge representation, by *semantic* compositionality we mean that composed symbols are interpreted with respect to e.g. a model-theoretic semantics that grounds symbols into an intended world.

- world structure: ***cutting a cake*** vs. ***cutting the grass***;
- metaphoric composition: *Breaking point: why the Kyrgyz* ***lost their patience***.

In all those cases, entailment, perspective, background or commonsense knowledge, or blending [14][23], need to be supplemented in order to finalise semantic composition.

In practice, those cases are more or less easily understood by people, despite their asymmetry: what is lacking to computational semantics to approximate that ability?

Currently, we have sophisticated logical compositionality within ontologies: classes are associated with other classes via properties or taxonomical relations, properties are associated with other properties through chains, SWRL, or SPIN rules, classes are associated to properties via domains and ranges, or restrictions. We have an ontology compositionality via ontology import. We even have vectorial compositionality in vector space models of semantics [60], now enriched by deep learning techniques.

But we do not have a straightforward compositional machinery, let alone an algebra, to compose knowledge patterns. We can represent KPs in ontology modules (or alternatively in named graphs), and import them in a new ontology, or merging them into a graph. We could use intersection of predicates for e.g. an *A cat is on the mat* situation, predicating both `Over` and `Contact` to it. But how to establish which role of `Over` maps to a respective role of `Contact`? We may use "layering" (Sect. 6) by reifying the ordering of roles in the two relations, as well as their mappings. But this is not straightforward. We may at least empirically study intersections of predicates in existing ontologies, and check what properties are shared, and if the potential composition (provided it is logically coherent) makes sense cognitively.

In fact, how to establish whether the result of composing a KP with another is a third KP? Our proposal is that we need a language to talk about intensional compositionality, jointly with a grounding into ontologies and off-the-shelf classes and properties. A beginning of such a composition style is demonstrated in Sect. 6. We exemplify here in more detail how KP compositionality provides a different view on well-known problems in natural language semantics and ontology engineering.

Framality Some compositionality effects on formal representations of meaning seem to derive from *framality* i.e. the hypothesis that KP (a.k.a. frames) are one of the motivating forces of contextual meaning. The hypothesis is supported by linguistic data, but also by neuropsychological studies (e.g. [63]) that report framal effects on selectional restrictions.

An example of framality can be given in adjectival semantics [25] as *the ability of an expression to evoke a KP from the joint evocation of KPs emerging during interpretation*. In the case of adjectives, a good example is the following pair of terms: *Extroverted Surgeon* vs. *Skillful Surgeon*. We might represent the two terms as a conjunction of predications, but while we can safely infer that all extroverted surgeons are extroverted in general, we are not safe at inferring that skilful surgeons are skilful in general. The likely reason is that `Being_skilled` is a possible value for the core aspects of the `Medical_professionals` KP, while `Being_extroverted` is not, therefore this tends to be interpreted as a frame composition. For comparison, a similar treatment for *extroverted comedian* does not allow a safe inference of being extroverted in general, while *alcoholic comedian* does.

Another example of framality can be done with reference to meta-properties proposed by the OntoClean methodology [33]. E.g. a property is traditionally called *rigid* when it is true for an entity during the entire course of its life, as with the `Student` property (in the sense of being enrolled at some educational institution) can hardly be true during the entire life of a person. However, this distinction is usually understood without taking into account locality conditions. For example, if an ontology is not interested in representing properties of entities in a foreverlasting perspective –as with a university enrolment ontology– what establishes rigidity is the temporal perspective of e.g. `Being_a_Student` frame, rather than the `Being_a_Person` frame. Within the university context, it is a property like `Enrolled_in_a_course` that is non-rigid, since the frame of that property has a shorter time span compared to that of `Being_a_Student`.

