*Formal Ontology in Information Systems*
*S. Borgo et al. (Eds.)*


*doi:10.3233/978-1-61499-910-2-159*

# Amnestic Forgery: An Ontology of Conceptual Metaphors

Aldo GANGEMI [a,b] Mehwish ALAM [b] Valentina PRESUTTI [b]
[a] *FICLIT, University of Bologna, Italy*
[b] *Semantic Technology Lab, ISTC-CNR, Rome, Italy*

**Abstract.** This paper presents Amnestic Forgery, an ontology for metaphor semantics, based on MetaNet, which is inspired by the theory of Conceptual Metaphor. Amnestic Forgery reuses and extends the Framester schema, as an ideal ontology design framework to deal with both semiotic and referential aspects of frame and role mappings. The description of the resource is supplied by a discussion of its applications, with examples taken from metaphor generation, and the referential problems of metaphoric mappings. Both schema and data are available from the Framester SPARQL endpoint.

**Keywords.** Metaphor Semantics, Knowledge Extraction, Formal Ontology

## 1. Introduction

A metaphor is a cognitive operation involving usage of natural language and cross-domain conceptual mapping. Its ontological interest is in principle purely cognitive and linguistic, i.e., how to model what humans do when a metaphorical mapping is activated by speech or text. However, ontology-based extraction and representation of knowledge needs to make the semantics of natural language explicit, establishing the referential aspects of a natural language construction as used in dialogues, descriptions, memorisation, fiction, poetry, instructions, emotional expression, etc. [16].

In this paper we took the bull by the horns, and straightforwardly designed an OWL ontology for metaphors, mappings, etc. This ontology is then populated with data from Berkeley's MetaNet [6]. MetaNet is the reference repository of conceptual metaphors, developed as a Semantic Wiki[1], maintained through collaborative editing by multiple conceptual metaphor experts (cf. Sect. 2).

The new ontology is called Amnestic Forgery[2]. We deploy it as an extension of Framester[3] [11], a knowledge graph represented as Linked Open Data (LOD), which integrates heterogeneous linguistic resources (OntoWordNet [8], VerbNet [17], FrameNet-OWL [1,25], BabelNet [24], etc.), factual datasets (DBpedia [21], YAGO [32], etc.), and foundational ontologies, by providing them a unified formal semantics. We give a brief introduction to MetaNet and FrameNet in Sect. 2 and then discuss the problem of creat-

[1] https://metaphor.icsi.berkeley.edu/pub/en/
[2] This is a recursive name, since FORGERY IS AMNESIA is a new metaphor generated by means of the ontology itself, cf. Sect. 5.
[3] http://etna.istc.cnr.it/framester2/sparql

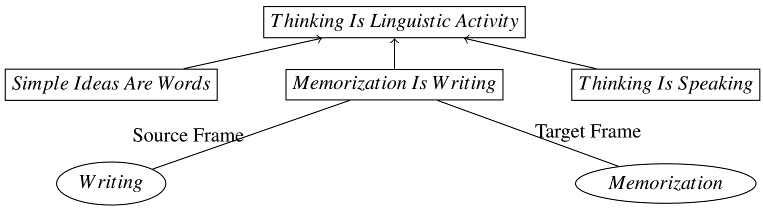


**Figure 1.** An example of network of metaphors contained in MetaNet. The arrows represent the *Inheritance* relations as defined in MetaNet.

ing a metaphor representation (Sect. 3), extracting MetaNet data and their schema, and aligning them to Framester (Sect. 4), as well as envisaging multiple use cases with practical examples (Sect. 5). We also discuss referential problems of metaphorically filtered situations (Sect. 6). We complement the paper with a survey of computational metaphor studies (Sect. 7), and conclusions.

## 2. MetaNet and FrameNet

MetaNet [6] is a repository of manually curated metaphors. It has its roots in the linguistic frames as available on FrameNet [1]. FrameNet is a resource containing conceptual frames, where each frame consists of a description of a situation as denoted by a text; e.g., the frame `Studying` depicts the situation where a `student` is performing an act of studying in some `institution` (student and institution are the "semantic roles" used to encode a studying situation).

In [6], the authors describe an automated system for extracting metaphors from text using the information from a manually built metaphor repository. The repository contains metaphors along with their conceptual frames, metaphor constructions, and metaphoric relational patterns. Figure 1 shows an example of a metaphor **Memorization is Writing** from MetaNet, which is related to another metaphor **Thinking is Linguistic Activity**. The other neighboring metaphors such as **Simple Ideas Are Words** and **Thinking Is Speaking** are also connected to the same generic metaphor. The ellipses represent the source and the target frames of one of the metaphors i.e., **Memorization Is Writing**. ‘‘Memorization" frame already exists in FrameNet while ‘‘Writing" frame is a MetaNet specific frames. In Sect. 4 we describe how the informal encoding of MetaNet has been extracted, refactored, and integrated into the Framester knowledge graph.

## 3. Representing metaphors and blending in linguistic-factual semantics

The range of metaphorically-laden linguistic constructions is large. In order to focus on the semantic aspects of metaphors, we mostly provide examples from a unique linguistic construction: the adjective-noun phrase modification. As we reported in [12], when an adjective modifies a noun, we can represent that modification as a frame composition. The composition can have a conservative effect upon the compositional semantics of the discourse (as in the case of a specification), or can act as a non-conservative extension, as in the case of metaphors.

