

# EMISSOR: A platform for capturing multimodal interactions as Episodic Memories and Interpretations with Situated Scenario-based Ontological References

Selene Báez Santamaría, Thomas Baier, Taewoon Kim, Lea Krause, Jaap Kruijt, Piek Vossen

Computational Linguistics and Text Mining Lab (CLTL)

Vrije Universiteit Amsterdam, Netherlands

s.baezsantamaria, t.baier, t.kim, l.krause, j.m.kruijt,  
piek.vossen@vu.nl

## Abstract

We present EMISSOR: a platform to capture multimodal interactions as recordings of episodic experiences with explicit referential interpretations that yield an episodic Knowledge Graph (eKG). The platform stores streams of multiple modalities as parallel signals. Each signal is segmented and annotated independently with interpretation. Annotations are eventually mapped to explicit identities and relations in the eKG. As we ground signal segments from different modalities to the same instance representations, we also ground different modalities across each other. Unique to our eKG is that it accepts different interpretations across modalities, sources and experiences and supports reasoning over conflicting information and uncertainties that may result from multimodal experiences. EMISSOR can record and annotate experiments in virtual and real-world, combine data, evaluate system behavior and their performance for pre-set goals but also model the accumulation of knowledge and interpretations in the Knowledge Graph as a result of these episodic experiences.

## 1 Introduction

Multimodal interaction in real-world settings using sensors between humans and agents is a complex process. Furthermore, it typically evolves over time and within a shared (physical) space, being bound yet remaining continuously dynamic. The fact that certain contextual factors are not physically present, such as past episodic encounters, background knowledge and intentions, adds to this complexity. Agents designed to behave intelligently need to handle this complexity and form teams with people to collaborate and achieve shared goals.

Within the Hybrid Intelligence framework,<sup>1</sup> we are specifically interested in such collaborative set-

tings and focus on analysing what causes such systems to succeed or fail. Collaboration requires shared grounding and partially shared understanding of situations, communications, and references across modalities. As humans and agents may have different beliefs and perceptions of these situations, we argued in previous work (Vossen et al., 2018, 2019a) that agents need a theory-of-mind (ToM) (Premack and Woodruff, 1978; Leslie, 1987) to handle conflicts, miscommunication and errors in referential grounding and interpretations.<sup>2</sup>

Although there are many initiatives for representing multimodal interactions, referential grounding is hardly handled in its full complexity. Most approaches to multimodal interaction data either *label* media such as video, images, or audio with annotations or simply present situated agent interactions in the form of dialogues or actions without labeling. In these approaches, annotations may be seen as interpretations of direct “behavioral” responses (utterances or actions) to the preceding signals. However, they lack a formalization of these interpretations into an explicit model that supports transparent reasoning. Such multimodal data sets can be seen as episodic experiences but not yet as knowledge-aware episodic memories that reflect the cumulative result. The latter requires interpretations of different multimodal signals to be combined in an explicit knowledge structure according to an ontological model that reflects our conceptualization of the world. In addition, this model needs to handle alternative interpretations, uncertainties and conflicts as the interpretations are not always correct or consistent.

We propose a generic model that can capture multimodal interactions as recordings of episodic experiences with explicit referential interpretations that also yield an episodic Knowl-

<sup>1</sup>[www.hybrid-intelligence-centre.nl](http://www.hybrid-intelligence-centre.nl)

<sup>2</sup>[makerobotstalk.nl](http://makerobotstalk.nl)

edge Graph (eKG) as a ToM . **EMISSOR** stands for **E**pisodic **M**emories and **I**nterpretations with **S**ituated **S**enario-based **O**ntological **R**eferences. The platform stores multiple streams of modalities as parallel signals. Each signal can be segmented and annotated independently with interpretation, providing robustness and simplicity. Signals can thus represent natural conversations in situated contexts in which (visual) actions and (verbal) utterances can be responses to each other but can also happen independently. EMISSOR can represent any (multimodal) interaction that takes place either in virtual or real-world settings, involving any virtual or real-world agent.