In other words, the context of meta-level properties is maximal, while framality requires contexts to be bound to requirements or local conditions. Interoperability requirements may change this sanity assumption: if university data are integrated with personal data, `Being_a_Student` would become non-rigid. Anyway, this may also apply to `Being_a_Person` if personal data are integrated with notarial data, which may include actions of a person even after death.[15].

6. A Framester Semantics for KP

Following from previous sections analysis of knowledge patterns, we provide here a more precise notion of KP as a reified relation that can be used to homogeneously represent *relations* as they can be evoked by natural language terms, logical constants, data modelling entities, informal terms such as those found in web formats: XML stylesheets, templates, microdata, infoboxes, JSON objects, etc. as well as "concept norms" as used in cognitive neuroscience [38], and "features" as used in machine learning. KP need to be independent from a particular representation or notation; in practice, we are talking about relational *senses* (Frege's *Sinn* [18]) of symbols, i.e. their intension. We will defend the intuition that a relational sense has no fixed arity, and we will discuss how KPs as intensional relations are approximated extensionally in OWL implementations.

The notion is presented here in a succinct formal notation, but it is equally implemented in OWL2 as a schema[16] for the Framester [22] factual-linguistic data hub. Framester is used here as evidence of a pragmatic and rigorous way of obtaining semantic interoperability at the schema level, across heterogeneous knowledge.

We start with defining *multigrade predicates* [46][57]. A multigrade predicate denotes a polyadic relation. The notion was firstly introduced by Leonard and Goodman [37]:

> a relation without any fixed degree may be called a "multigrade relation" (p. 50).[17]

This notion is nowadays scarcely known, but it is *de facto* used everywhere in formal linguistics and knowledge representation, following the neo-Davidsonian approach [11] of reifying event relations as individuals. In practice, most relations intuitively admit a non-fixed amount of arguments in their signature, e.g. *preparing a coffee* may express its maker, the coffee produced, the mix used, the machine employed, time, location, method, etc. It is in fact quite obvious that people tend to use a unique predicate for the same relation, taking the complexity of its context for granted. Trying to make all arguments explicit may even have deleterious effects e.g. in conversation (verbosity), in data modelling (local irrelevance), or machine learning (non-existent information for the expected feature).

Multigrade predicates have a signature including *argument labels*. Some arguments can be optional, and even implicit. In other words, a signature can be *expanded* once new knowledge becomes available because of some inferential process, or knowledge evolution. For example, given the `Over(x,y,t,...,s)` relation with the meaning of an object standing over another, we may expect that two of its arguments (`x` and `y`, with labels *above object* and *below object*) are not optional. We may also expect that other arguments are optional (e.g. a *time* `t`), and even implicit (e.g. the amount of *shadow* `s` projected by the above object onto the below object).

In other words, the intensional signature of a KP is ideally open, so that its implementations would typically result as approximations as per the *Open World Assumption* –values can be missing– and *Contextuality*: arguments can be different in different contexts.

Approximation obeys pragmatic principles: availability of knowledge, local/evolving requirements, etc. For example, `Over` could be implemented with different signatures in different ontologies, data models, machine learning features, etc., based on local requirements.

When ontology matching is applied in order to enable interoperability, signatures need to be morphed into one another, and interpretation **incompatibility** or **expansion** may arise. For example, given five ontologies A, B, C, D, E with a notion of `Over`, in A time might be instantaneous, while in B is based on intervals; in A shadow has a relevance, while in C does not care, but has a different argument considering magnetic forces between the objects; in D objects can be only of a certain size or type; in E objects can be non-physical, leading to metaphorical interpretations, etc.

As a consequence, semantic interoperability across heterogeneous data needs abstraction, which has to rely e.g. on a general intension of `Over`, independently from how it is logically represented (e.g. OWL class, property, individual, restriction). Unfortunately, current techniques for ontology matching do not address this problem in general, with the exception of those that do take into account KPs, as with Frame-

[15] In that case, the legal validity of `Being_a_Legal_Person` persists beyond the physical persistence of `Being_a_Person`.
[16] `https://w3id.org/framester/schema/`
[17] A predicate is the name for a relation that represents actual situations, but in the literature the two terms are often used interchangeably.