Examples for non-metaphorical adjectival modification include global intersective composition (*American woman*), local attributional composition (*skillful woman*), attitude-laden frame construction (*alleged woman*), or "privative" novel frame creation (*stony woman*). This can be easily understood when considering a fixed `Woman` frame, and composing it with a `Nationality` frame vs. a `Competence` frame vs. an `AttitudeTowards` frame, vs. a `ConstitutingMaterial` frame. These four composition types work in different ways, but what they basically do is to establish a new referential frame, which, in the case of conservative composition, (1) inherits from the core one (`Nationality+Woman`, `Competence+Woman`), or announces a meta-level predication over the core frame (`AttitudeTowards(Woman)`), or, in the case on non-conservative composition, (2) provides instructions to create a blending (`ConstitutingMaterial+Woman`), in which a new frame emerges that does not inherit from the core one, rather it reuses part of the roles in the core frame (`Woman`), while substituting others with the blended frame (`ConstitutingMaterial`).

When dealing with adjectives that are used metaphorically (e.g., *stony woman* in the sense of a woman that does not show feelings or sympathy), classical theories of metaphor (from structure mapping [14] to embodied conceptual metaphor [20]) tell us that two frames are composed in more or less fixed ways, where e.g., certain roles from the `Feeling` frame associated with a target frame *Woman* are substituted by roles from a source frame (*ConstitutingMaterial*).

Metaphoricity (the property of having a metaphorical interpretation) is sensible to modification structures in language [23], so that e.g., adjective-noun constructions tend to have a double interpretation when the two frames seem to be incompatible. In *stony woman*, we have specific frame incompatibility (stone as a material is incompatible with roles from the `Woman` frame), but depending on the modification semantics adopted by the interpreter, the referential semantics changes completely. In the privative case, the head of the construction (*woman*) is interpreted as a fake woman made of stone, e.g., a statue, while in the metaphoric case, the *stony* modifier is interpreted by blending the `ConstitutingMaterial` and `Feeling` frames.

Distinguishing between privative and metaphoric interpretation of apparently incompatible modification is by no means trivial. It requires a larger frame composition, supporting the need for deep knowledge extraction in order to approximate human-level Natural Language Understanding (NLU). For example, in sentence 1:

The Giza sphinx is a chimeric stony woman (1)

the type of the header *woman* is not Organism, but Statue, and lets us adopt the privative interpretation. On the contrary, in the sentence 2:

The few words from that stony woman gave me the creeps (2)

the type of the header is not changed, and a metaphoric interpretation must be applied. Anyway, more intermediate cases can be imagined: if deep extraction from full discourse parsing is needed to distinguish between these clear examples, much more background knowledge would be needed to achieve accuracy in arbitrary sentences. For this very reason, we propose to integrate metaphoric knowledge into a large-scale graph of linguistic and factual knowledge, Framester [11].

Empirically, previous work on framal adjective semantics [12] has been extended by generalising the frame semantics underlying any piece of linguistic content or data. It is formalised in the Framester knowledge graph [11], which adds frame annotations to millions of words, word senses, synsets, individuals, classes, relations, etc. Recently, we have converted the MetaNet wiki information as a RDF knowledge graph with an OWL ontology compatible to the Framester schema, linked MetaNet frames to existing Framester frames, and added the resource to Framester, in order to make it queryable with ordinary APIs for computational experiments. In this paper, we report about the conversion and integration of MetaNet into Framester along with the resulting ontology. We also show some knowledge graph operations on Framester, which allow to generate novel candidate metaphor extensions.

## 4. Amnestic Forgery: an ontology for MetaNet and beyond

*Elements of Framester for Metaphor Representation:* Framester ontology is based on Descriptions and Situations (D&S) [9,10]. It is a flexible ontology pattern framework that can be used with any reasoning pipeline in order to perform classification, partial matching, diagnosis, abstraction, or construction operations between a theoretical (Description, Frame) structure, and a factual (Situation, Frame Occurrence) structure. Typically, entities in a situation are typed by concepts from a description.

The D&S formal framework can easily represent frames. As informally exemplified in the FrameNet core schema: "For example, the **Apply heat** frame describes a common situation involving a Cook, some Food, and a Heating Instrument, and is evoked by words such as bake, blanch, boil, broil, brown, simmer, steam, etc. We call these roles frame elements and the frame-evoking words are lexical units in the **Apply heat** frame".

Once formalised in D&S, the FrameNet implicit schema becomes a semiotic passepartout: a frame $f$, as a description, can be the reification of any relation $\rho$ with arbitrarily variable arity, a frame element $fe$ is a binary projection of $\rho$, and a lexical unit $lu$ of $f$ is a symbol, for which $\rho$ (and its reified counterpart $f$), and its projections $fe_{1...n}$ act as intensional interpretations. A "common situation" $s$ described by $f$ is the extensional interpretation (aka *denotation*) of $lu$, whose intension is $f$.

Since D&S allows descriptions to be composed of unary concepts, and situations to be composed of arbitrary entities, the game becomes more interesting, and enables the formal representation for the dependency of $fe$ on $f$, and the formal construction of $s$ out of arbitrary entities $e_{1...n}$ corresponding to the projections $fe_{1...n}$ of $f$. An extensive explanation of the FrameNet-OWL resource designed according to D&S is presented in [25].