Annotated signals do not necessarily stand on their own, but can be mapped to explicit identities, relations, and properties in an eKG for capturing time-series of instances of situations. These mappings ground signal segments to formal instance representations and ground different modalities across each other. Time-bound experiences are thus captured as episodic experiences, i.e. as an explicit cumulative interpretation of streams of signals. The eKG models knowledge and interpretation shifts over time and supports reasoning over the interpretation. By keeping track of the provenance of signals and interpretations, our model reflects alternative ToM interpretations from different sources, modalities and experiences.

In the current paper, we set out the basic design and structure of our representation and our motivation. We first discuss in Section 2 representations of multimodal interaction proposed in various paradigms such as virtual games, agent interactions, multimodal dialogue systems. In Section 3, we describe the desiderata for our proposal for representing interactions, which combines aspects from the different approaches in the related work but adds a KG as the host of such transparent episodic memories for situated experiences. We elaborate on the different data layers and relations in our proposal. We discuss how different types of data sets can be converted, aligned and annotated such that segments get grounded to identities in a Knowledge Graph using an annotation tool. Future research and conclusions are presented in the final Section 4.

## 2 Related work

Interaction data takes many forms. Not only can it come in different modalities (visual, au-

dio, text, action), but we can also have different types of interactions, e.g., search, question-answering, command-action sequences, (task-based) dialogues, navigation tasks, games, graphical interfaces, plain video, and audio recordings. It is impossible to provide a comprehensive overview of representations in each separate modality and interaction type in this paper.

For our research on social communicating robots, we are interested in representations of multimodal interactions with *referential grounding* across modalities and the representations of these modalities as such. Therefore, we discuss mainly works on modalities aligned in time series representing interactions. This excludes data in single modalities, such as plain text corpora with dialogues (spoken or text) and image or video data without dialogues. It also excludes static data that does not represent temporal sequences of data. For example, visual data labeled with textual descriptions and textual data augmented with visual scenes do not necessarily represent interactions. In interactions, modalities partially complement each other, such as speech responding to speech or to scenes and actions following speech. Such sequences often react to and complement each other and reflect some degree of causality and coherence, but not entirely. Augmented modalities, on the other hand, mainly represent paired data where one modality describes or illustrates the other. Collections of augmented data, e.g., labeled Flickr images, do not exhibit coherence across data points and do not reflect causal interaction. Nonetheless, single modality data and non-interactive data can still aid in processing multimodal interaction. Models and classifiers trained on such data can automatically annotate scenes in a time series of multimodal interaction data. An interesting research question would be whether static or single modality annotations also model collaborative interactions without considering temporal, causal, and coherence relations across data.

A recent survey of interactive dialogue datasets is given in (Serban et al., 2018). Serban et al. (2018) differentiate dialogue systems by the types of interaction (goal(s), non-goal chat, topical); their modalities (written, spoken, video); the participants (Human, Agent); being constrained, spontaneous, scripted, fictional; being goal-oriented, domain-specific or open. Most of the data described is, however, non-situated or not situation-grounded. There is hardly any reference to and interaction with phys-

ical contexts. Many of these datasets and tasks have been developed and presented in SigDial<sup>3</sup>, the ACL Special Interest Group, for research on dialogues structures and models. A more applied perspective is taken by the Dialogue System Technology Challenge (DSTC<sup>4</sup>). DSTC provides a platform for researchers and industry to develop and evaluate agent-interaction systems. Older datasets mostly contain conversational data for chatbots. A more recent challenge, Audio Visual Scene-Aware Dialog (AVSD, (Alamri et al., 2019)), contains short video clips with audio, descriptive captions, and a dialogue history. Participating systems need to answer a follow-up question or ground an answer to an image from the video or an audio fragment. This challenge represents situated-references and grounding across modalities, but the conversation is very descriptive: the conversations describe situations rather than being embedded in them. Below, we briefly describe some well-known datasets and challenges that represent various types of interaction data.