Base [54], and Framester [22]. The latters do not entirely solve the task of large-scale interoperability with the detail that e.g. Formal Ontology has been dreaming of since 25 years ago [32], but at least they take a step into that direction, at Web scale.

KP as multigrade predicates can be easily represented in OWL2 by using *punning* [40], and a vocabulary to talk about their *argument places* (a.k.a. *roles*), and the *types* of things denoted by argument values.

This vocabulary already exists since 2003, as a knowledge pattern framework called Descriptions and Situations (D&S) [24,21,27]. D&S was originally intended as a two-tier modelling of the extensional and intensional semantics of predicates, with a focus on events or situations. The motivating use cases were in legal and medical ontologies, where we need to talk both about the world (e.g. organic or social facts), and about the way we observe or categorize it (e.g. a clinical condition or a legal norm).

In Framester D&S-inspired semantics, a KP is defined as a multigrade predicate $\phi(e, x_1, ..., x_n)$, where ϕ is a first-order relation, e is a situation described by a KP, and x_i is a variable for any argument place. Now, using D&S-style OWL2 punning, ϕ is both an individual denoting a knowledge pattern from the class `KP` (axiom 1), i.e. its intension, and a subclass of the class `SIT` (axiom 2) of situations occurred or observed by using that KP, i.e. its extension. For example, `PreparingCoffee` is both a class of situations, and an individual belonging to the class of knowledge patterns. Once punning is established for multigrade predicates, Framester semantics introduces punning for arguments and types. Axioms 4-6 introduce ρ, the class of reified projections of multigrade predicates, which was originally employed [43] to formalise FrameNet [3] conceptual frames.

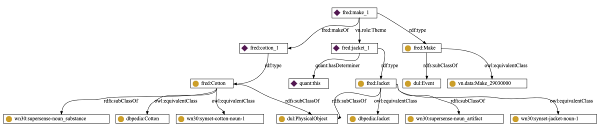

Fig. 2.: An automatically generated knowledge graph based on Framester semantics.

Axiom 3 is for *semantic roles* ρ^r, expressing binary relations between the reified predicate, and one of its reified arguments, e.g. between `PreparingCoffee` the `CoffeeMix`.

Axiom 4 is for *co-participation relations* ρ^c, expressing binary relations between between any two arguments of a same predicate, e.g. a `prepares` relation holding between a `CoffeeBrewer` and a `Coffee`.

Axiom 5 is for *types* ρ^t, expressing unary relations for the type of things that are used as values for an argument, e.g. `Agent` for `CoffeeBrewer`.

$$\forall(\phi)\mathtt{KP}(\phi) \tag{1}$$

$$\forall(s)\phi(s) \leftrightarrow \mathtt{SIT}(s) \tag{2}$$

$$\forall(s, x_i)\rho^r(s, x_i) \rightarrow \phi(s, x_1, ..., x_n) \tag{3}$$

$$\forall(s, x_j, x_k)\rho^c(x_j, x_k) \rightarrow \phi(s, x_1, ..., x_n) \tag{4}$$

$$\forall(s, x_m)\rho^t(x_m) \rightarrow \phi(s, x_1, ..., x_n) \tag{5}$$

$$(i,j,k,m \geq 1 \leq n)$$

$$\mathtt{Clothing} \in \mathtt{KP} \tag{6}$$

$$\mathtt{Causation} \in \mathtt{KP} \tag{7}$$

$$\mathtt{Jacket.n.1} \sqsubseteq \mathtt{Clothing} \tag{8}$$

$$\mathtt{Make\_26010000} \sqsubseteq \mathtt{Causation} \tag{9}$$

$$\mathtt{Cotton.n.1} \sqsubseteq \mathtt{Substance} \tag{10}$$

$$\mathtt{Make\_26010000.Theme} \in \rho^r \tag{11}$$

$$\mathtt{Make\_26010000.Theme.Material} \in \rho^c \tag{12}$$

$$\mathtt{Cotton.n.1} \in \rho^t \tag{13}$$

$$\mathtt{J.M.C.} \equiv \tag{14}$$

$$(\mathtt{Jacket.n.1} \otimes \mathtt{Make\_26010000} \otimes \mathtt{Cotton.n.1})$$