Later, this approach to abridge semiotic and model-theoretical representation of frame semantics has been broadened in order to encompass any linguistic or factual resource, and opened the way to Framester [11], a large knowledge graph containing more than 50 million triples linking millions of linguistic, conceptual, or real world entities. In Framester, D&S-based frame semantics allows to reduce the heterogeneity of the targeted resources as follows:

- Framester implements a dual frame semantics by means of OWL2 [26] punning, so that each instance of the `Frame` class is also a subclass of the `FrameOccurrence` class.

- Furthermore, the *FrameProjection* class allows to integrate any predicate defined either intensionally or extensionally in ontologies, lexical resources, or other vocabularies or web formats, For example, the FrameNet frame `Activity_start` as well as the VerbNet verb class `verbclass-begin-55.1-1` are linked as intensionally equivalent to the Framester frame `ActivityStart`, while the synset `synset-newcomer-noun-1` from WordNet, which is intensionally mapped to FrameNet `Activity_start`, is extensionally represented as a class of newcoming entities, and linked as a *unary projection* of a Framester class of newcoming situations, `Newcomer.n.1`, which in-turn is represented as a subclass of `ActivityStart`.
- An `InternalBinaryProjection` of a frame, such as semantic roles from FrameNet, VerbNet, PropBank, the Preposition Project, etc., as well as properties from factual resources as e.g., the DBpedia Ontology, are aligned to generic roles in Framester, leading to better interoperability when extensionally represented as internal binary projections of a frame.
- An `ExternalBinaryProjection` of a frame (e.g., a relation between the agent and the location of a situation for the frame `ActivityStart`) can also be generalised by mapping them to pairs of internal binary projections.
- Finally, individuals classified with types that can be disambiguated as unary projections of a frame (e.g., "a newcomer"), can be formalised as potential evokers of that frame in the context of a discourse.

In D&S, higher-level descriptions can also be defined, e.g., for meta-norms that describe priority between other norms. Accordingly, frame composition (see Sect. 3 can be formalised in D&S as a higher-level description that creates a new frame by assigning roles to the composing frames. A metaphor itself can be now represented as a kind of description, which incorporates roles for two more descriptions (the source and target frames), as well as mapping rules between the respective roles.

*Representing and Linking Metaphors:* Equipped with this intuition, originally sketched in [10], and already used in several projects to formalise linguistic resources (OWL versions of FrameNet, VerbNet, PropBank, etc.), we have the possibility to reuse Framester [11] when formalising MetaNet.

The design approach taken to formalise MetaNet is to use D&S in order to extract and formalise the MetaNet schema, then to extract data, formally represent them according to the new schema, and finally to align both schema and data to elements in the Framester knowledge graph.

We firstly scraped tabular data from the MetaNet wiki[4], and designed a preliminary MetaNet schema that catches the intended meaning of the interface used to populate the MetaNet wiki. Secondly, we refactored the extracted data according to this preliminary schema, and fine-tuned it against features deriving from the data entry variety in the wiki. The result is a refined schema, the Amnestic Forgery ontology, and its MetaNet data. Fig. 2 depicts a subgraph of MetaNet for the metaphor CRIME_IS_A _DISEASE. The subgraph contains examples of the core relations in MetaNet, linking metaphors to their source and target frames, their role mappings, entailments, and other relations contributed by the users of the wiki, which deserve further investigation to be fully axiomatized.

[4]The MetaNet wiki is a SemanticMediaWiki instance, but its data querying facility is not accessible.

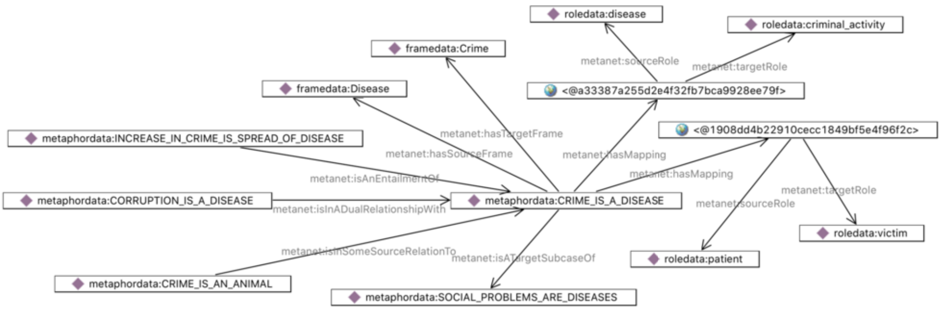


**Figure 2.** The subgraph for the metaphor CRIME_IS_A _DISEASE.

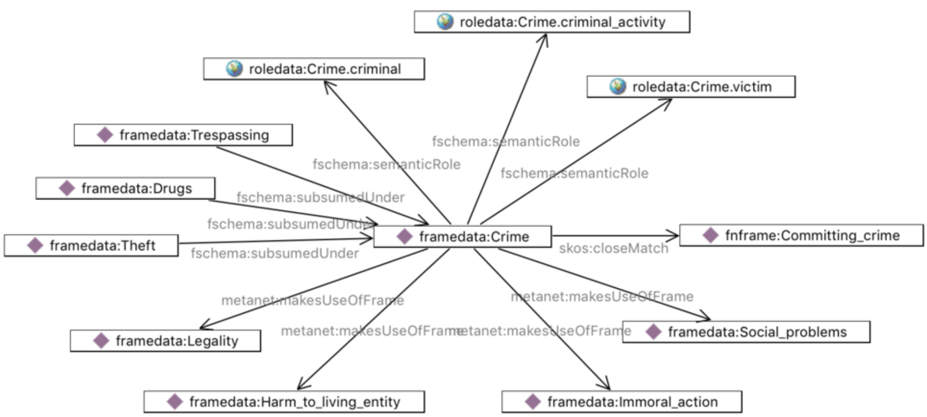


**Figure 3.** The subgraph for the frame `Crime`.

Fig. 3 depicts a subgraph of MetaNet for the frame `Crime`. The subgraph contains examples of the core relations in MetaNet, linking frames to their roles, their more specific frames, their alignments to frame from other resources, and other frames bearing dependency relations.