ParLAI<sup>5</sup> released more than a hundred conversational datasets covering a wide range of topics, but most are single modality chat (Miller et al., 2017). SIMMC<sup>6</sup> is Facebook’s sequel to ParLAI with Situated and Interactive Multi-Modal Conversation (Crook et al., 2019). It consists of task-oriented dialogues in multimodal contexts represented by collections of images. The data contains referential relations between the dialogues and situations, but the current challenge is restricted to the e-commerce contexts of buying furniture and fashion items. There is no grounding to complex situations but intentions and goals are explicitly represented. Facebook-research also launched various other related tasks, among which RECCON<sup>7</sup>: Recognizing Emotion Cause in Conversations (Poria et al., 2020), and MINIRTS: Hierarchical Decision Making by Generating and Following Natural Language Instructions in Real-time strategy game environments<sup>8</sup> (Hu et al., 2019). The former grounds dialogues to emotions and their causes but not to visual or audio data. The latter grounds language to a closed virtual world by references to objects, agents, and actions. The dialogues are limited to

commands and instructions to operate the game.

Google developed Schema-Guided-Dialogue (SGD<sup>9</sup>) for task-oriented conversational agents (Rastogi et al., 2020). In addition to e-commerce services, the tasks involve intent prediction, slot filling, dialogue state tracking, policy imitation learning, language generation, and user simulation learning. The goals are defined, but there is no multimodal situational grounding.

An older comprehensive robot platform is provided by openEASE<sup>10</sup>, which is a web-based knowledge service providing robot and human activity data constituting episodic memories (Beetz et al., 2015). It produces semantically annotated data of manipulation actions, including the agent’s environment, the objects it manipulates, the task it performs, and the behavior it generates. EASE uses so-called NEEMS (Narrative Enabled Episodic Memories) as episodic memories. NEEMS consist of a video recording by the agent of the ongoing activity. These videos are enriched with a story about the actions, motion, their purposes, effects, and the agent’s sensor information during the activity. EASE is not data-centric but a service platform that uses a knowledge database as a back-end. The database can be explored through prolog queries. The focus of openEASE is on physical interactions and not on conversations with complex referential relations between expressions and situations.

Microsoft created a Platform for Situated Intelligence, PSI<sup>11</sup> (Bohus et al., 2017). PSI offers multimodal data visualization and annotation tools, as well as processing components for various sensors, processing technologies, and platforms for multimodal interaction. PSI models multimodal situations and interactions within and comes close to a comprehensive solution. However, PSI is a software integration platform through which developers can share modules using a streaming architecture for signal annotation. Interactions are not stored in a shared representation and the platform cannot be used for sharing experimental data independently.

Action Learning From Realistic Environments and Directives (ALFRED<sup>12</sup>) is a recent benchmark for mapping natural language instructions and ego-

<sup>3</sup>[www.sigdial.org](http://www.sigdial.org)

<sup>4</sup>[dstc9.dstc.community](http://dstc9.dstc.community)

<sup>5</sup>[github.com/facebookresearch/ParLAI](https://github.com/facebookresearch/ParLAI)

<sup>6</sup>[github.com/facebookresearch/simmc](https://github.com/facebookresearch/simmc)

<sup>7</sup>[github.com/declare-lab/RECCON](https://github.com/declare-lab/RECCON)

<sup>8</sup>[github.com/facebookresearch/minirts](https://github.com/facebookresearch/minirts)

<sup>9</sup>[github.com/google-research-datasets/dstc8-schema-guided-dialogue#dialogue-representation](https://github.com/google-research-datasets/dstc8-schema-guided-dialogue#dialogue-representation)

<sup>10</sup>[www.open-ease.org](http://www.open-ease.org)

<sup>11</sup>[github.com/Microsoft/psi](https://github.com/Microsoft/psi)

<sup>12</sup>[askforalfred.com](http://askforalfred.com)

centric vision to sequences of actions for household tasks (Shridhar et al., 2020). ALFRED releases challenges and leader boards for trajectory tasks in which a human instructs an agent through natural language to carry out specific household tasks in a virtual world. The task is grounded in situations and combines video, audio, and text for clear goals. There is no natural dialogue that is independent of the task.