A simple example of Framester KP semantics is provided in Axioms 6-14 with a description logic representation of KPs and a sample of their projections evoked by the sentence: *this jacket is made of cotton*[18]. Framester semantics is also adopted in the knowledge graphs extracted by FRED [29],[19] see also Fig. 2.

An advantage of this intensional generalisation is its neutrality: no special knowledge representation language is required. E.g. a *make* concept may be represented in different ways: an object property from an ontology, a datatype property from a database refactoring, a class from another ontology, an individual from a linguistic ontology, etc. In the KP view to in-

[18]The used predicates are all from the Framester knowledge graph, which can be downloaded and queried from `https://github.com/framester/Framester`

[19]Use this API for an example: `http://wit.istc.cnr.it/stlab-tools/fred/demo`

teroperability, they can be all aligned using both ontology and linguistic matching techniques, once the intensional disambiguation has been performed.

Based on the two-tier semantics sketched here, we exemplify the formal problem of intensional compositionality (Sect. 5) with the sentence *this jacket is made of cotton*, which evokes a KP (`J.M.C.`) that is composed as in axiom 14 by using an $\otimes$ operator.

However, the formal semantics for the $\otimes$ operator needs to be investigated, and a comfortable solution for its extensional representation is yet to come. As anticipated in Sect. 5, a simple class intersection may fail to catch the actual situation. The actual situation is in fact a context to jointly interpret a piece of clothing, a causation event, and some substance. Representing the situation as an instance of an intersection of `Clothing`, `Causation` and `Substance` may be superficially valid, but fails to catch the underlying cognitive pattern reflected in the ordering of the arguments for the three composed KP.

7. Discussion

We have presented a blended essay-position paper on the current state of knowledge pattern (KP) research. We have proposed an intensional abstraction for knowledge patterns (KP) on top of heterogeneous representation formats and languages. This is needed at a time when data is an essential fuel of society, but semantics is still a hodgepodge of approaches from different communities, and multiple computational and reasoning paradigms co-exist without much interaction, as with symbolic vs. sub-symbolic methods, scruffy vs. neat attitudes in modeling, proprietary vs. open representations, etc.

While the Semantic Web (SW) has induced a substantial leap forward in bringing a knowledge-level paradigm to computer science, its modelling practices have barely touched the problem of cognitive meaning processing. Cognition heavily relies on intensional schemas, and uses sensory data/entities to activate, and possibly adapting, them in order to make sense of the environment, store memories, resolve a problem, etc. (cf. Sect. 4).

In the SW, ontology evolution and matching are the main areas where *semantic interoperability* is addressed, and can be seen as the empirical counterpart to Minsky's and Barsalou's (cf. Sect. 4) notions of *adaptation* and *need for coordination* as forces for human knowledge sharing.

However, those authors showed us that adaptation and coordination are not casually emerging, but happen on top of stable patterns. The research hypothesis of ODP is that knowledge patterns can be used (once discovered and represented) as heuristics to coordinate evolution and matching, aiming at a simplicity that naturally matches human cognitive interpretation of the world. The fairly large literature on KP, and their related adoption, prove that the ambition is justified, even though discovering, representing, and sharing KP is less simple than expected. Firstly, because of the said dynamics of evolution/stability; secondly, because social, economic, and political reasons often work against it.