Notice that the subgraphs showed here are examples of the intensional view of Amnestic Forgery (metaphors, frames, mappings, and roles are all represented at the Description level), while the extensional view adds axioms that have metaphors and frames as classes of situations, and relations to roles as OWL restrictions (quantified clauses as superclasses of frames).

A summary diagram of the axiomatisation for Amnestic Forgery is showed in Fig.

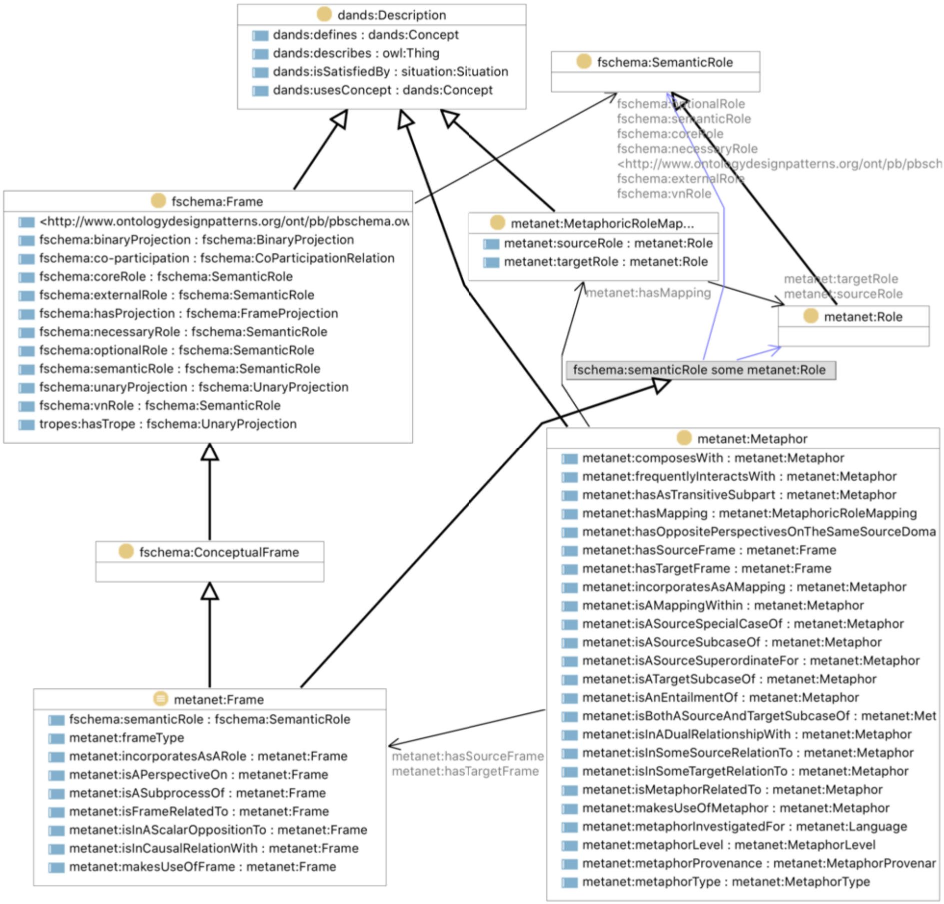


**Figure 4.** A class-diagram profile for the OWL axiomatisation of Amnestic Forgery.

4. The diagram uses a UML-class-diagram-oriented profile to sketch the core axioms for the `Metaphor` class, shown either as "attributes" within class boxes, or as either "associations" or "generalisations" (subsumption) across class boxes. The diagram summarizes the reuse of the `Description` class from D&S, which subsumes the `Metaphor`, `Frame`, and `MetaphoricRoleMapping` classes. A hierarchy of frame and role notions exemplify Framester schema alignements, and the treatment of semantic roles as both binary projection of frames, and OWL properties (binary relations). Association-like edges derive from either domain or range restrictions in the OWL encoding of Amnestic Forgery, or from existential restrictions. The OWL resource with full axiomatization including imports, alignments, disjointness, and documentation axioms is downloadable from Github[5] and queried through the Framester SPARQL endpoint[6].

In order to take full advantage of Framester, alignment of metaphors and frames to other resources such as FrameNet and WordNet is necessary. However, only about

[5] https://github.com/alammehwish/AmnesticForgery
[6] http://etna.istc.cnr.it/framester2/sparql

25% of MetaNet frames are aligned to FrameNet frames, therefore an alignment completion is needed. An example shows the non-triviality of this completion. The metaphor ABUSIVE_POLITICAL_LEADERS_ARE_PHYSICAL_BULLIES could be aligned by using the alignments already existing in Framester, for example, *abusive*, *(political) leader*, and *bully* are all present in WordNet, and are aligned to the following FrameNet frames (as unary projections) respectively: `Abusing`, `Leadership`, and `Manipulate_into_doing`. However, some MetaNet frames can only be aligned to FrameNet if we consider them as frame compositions: `Abusing+Leadership`, while others, like `Manipulate_into_doing`, require a specialisation to a further feature (here: physical). The alignment completion activity becomes then a discovery activity, where new frames can be proposed, and defined as compositions of existing frames.