The previous datasets all involve agents. Video recordings of interacting people can be seen as another source of data. MELD<sup>13</sup> and COSMIC<sup>14</sup> represent the Friends sitcom through videos, time-stamped dialogues and emotion annotations (Poria et al., 2019; Ghosal et al., 2020). Two seasons from the Friends dialogues were also annotated with person references and identities by (Choi and Chen, 2018) for the SemEval2018-task4 on and for Q&A on open dialogues (Yang and Choi, 2019). IEMOCAP<sup>15</sup> created detailed multimodal recordings of scripted human-human conversations with annotations of participants, gaze, gestures, etc. (Busso et al., 2008). Both datasets do not provide any further situated-references, and there is no specific goal set for the conversation. Other similar smaller datasets with audio-visual emotion expression are RAVDESS (Livingstone and Russo, 2018), TESS (Dupuis and Pichora-Fuller, 2010) and SAVEE (Haq and Jackson, 2010).

## 2.1 Annotation schemes

Interaction data come in different formats and follow different schemes. DiaML (Bunt et al., 2012) is a modeling language for the annotation of dialogues as a discourse. However, it does not tackle the grounding problem. It does not make use of identifiers that represent identities of entities, contexts, and situations independently of their mentions as it targets single modality data.

VOXML (Pustejovsky and Krishnaswamy, 2016) is a formal modeling language for capturing spatial semantics of object entities in 3D simulations. VOXML tackles grounding but does not model dialogue interaction nor the sequential alignment of cross-modality segments. It is a model for defining the semantics of linguistic expressions through physical world simulations, and it does not model this world per se independently of these expressions. Furthermore, it is a formal symbolic

<sup>13</sup>[affective-meld.github.io/](https://affective-meld.github.io/)

<sup>14</sup>[github.com/declare-lab/conv-emotion](https://github.com/declare-lab/conv-emotion)

<sup>15</sup>[sail.usc.edu/iemocap/](https://sail.usc.edu/iemocap/)

representation that relies on a fully descriptive relation between language expressions and situation modeling.

The Simple Event Model or SEM is a Resource Description Framework (RDF) model for capturing situations following semantic web principles (Van Hage et al., 2011). SEM represents situations as event instances through URIs, with actors, places, and temporal relations to OWL-Time objects<sup>16</sup> either defined as time points or as periods. Situations can be related as sequences in time series through OWL-Time grounding, as well as through explicit temporal and causal relations between events. SEM can be used to construct event-centric KGs rather than entity-centric KGs. Event-centric KGs are well-suited for representing temporal properties of situations and entities within these. Furthermore, they are not limited to the predefined properties of entity-centric graphs but exploit abstract event-participant relations that can be further modeled in additional ontologies (Segers et al., 2018).

The Grounded Representation and Source Perspective (GRaSP) model (Fokkens et al., 2017), augments SEM with *grasp:denotes* relations between linguistic expressions (so-called mentions) and their referential identities. Through GRaSP, any segment in a signal (verbal, audio, or video) can be mapped to an instance in a SEM model, as such providing a flexible framework for referential grounding. Although SEM and GRaSP can be used for any modality, they have mostly been used for representing events in text (Vossen et al., 2016). In (Vossen et al., 2018, 2019a), we have shown that GRaSP can also be used for modeling multimodal situations with unaligned signals to model a theory-of-mind or ToM in which different modalities and different sources can generate alternative interpretations that can co-exist in the robot’s eKG. Similar to (Kondratyuk and Kennington, 2017), our robot eKG reflects the episodic accumulation of knowledge through interaction over time. Our model differs from theirs in that our model allows for alternative facts and properties.