As an example, the current practices of accumulating data, and their automated analysis without an explicit shareable knowledge accessible to final users of semantically-enhanced systems, as argued in [15], hints at general policies that only partly work for coordination, focusing instead on local adaptations motivated e.g. by revenue optimisation.

The proposal for an intensional abstraction is therefore a call towards cognitive transparency in society: stable patterns are made available for the sake of interoperability across arbitrary knowledge organisation systems, which are often far from rigorous logical design, and their representation requires attention to how humans *compose* meaning It is a responsibility of the Semantic Web community to find scientific solutions and incentives to foster attention to those human factors in the adoption of semantic technologies.

Far from being complete, we have made an assessment of the state of play about knowledge patterns in the Semantic Web (Sect. 2-3), jointly with a review of why cognitive science can give us directions to design, extract, and use data and ontologies in the most efficient way when the problem is interoperability in computational treatment of meaning (Sect. 4).

We have presented a simple, layered semantics for KP as *multigrade predicates* (Sect. 6), as currently implemented in the Framester data hub, and touched research issues concerning KP compositionality (Sect. 5).

The variety of situation types addressed by existing ontologies, let alone the larger societal scenarios, require a stronger push towards cross-disciplinary systematization of KP research. This can be helped by more semantics (i.e. more detailed axiomatisation in ontologies), but it requires the availability of a large-scale repository of cognitive patterns that enable complex similarity reasoning (KP matching, (un)predictable KP evolution) among ontologies.

Given the amount of work carried out in multiple communities in order to gather KP or KP-like data, we can conclude that the situation is not so bad. The difficulty of knowing our own ways of representing knowledge, and the effects that emerge out of a continuous activity of data production and coordination, constitute a formidable challenge.

We assume a positive lookout for the next years, when we may be able to make currently unrelated approaches converge, eventually gathering wider coverage and more detailed understanding of how KP emerge, change, and compose in our computationally-enriched societies.

References

[1] M. Alam, D. R. Recupero, M. Mongiovì, A. Gangemi, and P. Ristoski. Event-based knowledge reconciliation using frame embeddings and frame similarity. *Knowl.-Based Syst.*, 135:192–203, 2017. 10.1016/j.knosys.2017.08.014.

[2] American Society of Mechanical Engineers. *Creative Engineering: A Series of Discussions at ASME Meetings Emphasizing the Importance of Ingenuity, Intuition, and Creative Ability in the Engineering Profession: Suggesting how These Qualities Might be Encouraged and Developed*. 1956.

[3] C. F. Baker, C. J. Fillmore, and J. B. Lowe. The Berkeley FrameNet project. In C. Boitet and P. Whitelock, editors, *Proceedings of the Thirty-Sixth Annual Meeting of the Association for Computational Linguistics and Seventeenth International Conference on Computational Linguistics*, pages 86–90, San Francisco, California, 1998. Morgan Kaufmann Publishers.

[4] L. Banarescu, C. Bonial, S. Cai, M. Georgescu, K. Griffitt, U. Hermjakob, K. Knight, P. Koehn, M. Palmer, and N. Schneider. Abstract meaning representation for sembanking, 2013.

[5] L. W. Barsalou. Perceptual symbol systems. *Behavioral and Brain Sciences*, 22:577–609, 1999.

[6] L. W. Barsalou. Cognitively plausible theories of concept composition. In *Compositionality and concepts in linguistics and psychology*, pages 9–30. Springer, Cham, 2017. 10.1007/978-3-319-45977-6_2.

[7] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.

[8] D. Calvanese, G. De Giacomo, and M. Lenzerini. Identification constraints and functional dependencies in description logics. In *IJCAI*, pages 155–160. Citeseer, 2001.

[9] V. A. Carriero, A. Gangemi, M. L. Mancinelli, L. Marinucci, A. G. Nuzzolese, V. Presutti, and C. Veninata. ArCo: the Italian Cultural Heritage Knowledge Graph. In *Proceedings of the International Semantic Web Conference, ISWC2019*, 2019. 10.1007/978-3-030-30796-7_3.