## 5. Generating new metaphors with MetaNet and Framester

In order to prove the advantages of having a large and formally rigorous knowledge base, we report here a query to Framester extended with Amnestic Forgery and MetaNet data. Given a MetaNet metaphor (CRIME_IS_A_DISEASE in the example), the query is able to generate hundreds of novel intensional metaphors. A web application is available on-line[7], where the user can input a metaphor from MetaNet, and the API generates new metaphors. The results may sound "strange" and difficult to understand, but this is exactly what happens when a novel metaphor appears. One such example is the FORGERY_IS_AMNESIA metaphor, from which we have generated the eponymous *amnestic forgery* term as a linguistic rendering. *amnestic forgery* appears to be novel: no realisations can be found e.g. on the Web (based on Google searching).

In order to make sense of it, we need to focus on the abstraction provided by the general metaphor CRIME_IS_A_DISEASE, which maps roles of the Crime frame into roles of the Disease frame. Some known examples include *corruption is infecting our administration*, or *Mafia is an epidemic*. But since the general metaphor is cognitively active, we may expect that any evocation of the Crime and Disease frames will instantiate the metaphor, probably with unpredictable (inspiring, poetic, hilarious, obscure) effects.

FORGERY_IS_AMNESIA creates the mapping of a type of crime (forgery) to a type of Disease (amnesia): it is basically well-formed according to the hypothesis of general metaphors, but it is more difficult to understand. We can say something like *forgery is a gangrene in Argentina's economy*, but can we say also *forgery is an amnesia in Argentina's economy*? probably not.

The example leads us to revisit our metaphor generation hypothesis: a general metaphor shows that there is a cognitive disposition to understand certain metaphors, but their acceptability seems to require additional conditions, which are not entirely known. In fact, typical examples of CRIME_IS_A_DISEASE are very basic: diseases are used to convey the idea of malfunctioning and/or spreading, while subtler pathologies, such as amnesia, seem to disrupt the basic intuition, and might be more adequate for dadaist poetry than as regular conversational means.

A closer analysis shows that metaphors might be richer than what CMT and MetaNet tell us. In the FORGERY_IS_AMNESIA metaphor, the role mapping is correct: a community is affected by forgery, and such affection is metaphorised as amnesia. Then what

[7] https://lipn.univ-paris13.fr/framester/en/metanet/

is the showstopper here? maybe (fictive) causality: while metaphorising forgery as a gangrene evokes damage and loss, which are active in the semantic space of forged things such as money, metaphorising forgery as an amnesia requires evoking a community that starts forgetting because of the forgery, but it is very hard to imagine e.g. that money faking induces amnesia.

If this closer introspection could be generalised, we should add a new dimension to CMT, i.e. the (fictive) plausibility of a causal relation. Maybe there are also other relations that act as conditions for accepting novel metaphors in non-dadaist contexts.

A second query type could be conceived in order to search for MetaNet metaphors that explain a metaphoric expression, such as *bright coach*, taken from a large repository of metaphoric adjective-noun phrases[8].

The queries, which can be tested on the Framester SPARQL Endpoint given above, are designed to use the adjective-noun phrase construction, and their related senses and frames in Framester. Alternative choices can lead to different metaphor-evoking phrases.

```
prefix metanet: <https://w3id.org/framester/metanet/schema/>
prefix framedata: <https://w3id.org/framester/metanet/frames/>
prefix metaphordata: <https://w3id.org/framester/metanet/metaphors/>
SELECT DISTINCT ?ssyn ?tsyn
WHERE {
metaphordata:CRIME_IS_A_DISEASE metanet:hasSourceFrame ?s ;
     metanet:hasTargetFrame ?t .
?s skos:closeMatch ?fns . ?fns a fn15schema:Frame .
?t skos:closeMatch ?fnt . ?fnt a fn15schema:Frame .
?fns skos:closeMatch ?ssyn .
?fnt skos:closeMatch ?tsyn .
{?ssyn a wn30schema:AdjectiveSynset}
UNION
{?ssyn a wn30schema:AdjectiveSatelliteSynset}
?tsyn a wn30schema:NounSynset }
```

## 6. Referential aspects of metaphorical mappings: a case for quasi-truth

As anticipated in Sect. 1, we want to represent not only conceptual metaphors as frame mappings, but also the factual knowledge that is possibly affected by those mappings. Does metaphor involve an actual referential movement in the domain of discourse, or is it just a "false movement", a sort of distorting lens that does not affect referential aspects at all? In the second case, the ontological relevance of metaphors would be confined to intensional aspects impacting cognitive or linguistic worlds, as Davidson would have probably subscribed due to his belief in purely linguistic relevance of metaphors [5]. In the first case, we should admit that frame mapping has a factual import.

Let's consider the trailing example: the CRIME_IS_A_DISEASE metaphor as a frame mapping. Intensionally, it is pretty clear that after applying the metaphor, some roles from the `Crime` frame accept values from the `Disease` frame. However, what is happening extensionally? What are corresponding metaphoric situations like?