DIAML, VOXML and GRaSP are complex XML representations. Most of the DSTC datasets, however, follow a more basic schema in JSON format. Data elements represent sequences, possibly including time stamps and pointers to separate media files, possibly including bounding box

<sup>16</sup>[www.w3.org/TR/owl-time/](https://www.w3.org/TR/owl-time/)

coordinates, scene interpretations, participants, utterances, and goals to achieve. Google’s Schema-Guided-Dialogue, Amazon’s Alexa Topical-Chat and Facebook’s SIMMC provide comparable JSON formats for capturing simple situations, goals, participants and dialogues with communication.

## 2.2 Our contribution

Although we can use many aspects of the previously discussed data and representation models, none of these completely provide what we need for modeling streaming data in different modalities representing parallel sequences of signals within physical or virtual world contexts while allowing for alternative interpretations. Most of the described data and models do not entirely represent the contextual situation in which conversations are embedded. Furthermore, they lack the means to represent the cumulative result of interpreting streams of multimodal signals over time.

In our representation, we combine the best of two worlds. On the one hand, we use light-weight JSON-LD<sup>17</sup> for data in different modalities, their segmentation, alignment, and annotations, which provides us with the flexibility to easily represent any data streams; on the other hand, we use RDF to represent interpretations of such multimodal situations and the referential relations to explicit identities following GRaSP, which allows us to reason over the data in a robust way. Likewise, our framework can model the interpretation from raw signals to interpreted segments up to the situation-centric aggregation of triples over time as an episodic memory.

Our approach to connecting multimodal situated interactions to an explicit Knowledge Graph comes close to openEASE (Beetz et al., 2015), except that we focus on complex *reference* relations between conversations and situations rather than on robot *actions* only. Furthermore, we follow a data-centric approach in which an open representation of the interaction forms the basis for sharing data, tools, and solutions, whereas interactions in openEASE can only be accessed through queries. In contrast to the episodic triples from conversation generated by (Kondratyuk and Kennington, 2017) and the episodic NEEMS of openEASE, our eKG incorporates a ToM model based on SEM and GRaSP which supports reasoning over conflicting information, knowledge gaps and uncertainties across

modalities and sources.

Nevertheless, EMISSOR is not restricted to a specific annotation scheme nor a specific formal model of situations. The use of JSON-LD allows seamless integration with any eKG, i.e. from raw signal, to annotation, to explicit symbolic representation. Our EMISSOR platform supports converting different datasets to a shared model that imposes alignments of spatial and temporal signals within a referential grounding framework and an episodic memory with understanding. The framework combines modalities, generates segmentation for each, and creates referential grounding and corresponding triple representations that capture identities and relations. We furthermore provide an annotation tool to create gold annotations for referential grounding, both from dynamic streams and manually designed scenarios using controlled static data.

## 3 EMISSOR: design and specification

In this section, we describe the different data layers of our model and their interrelations. We summarise the design desiderata for our model as follows:

1. Support parallel unaligned streams of multimodal signals
2. Detect sequences of segments within signals grounded in time and space
3. Allow segment alignment, overlap, and disjointedness across modalities
4. Model situated-references in segments to unique identities in a Knowledge Graph
5. Model causal coherent relations between subsets of segments across modalities
6. Model physical and virtual real-world contexts
7. Model streams of signals as a transparent cumulative symbolic interpretation of experiences
8. Provide an episodic memory of situated references and interpretations in a knowledge graph that supports reasoning

To meet these requirements, we use GRaSP as a referential framework to connect segments in signals to identities in an eKG. These identities are individual people, objects, and places, or relations and properties of these. In the former case, identities are represented as instances through their URIs, possibly with names as labels and instances of a

<sup>17</sup>[json-ld.org/](http://json-ld.org/)

particular type. In the latter case, constellations of instances are interpreted as relations or properties represented by RDF triples. In addition, media streams are stored as separate data files for each modality. Separate JSON-LD files are provided for each modality that 1) divide this modality signal into segments grounded in time and space and 2) annotate these segments as representing identities and relations registered in the eKG. A stream of media interpretations emitting triples over time then results in the cumulative growth of an eKG representing an episodic memory. With the use of JSON-LD elements, signal metadata can be mapped to our referential framework and included in the eKG.