[10] P. Clark, J. Thompson, and B. Porter. Knowledge Patterns. In A. G. Cohn, F. Giunchiglia, and B. Selman, editors, *KR2000: Principles of Knowledge Representation and Reasoning*, pages 591–600, San Francisco, 2000. Morgan Kaufmann.

[11] D. Davidson. The Logical Form of Action Sentences. In *The Logic of Decision and Action*. University of Pittsburgh Press, Pittsburgh, 2nd edition, 1967.

[12] R. de Beaugrande. *Text, discourse, and process: toward a multidisciplinary science of texts*, volume 4 of *Advances in discourse processes*. Ablex Pub. Corp., 1982. 10.2307/1772216.

[13] P. A. M. Dirac. *Die Prinzipien der Quantenmechanik. (German) [The Principles of Quantum Mechanics]*. Verlag von S. Hirzel, Leipzig, Germany, 1930. Translated by Werner Bloch from the English original.

[14] M. Eppe, E. Maclean, R. Confalonieri, O. Kutz, M. Schorlemmer, E. Plaza, and K.-U. Kühnberger. A computational framework for conceptual blending. *Artificial Intelligence*, 256:105–129, 2018. 10.1016/j.artint.2017.11.005.

[15] J. Euzenat. A map without a legend: The semantic web and knowledge evolution. *Semantic Web*, to appear.

[16] R. d. A. Falbo, G. Guizzardi, A. Gangemi, and V. Presutti. Ontology patterns: clarifying concepts and terminology. In *Workshop on Ontology and Semantic Web Patterns*, 2013.

[17] C. J. Fillmore. *Frame semantics*, pages 111–137. Hanshin Publishing Co., Seoul, South Korea, 1982.

[18] G. Frege. Über sinn und bedeutung("On Sense and Denotation"). *Zeitschrift für Philosophie und philosophische Kritik, C*, 1892.

[19] A. Galton. Time and Change for AI. In *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 4*, pages 175–240. Clarendon Press, Oxford, 1995.

[20] A. Gangemi. Ontology Design Patterns for Semantic Web Content. In *M. Musen et al. (eds.): Proceedings of the Fourth International Semantic Web Conference, Galway, Ireland*, Berlin, 2005. Springer. 10.1007/11574620_21.

[21] A. Gangemi. Norms and Plans as Unification Criteria for Social Collectives. *Journal of Autonomous Agents and Multi-Agent Systems*, 16(3), 2008.

[22] A. Gangemi, M. Alam, L. Asprino, V. Presutti, and D. R. Recupero. Framester: A wide coverage linguistic linked data hub. In *Knowledge Engineering and Knowledge Management - 20th International Conference, EKAW 2016*, pages 239–254, 2016. 10.1007/978-3-319-49004-5_16.

[23] A. Gangemi, M. Alam, and V. Presutti. 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*, pages 159–172, 2018. 10.3233/978-1-61499-910-2-159.

[24] A. Gangemi, N. Guarino, C. Masolo, and A. Oltramari. Sweetening WORDNET with DOLCE. *AI Magazine*, 24(3):13–24, 2003.

[25] A. Gangemi, A. G. Nuzzolese, V. Presutti, and D. R. Recupero. Adjective semantics in open knowledge extraction. In *Formal Ontology in Information Systems - Proceedings of the 9th International Conference, FOIS 2016*, pages 167–180, 2016. 10.3233/978-1-61499-660-6-167.

[26] A. Gangemi and V. Presutti. Towards a pattern science for the semantic web. *Semantic Web*, 1(1-2):61–68, 2010. 10.3233/SW-2010-0020.

[27] A. Gangemi and V. Presutti. Multi-layered n-ary patterns. In Hitzler et al. [35], pages 105–131. 10.3233/978-1-61499-676-7-105.