[8] http://pages.ucsd.edu/~e4gutier/m4p/AN-phrase-annotations.csv

Following Amnestic Forgery, a metaphoric situation is the counterpart of a metaphoric description, where entities from two frame occurrences are blended when playing the particular role mappings enabled by the metaphor. For example, since the CRIME_IS_A_DISEASE axioms dictate that the `criminal_activity` role from `Crime` is swapped for the `disease` role from `Disease`, and the `victim` role from `Crime` is swapped for the `patient` role from `Disease`, the denotation of sentence 3:

Corruption has infected our community. (3)

would not be assigned within the literal domain of discourse.

The literal assignment would work like this:

- a specific series of corruption events *ce* is the value for the `criminal_activity` role from a `Crime` situation *cs*: $CA(cs, ce)$
- a specific series of infection events *ie* is the value for the `disease` role from a `Disease` situation *ds*: $D(ds, ie)$
- a specific community *com* is the value for the `victim` role from a `Crime` situation *ds*: $V(cs, com)$

When the metaphor is activated, the mapping produces the following blended frame occurrence:

- a specific series of corruption events *ce* is the value for the `disease` role from a `Crime+Disease` situation *cds*: $D(cds, ce)$
- a specific series of infection events *ie* is the value for the `disease` role from a `Crime+Disease` situation *cds*: $D(cds, ie)$
- a specific community *com* is the value for the `patient` role from a `Crime+Disease` situation *cds*: $P(cds, com)$

What happened? In this metaphor occurrence, once the metaphorical mapping generates the blended frame, the `criminal_activity` role value either (a) results to be the same thing as the `disease` role value: $ce = ie$, or (b) melts into a new hybrid entity $ce + ie$. That hybrid entity is a strange beast, but in formal ontology we are accustomed to such things: aggregates, amalgams, qua-entities, etc.

The actual problem is if our design could satisfy application requirements by leaning to an economic commitment to our domain of interpretation $\Delta$: $ce = ie \in \Delta$, or else if we commit to a multiplicative design style, which accepts commitments also to a new entity $ice \in \Delta$, generated via metaphor.

The decision is not trivial. On one hand, we can safely claim that no physical substrate of entities involved in the blended frame are changed by the metaphor, and that on the contrary, the space of cognitive entities is most probably affected. On the other hand, it is not obvious at all what are the social entities possibly involved in the blending.

A paradigmatic case of the social impact of metaphor is fake news, and more specifically what we call *quasi-true facts*, i.e. sentences that distort facts in a way that make them not really false, but "alternative" to (supposedly true) ones.[9] Spin doctors speech heavily adopts quasi-truth in order to obtain the best cognitive impact on citizens [19], e.g., when one political party narrates socially-relevant criminal cases by using the CRIME IS A DISEASE metaphor (in this case, typical entailment includes that that

[9] *Alternative facts* are seriously used as a category in communication science talk, cf. [22].

crime spreads like disease in human societies, and its etiology must be destroyed), while another uses the CRIME IS A PHYSIOLOGICAL PROCESS metaphor (in this case, the typical entailment is that crime is organic to human societies).

When addressing the referential aspects of quasi-truth, the physical substrate of facts does not change, the cognitive impact is high (and can even lead to massive social change), while social reality still needs an appropriate characterisation: as a social entity, are corruption events – seen like disease spreading – a new entity in $\Delta$?

We do not propose a preferred design pattern for the social ontology of metaphors, since we'd rather defer conclusions after an ongoing empirical study. However, we do hope that the problems at hand are sufficiently clear after the application of Amnestic Forgery.[10]

## 7. Related work

Occurrence of a metaphor in language is very frequent. It is not only what the meaning of a word is used in certain context, it is also deeply embedded in the fact that the words used in different domains are re-used differently in another context based on the knowledge a particular word or lexical unit reflects. The reference theory is Conceptual Metaphor Theory (CMT) proposed by Lakoff and Johnson in [20].

Despite computational work started many years ago [28], the computational metaphor literature is rather small, cf. [36] for a survey. Recently the problem seems to have caught attention though. Computational applications nowadays include *(i)* Machine Translation: the metaphors vary across cultures and languages, *(ii)* Sentiment Analysis: Linguistic units may look positive or negative, but polarity could be inverted because of irony or figurative speech, *(iii)* Information retrieval: faulty information may be found in textual data without properly attaching meaning to metaphorical phrases, *(iv)* Computational Creativity: reverse engineering human ability to generate appropriate metaphors is key to computational creativity. Following these lines, three major tasks can be singled out: Metaphor Detection, Interpretation, and Generation. Where appropriate, we compare existing work to the ontology and themes described in this paper.

*Metaphor Detection.* Over the past few years there has been a rise in the development of statistical methods for detecting metaphors. Many of these techniques take advantage of vector-space models, and perform a binary classification of metaphorical vs. literal occurrences in text [34,33], which may come from compositions of word pairs such as "sweet" and "person", where *sweet* is only metaphorical when composed in phrases with words that do not denote tastable entities, such as "person". [15] proposes a Compositional Distributional Semantic Model (CDSM), which generates a vector representation of the phrases. [2] also introduces a framework based on CDSM, targeting adjective-noun constructions. In these cases, the meaning of the phrase is derived by composing the representations of adjectives and nouns. As a contrast, [31] uses clustering techniques over nouns and verbs to perform metaphor identification. It takes manually annotated metaphors as a starting point, and then, based on syntactic similarity, detects a large number of metaphors from a corpus. [18] uses binary classification of the verbs

[10] Funnily enough, after naming the new ontology, the authors realised that Amnestic Forgery is a good literal description for quasi-truth as well.

into metaphorical or literal using semantic classes of the verbs such as grammatical, resource-based, or distributional. While in [3] the authors propose a two step approach for detecting if the an adjective-noun pair is a metaphor or a literal. They use pre-trained word vectors of the AN-pairs as input vectors and then propose a neural network for composing AN phrases. Finally, [4] uses eventive information in detecting metaphors in Chinese. Metaphor detection can be nicely paired with metaphor-oriented knowledge graph processing. For example, we plan to pre-process corpora with a model trained with combinatorially-generated metaphors, in order to detect if such metaphors are used, and in what context, thus generating additional knowledge.