**Motivating example** Let us consider a simple example (1). A face detector module detects a human face in a video frame at time  $t1$  which results in a box segment. Face recognition cannot recognize this person, so it is identified as a new instance of the type *PERSON*. Next, at  $t2$ , another face is detected as a box segment. This face is identified as a known person in the eKG: a URI with the name ‘Carl’. At  $t3$  a speaker is detected whose voice is mapped to the same identity of “Carl”. The speaker says “Do you know my daughter Carla?”. A text understanding module processes the text representation of the audio at  $t3$ . Words such as “I” and “me” are linked to the same URI as the speaker, while “Carla” will be mapped to another identity (not yet visually grounded) and related to “Carl” as his daughter. Through contextual reasoning, the model may ground the first unidentified face perception to this newly created identity of “Carla” in hindsight. Let us consider an alternative variation on this scenario in which the detected speaker at  $t3$  is mapped to “Alice” rather than “Carl”. Alice says, “That is Carl and his daughter Carla”. The deictic references to “Carl” and “Carla” and the pronominal reference “his” can only be resolved by combining the earlier perceptions with the semantics of this utterance. On the other hand, in both scenarios, “being a daughter” is knowledge that cannot be concluded from visual and audio signals and is solely conveyed by interpreting the semantics of the utterance. This demonstrates that neither audio-visual nor textual segments contain all the information needed to come to the correct interpretation: language utterances due to their referential nature rather than being descriptive, audio-visual segments due to their limitation to signal social, conceptual, and cultural framing.

```
video, t1 -> segment -> [PERSON]
video, t2 -> segment -> URI(:Carl)
audio, t3 -> segment -> URI(:Carl)
text, t3 -> segment -> URI(:Carl)
text, t3 -> segment -> URI(:Carla)
text, t3 -> segment -> triple
    (:Carl :daughter :Carla)
reinterpretation:
video, t1 -> segment -> URI(:Carla)
```

Example 1: Carl - Carla scenario

### 3.1 Model description

Figure 1 shows an overview of our representation’s different data layers in terms of an entity-relationship model. For grounding data, we define different layers for segments, rulers, and containers. Segments can have complex causal coherence relations across modalities. Since the segmentation of these modalities is done separately, we need temporal and spatial containers to order and connect segments across modalities, following the principles of TimeML (Pustejovsky and Stubbs, 2011) and VOXML. The temporal and spatial containers form the basis for constituting potential causal relations (forward and backward and across distances). Therefore, each container will consist of a ruler that defines the granularity of segments across modalities (e.g., a sequence or region). The ruler positions the segments relative to each other and makes them conditional for defining relations and for predictive models.

A scenario (bottom left) is an instance of a context in a specific time and space and acts as a container for parallel streams of multimodal signals, which are divided into segments. Scenarios can have specific attributes to qualify them, including names, overall scenario type and location, the purpose or intention of (specific) participants. When segments get annotated, mentions (linguistic) or perceptions (video or audio) are created, pointing to one or more segments and an annotation value. Annotations can be added freely and there are no restrictions on the values for annotations in the JSON structure. Figure 1, shows a few examples of typical annotation values such as *Face* for boxes in images, *Tokens* in texts or *NamedEntity* for Named Entity expressions. However, EMISOR additionally uses JSON-LD to also support the direct linking of interpretations to the eKG in which people, objects and situations are modeled through explicit URIs. We therefore allow explicit URIs as annotation values to ground the segments to these identities. In that case, a box segment, a