[28] A. Gangemi, V. Presutti, and D. R. Recupero. Frame-based detection of opinion holders and topics: A model and a tool. *IEEE Comp. Int. Mag.*, 9(1):20–30, 2014.

10.1109/MCI.2013.2291688.

[29] A. Gangemi, V. Presutti, D. R. Recupero, A. G. Nuzzolese, F. Draicchio, and M. Mongiovi. Semantic Web Machine Reading with FRED. *Sermantic Web*, 2016. 10.3233/SW-160240.

[30] D. Gentner and L. Smith. *Analogical learning and reasoning*, pages 668–681. Oxford University Press, New York, 2013. 10.1093/oxfordhb/9780195376746.013.0042.

[31] J. J. Gibson. *The ecological approach to visual perception*. Laurence Erlbaum, Hillsdale, N.J., 1979.

[32] N. Guarino. The ontological level: Revisiting 30 years of knowledge representation. In *Conceptual modeling: Foundations and applications*, pages 52–67. Springer, 2009.

[33] N. Guarino and C. A. Welty. Evaluating ontological decisions with OntoClean. *Communications of the ACM*, 45(2):61–65, 2002.

[34] J. A. Hampton and Y. Winter. *Compositionality and concepts in linguistics and psychology*, volume 3. Springer, 2017. 10.1007/978-3-319-45977-6.

[35] P. Hitzler, A. Gangemi, K. Janowicz, A. Krisnadhi, and V. Presutti, editors. *Ontology Engineering with Ontology Design Patterns - Foundations and Applications*, volume 25 of *Studies on the Semantic Web*. IOS Press, 2016.

[36] P. Hitzler, A. Gangemi, K. Janowicz, A. A. Krisnadhi, and V. Presutti. Towards a simple but useful ontology design pattern representation language. In *WOP@ ISWC*, 2017.

[37] H. S. Leonard and N. Goodman. The calculus of individuals and its uses 1. *The journal of symbolic logic*, 5(2):45–55, 1940.

[38] D. Lynott and L. Connell. Modality exclusivity norms for 400 nouns: The relationship between perceptual experience and surface word form. *Behavior research methods*, 45(2):516–526, 2013. 10.3758/s13428-012-0267-0.

[39] M. Minsky. A Framework for Representing Knowledge. In P. Winston, editor, *The Psychology of Computer Vision*. McGraw-Hill, 1975.

[40] B. Motik, B. Parsia, and P. F. Patel-Schneider. OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax. W3C recommendation, W3C, October 2009.

[41] M. A. Musen and Ó. Corcho, editors. *Proceedings of the 6th International Conference on Knowledge Capture (K-CAP 2011), June 26-29, 2011, Banff, Alberta, Canada*. ACM, 2011.

[42] R. Nozick. *Invariances*. The Belknap Press at Harvard University Press, Cambridge, Mass., 2001.

[43] A. G. Nuzzolese, A. Gangemi, and V. Presutti. Gathering lexical linked data and knowledge patterns from framenet. In Musen and Corcho [41], pages 41–48. 10.1145/1999676.1999685.

[44] A. G. Nuzzolese, A. Gangemi, V. Presutti, and P. Ciancarini. Encyclopedic knowledge patterns from wikipedia links. In *ISWC 2011 - 10th International Semantic Web Conference*, pages 520–536, 2011. 10.1007/978-3-642-25073-6_33.

[45] A. G. Nuzzolese, V. Presutti, A. Gangemi, A. Musetti, and P. Ciancarini. Aemoo: exploring knowledge on the web. In *Web Science 2013 (co-located with ECRC), WebSci '13, Paris, France, May 2-4, 2013*, pages 272–275, 2013. 10.1145/2464464.2464519.

[46] A. Oliver and T. Smiley. Multigrade predicates. *Mind*, 113(452):609–681, 2004.