*Metaphor Interpretation.* Metaphor interpretation requires complex analogical comparison and inferencing, because it performs cross-domain projection of knowledge. One way to assign interpretations to metaphors relies on MetaNet [6]. Each metaphor defined in MetaNet is manually encoded, and is connected to a combination of linguistic frames, often aligned to those available in FrameNet [1]. A detailed survey about metaphor processing systems is [30]. As far as our ontology is concerned, we intend to enrich knowledge extraction pipelines such as FRED [13] with metaphoric sensitivity, thus contributing to automated metaphor interpretation.

*Metaphor and Blending Generation.* Tony Veale has devoted substantial effort to make automated metaphor generation a reality, cf. [35] for a summary of his recent attempts to make it a creative agent on the social media. The most advanced computational framework for automated conceptual blending is Eppe et al. [7]. It is able to implement most of the nuances of the theory, but does not make use of knowledge graphs or public ontologies (they only mention ongoing work within the OntoHub repository). Amnestic Forgery is obviously reusable by them in order to experiment with new case studies involving metaphorical blending, as well as by reusing the whole Framester as background knowledge.

Interesting formal ontology work has also been conducted with reference to cognitive conceptual spaces [27], or description logics [29], but none of these works attempts to build a cognitively valid metaphor ontology that can be also exploited empirically in the current huge knowledge graphs that are populating the Web.

The work reported in this paper is to our knowledge the first attempt to position metaphor and blending theories within an open large graph that can be used to test or extend existing theories, as well as classification, interpretation, or generation algorithms.

## 8. Conclusions

We have presented Amnestic Forgery, an ontology that enables the integration of data related to Conceptual Metaphors into large knowledge graphs. The ontology is designed starting from the D&S ontology pattern framework, and inherits the Framester schema, which already integrates multiple linguistic and factual data and schemas, besides some foundational ontologies. The resource is intended to be used as background knowledge for empirical formal research on metaphor phenomena in discourse.

The benefits of integrating a large metaphor dataset into the Framester linguistic-factual graph may spread to any research work on metaphor: metaphoric sentence classification, metaphor generation, as well as empirical research on cognitive or formal theories of metaphor.

The paper also includes references to the current computational research on metaphors, describes a simple SPARQL-based method to generate novel metaphors, and discusses the challenging referential problems of metaphorically-induced situations in social reality. An algorithm for aligning FrameNet frames with MetaNet frames is under development.

Ongoing and future work bears multiple directions: using deep learning techniques to both detect metaphorical sentences in text, and automatically enriching metaphoric knowledge graphs; deepening the axiomatisation of Amnestic Forgery, and empirically testing it on large focused corpora; integrating Amnestic Forgery with computational blending and creativity platforms. Finally, fully aligning MetaNet frames to FrameNet frames.

## References

[1] Collin F. Baker, Charles J. Fillmore, and John B. Lowe. The Berkeley FrameNet Project. In *Proc. of COLING 1998*, pages 86–90, Montreal, Quebec, Canada, 1998. ACL.

[2] Marco Baroni and Roberto Zamparelli. Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, EMNLP 2010, 9-11 October 2010, MIT Stata Center, Massachusetts, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1183–1193. ACL, 2010.

[3] Yurri Bizzoni, Stergios Chatzikyriakidis, and Mehdi Ghanimifard. âĂIJdeepâĂİ learning: Detecting metaphoricity in adjective-noun pairs. In *Proceedings of the Workshop on Stylistic Variation*, pages 43–52, 2017.

[4] I-Hsuan Chen, Yunfei Long, Qin Lu, and Chu-Ren Huang. Leveraging eventive information for better metaphor detection and classification. In Roger Levy and Lucia Specia, editors, *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017), Vancouver, Canada, August 3-4, 2017*, pages 36–46. Association for Computational Linguistics, 2017.

[5] Donald Davidson. What metaphors mean. *Critical Inquiry*, 5(1):31–47, 1978.

[6] Ellen Dodge, Jisup Hong, and Elise Stickles. Metanet: Deep semantic automatic metaphor analysis. In *Proceedings of the Third Workshop on Metaphor in NLP*, pages 40–49. Association for Computational Linguistics, 2015.

[7] Manfred Eppe, Ewen Maclean, Roberto Confalonieri, Oliver Kutz, Marco Schorlemmer, Enric Plaza, and Kai-Uwe Kühnberger. A computational framework for conceptual blending. *Artif. Intell.*, 256:105–129, 2018.

[8] Christiane Fellbaum, editor. *WordNet: an electronic lexical database*. MIT Press, 1998.

[9] Aldo Gangemi. Norms and plans as unification criteria for social collectives. *Autonomous Agents and Multi-Agent Systems*, 17(1):70–112, 2008.

[10] Aldo Gangemi. *What's in a schema? A formal metamodel for ECG and FrameNet*. Studies in Natural Language Processing. Cambridge University Press, April 2010.