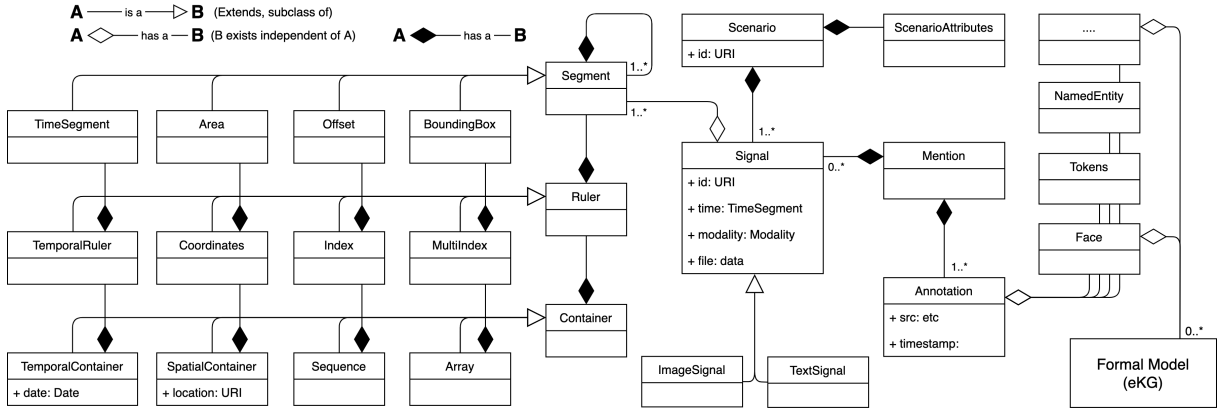


Figure 1: Entity-relationship overview of data elements and relations within EMISSOR. The right-side box functions as a placeholder for any formal model of situations that can be linked to the annotations of segments in the multimodal data streams. We assume here that identities in these models can be defined according to any set of ontologies to reason over the interpretations.

name or pronoun in the text is annotated with the unique URI of a person rather than a conceptual label as a value.

In Figure 1 due to limits of space, we only show a place holder for the eKG as a formal model of situations at the right side. This place holder stands for any ontological model and its population with instances. In practice, we use a wide range of ontologies to model situations as populations of instances, as described in (Vossen et al., 2019b). Note that EMISSOR allows for both types of annotations next to each other, e.g. a segment can be annotated as a human face without a specific identity but the same segment can have an additional annotation with the URI from the eKG.

**Scenario structure** We consider an interaction as a scenario. Scenarios are organized in folders on disk. Within a scenario folder, we store the source data as separate files in a modality subfolder, e.g. text, video, image, audio. Furthermore, one JSON-LD file per modality defines the metadata and the segments present for each signal in the modality. In addition to these segments, there may be lists of mentions or perceptions as JSON-LD elements. Each mention or perception specifies a range of segments (at least one) and the interpretations as annotations representing the instances and concepts in Figure 1. Next to the modality JSON-LD file, a specific folder contains the RDF triples extracted from the annotated signals. For example, an utterance in a conversation may mention somebody’s age, which yields an RDF triple with the person’s URI as the subject, the *has-age* property, and the actual age as a value.

Finally, there is a separate JSON file with meta-data on the complete scenario. This scenario JSON defines the temporal and spatial ruler within which the scenario is located (date, begin and end time, geo-location, place-name), the interacting participants (e.g., the agent and the human speaker(s)), and any other people and objects that participate in the scene. The specification of participants and props can be based on the instances from the eKG. This scenario JSON file has the same name as the folder name of the scenario.

**Modalities** The different modalities are represented in parallel streams of signals that are aligned by temporal and spatial rulers in the containers (Figure 2). We currently support text, audio, and visual modalities enriched by the knowledge layer as extracted by annotations.