[47] H. Paulheim and A. Gangemi. Serving DBpedia with DOLCE - more than just adding a cherry on top. In *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part I*, pages 180–196, 2015. 10.1007/978-3-319-25007-6_11.

[48] J. Piaget. *Le langage et la pensée chez l'enfant*. Delachaux et Niestlé, 1923.

[49] V. Presutti, L. Aroyo, A. Adamou, B. A. C. Schopman, A. Gangemi, and G. Schreiber. Extracting core knowledge from linked data. In *Proceedings of the Second International Workshop on Consuming Linked Data (COLD2011), Bonn, Germany, October 23, 2011*, 2011.

[50] V. Presutti, E. Blomqvist, E. Daga, and A. Gangemi. Pattern-based ontology design. In *Ontology Engineering in a Networked World.*, pages 35–64. 2012. 10.1007/978-3-642-24794-1_3.

[51] V. Presutti, E. Daga, A. Gangemi, and E. Blomqvist. eXtreme design with content ontology design patterns. *Proc. Workshop on Ontology Patterns, Washington, DC, USA*, 2009.

[52] V. Presutti and A. Gangemi. Dolce+d&s ultralite and its main ontology design patterns. In Hitzler et al. [35], pages 81–103. 10.3233/978-1-61499-676-7-81.

[53] I. Robinson, J. Webber, and E. Eifrem. *Graph databases*. O'Reilly Media, Inc., 2013.

[54] J. Rouces, G. de Melo, and K. Hose. Framebase: Representing n-ary relations using semantic frames. In F. Gandon, M. Sabou, H. Sack, C. d'Amato, P. Cudré-Mauroux, and A. Zimmermann, editors, *12th European Semantic Web Conference, ESWC*, volume 9088 of *Lecture Notes in Computer Science*, pages 505–521. Springer, 2015. 10.1007/978-3-319-18818-8_31.

[55] J. Ruppenhofer, M. Ellsworth, M. Petruck, C. Johnson, and J. Scheffczyk. *FrameNet II: Extended Theory and Practice*. http://framenet.icsi.berkeley.edu/book/book.pdf, 2006.

[56] M. Sap, R. LeBras, E. Allaway, C. Bhagavatula, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, and Y. Choi. Atomic: An atlas of machine commonsense for if-then reasoning, 2018. 10.1609/aaai.v33i01.33013027.

[57] J. R. Scally, N. L. Cassimatis, and H. Uchida. Worlds as a unifying element of knowledge representation. In *2011 AAAI Fall Symposium Series*, 2011.

[58] R. C. Schank and R. P. Abelson. Scripts, plans, and knowledge. In *Proceedings of the 4th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'75, pages 151–157, San Francisco, 1975. Morgan Kaufmann Publishers Inc.

[59] S. Swayamdipta, S. Thomson, C. Dyer, and N. A. Smith. Frame-semantic parsing with softmax-margin segmental rnns and a syntactic scaffold, 2017.

[60] P. D. Turney and P. Pantel. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37:141–188, 2010. 10.1613/jair.2934.

[61] T. van Dijk. *Macrostructures: An interdisciplinary study of global structures in discourse, interaction, and cognition*. Routledge, 1980. 10.4324/9780429025532.

[62] D. Vrandečić. Wikidata: A new platform for collaborative data collection. In *Proceedings of the 21st international conference on world wide web*, pages 1063–1064. ACM, 2012.

[63] T. Warren, E. Milburn, N. Patson, and M. Dickey. Comprehending the impossible: what role do selectional restriction violations play? *Lang Cogn Neurosci.*, 30:932–939, 2015. 10.1080/23273798.2015.1047458.

[64] D. Wisniewski. Context-dependence and context-invariance in the neural coding of intentional action. *Frontiers in Psychology*, 9:2310, 2018. 10.3389/fpsyg.2018.02310.

[65] L. Wittgenstein. Tractatus logico-philosophicus. *London: Routledge, 1981*, 1922. 10.4324/9781315884950.