[11] Aldo Gangemi, Mehwish Alam, Luigi Asprino, Valentina Presutti, and Diego Reforgiato Recupero. Framester: A wide coverage linguistic linked data hub. In *20th International Conference on Knowledge Engineering and Knowledge Management - Volume 10024*, EKAW 2016, pages 239–254, New York, NY, USA, 2016. Springer-Verlag New York, Inc.

[12] Aldo Gangemi, Andrea Giovanni Nuzzolese, Valentina Presutti, and Diego Reforgiato Recupero. Adjective semantics in open knowledge extraction. In *Formal Ontology in Information Systems - Proceedings of the 9th International Conference, FOIS 2016, Annecy, France, July 6-9, 2016*, pages 167–180, 2016.

[13] Aldo Gangemi, Valentina Presutti, Diego Reforgiato Recupero, Andrea Giovanni Nuzzolese, Francesco Draicchio, and Misael Mongiovì. Semantic web machine reading with FRED. *Semantic Web*, 8(6):873–893, 2017.

[14] Dedre Gentner. Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7:155–170, 1983.

[15] E. Dario Gutiérrez, Ekaterina Shutova, Tyler Marghetis, and Benjamin Bergen. Literal and metaphorical senses in compositional distributional semantic models. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics, 2016.

[16] Roman Jakobson. *Linguistics and Poetics*, pages 350–377. MIT Press, 1960.

[17] Karin Kipper, Hoa Trang Dang, and Martha Stone Palmer. Class-based construction of a verb lexicon. In Henry A. Kautz and Bruce W. Porter, editors, *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on on Innovative Applications of Artificial Intelligence, July 30 - August 3, 2000, Austin, Texas, USA.*, pages 691–696. AAAI Press / The MIT Press, 2000.

[18] Beata Beigman Klebanov, Chee Wee Leong, E. Dario Gutiérrez, Ekaterina Shutova, and Michael Flor. Semantic classifications for detection of verb metaphors. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer Linguistics, 2016.

[19] George Lakoff. *The Political Mind*. Penguin Books, 2008.

[20] George Lakoff and Mark Johnson. *Metaphors we Live by*. University of Chicago Press, Chicago, 1980.

[21] Jens Lehmann, Chris Bizer, Georgi Kobilarov, Sören Auer, Christian Becker, Richard Cyganiak, and Sebastian Hellmann. DBpedia - A Crystallization Point for the Web of Data. *Journal of Web Semantics*, 7(3):154–165, 2009.

[22] Felicitas Macgilchrist. Positive discourse analysis:contesting dominant discourses by reframing the issues. *Critical Approaches to Discourse Analysis Across Disciplines*, 1(1):74–94, 2007.

[23] Marcin Morzycki. *Modification*. Cambridge University Press, 2015.

[24] Roberto Navigli and Simone Paolo Ponzetto. Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artif. Intell.*, 193:217–250, December 2012.

[25] Andrea Giovanni Nuzzolese, Aldo Gangemi, and Valentina Presutti. Gathering lexical linked data and knowledge patterns from framenet. In *Proceedings of the Sixth International Conference on Knowledge Capture*, K-CAP '11, pages 41–48, New York, NY, USA, 2011. ACM.

[26] W3C OWL Working Group. *OWL 2 Web Ontology Language: Document Overview*. W3C Recommendation, 27 October 2009. Available at http://www.w3.org/TR/owl2-overview/.

[27] Martin Raubal and Benjamin Adams. The semantic web needs more cognition. *Semantic Web*, 1(1-2):69–74, 2010.

[28] Terry Regier. *The Human Semantic potential: Spatial Language and Constrained Connectionism*. MIT Press, 1996.

[29] Marco Schorlemmer, Roberto Confalonieri, and Enric Plaza. Coherent concept invention. In *Proceedings of the Workshop on Computational Creativity, Concept Invention, and General Intelligence (C3GI 2016) co-located with the 28th European Summer School in Logic, Language and Information (ESSLLI 2016), Bozen-Bolzano, Italy, August 20-22, 2016.*, 2016.

[30] Ekaterina Shutova. Design and evaluation of metaphor processing systems. *Computational Linguistics*, 41(4):579–623, 2015.

[31] Ekaterina Shutova, Lin Sun, and Anna Korhonen. Metaphor identification using verb and noun clustering. In Chu-Ren Huang and Dan Jurafsky, editors, *COLING 2010, 23rd International Conference on Computational Linguistics, Proceedings of the Conference, 23-27 August 2010, Beijing, China*, pages 1002–1010. Tsinghua University Press, 2010.

[32] Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: A large ontology from wikipedia and wordnet. *Web Semant.*, 6(3):203–217, September 2008.

[33] Yulia Tsvetkov, Leonid Boytsov, Anatole Gershman, Eric Nyberg, and Chris Dyer. Metaphor detection with cross-lingual model transfer. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 248–258. The Association for Computer Linguistics, 2014.

[34] Akira Utsumi. Computational exploration of metaphor comprehension processes using a semantic space model. *Cognitive Science*, 35(2):251–296, 2006.

[35] Tony Veale. Round Up The Usual Suspects: Knowledge-Based Metaphor Generation. In *Proceedings of the Fourth Workshop on Metaphor in NLP*, pages 34–41, 2016.

[36] Tony Veale, Ekaterina Shutova, Beata Beigman Klebanov, Graeme Hirst, and Carlo Strapparava. Metaphor: A computational perspective. *Computational Linguistics*, 44(1), 2018.