Within each modality, a signal is broken down into *segments* positioned relative to the temporal and spatial ruler through begin and end points or box coordinates respectively. The granularity can vary but depends on the minimal unit of the rulers. Figure 2 shows an example of a scenario with layers for these four modalities with a temporal ruler on the horizontal axis. In this scenario, a person, “Carl”, tells a robot, “Leolani”, that he cannot find a pillbox. The robot spots the box under the table and communicates this to Carl, who confirms finding it. At every turn in the conversation, we see interaction data as segments (bars) aligned through the temporal ruler and its corresponding subgraphs generated from the interpretations as added to the eKG. The triples in the subgraph not only contain the representations for the participants and the pill-

box but also representations of their mentions in the lower layers by specific sources. These mentions specify the offsets and box coordinates in the segments so that the graph can be related to the signal in time. The graphs also show that sources express denial and uncertainty through mentions. Our ToM model supports reasoning over the status of the triples derived from the signals: who said what, what sensor perceived what, etc.

Figure 2 furthermore shows that segments are not necessarily fully aligned but are always temporally ordered. At the second turn, the agent first perceives a pillbox under the table (white bar in the ruler), after that the agent reports this in audio and text (gray and black bar in the ruler).

It is possible to represent a scenario from recordings, as was done for the CarLani scenario in Figure 2, but also to create these manually by simply adding series of audiovisual and/or text content as files to the scenario folder. However, every data representation needs to have a temporal ruler to ground all units in each modality to the same time period, which needs to be done in the corresponding JSON file for each modality. These JSON files can be created through scripts and an annotation tool.

**Annotations and mentions** Any segment can be annotated, which results in mentions or perceptions added to the JSON-LD file for the specific media. Mentions define a relation between text segments and interpretations, whereas perceptions relate audiovisual segments to interpretations. Each annotation has the following attributes: 1) type: kind of annotation; 2) value: the actual interpretation (e.g. label, reference); 3) source: software or person that created the annotation; 4) timestamp: when the annotation was created. We can have any number of segments with any number of annotations defined per mention/perception. Furthermore, annotations can be added on top of other annotations, following the Layered Annotation Framework (Ide and Romary, 2007). Finally, annotations of segments in different modalities with the same identifier will automatically create cross-modality co-reference; for example in the CarLani scenario, perception of the pillbox and its mention in the utterances are mapped to the same instance URI in the eKG.

**Identities, properties, and relations** The core idea is to create a mapping between a segment of a *signal* (e.g., a bounding box in an image or an

offset position and length in a text) and the *signal's interpretation* (e.g., a person's face or a pronoun). Through referential grounding, we generate triples expressing properties and relations across different modalities. Shared identifiers (URIs) aggregate these properties and relations, resulting in a world model over time. In Figure 2, this is demonstrated by the sequence of subgraphs at the top showing different states of interpretation going from lack of knowledge about the location of the pillbox (negative polarity) to the perception and having it in possession (positive polarity). The triples stored in an eKG likewise reflect this accumulation over time, while each triple is still grounded to a segment in a modality.

As explained in previous work (Vossen et al., 2018, 2019a), our framework focuses on storing information related to episodic experiences and their interpretations as perspectives that agents have. By nature, our framework is flexible in dealing with incomplete or contradicting information and can reason over knowledge with uncertainty while considering the sources' trustworthiness. For the scenario in Figure 2, the model represents two realities at the same time point: "Carl" not knowing the location and "Leolani" knowing the location. Querying the model for the location of the pillbox at that time generates an answer according to "Leolani". Reasoning is thus not only used to derive knowledge or answer factual questions, but also to evaluate the quality of the knowledge itself. Agents can use such qualitative evaluations to formulate strategies and actions to improve knowledge states.

Following the principles of Linked Data, our framework reuses existing ontologies for provenance (PROV-O), text processing (NAF), event (SEM) and perspective modeling (GRaSP, GAF). The usage of RDF allows us to integrate information from other existing open Knowledge Graphs, e.g. WikiData or DBpedia, to include prior knowledge. Using JSON-LD elements in our representation enables us to directly attach the referential grounding to the eKG by mapping elements of our JSON structure described in Figure 1 to elements of the underlying ontologies of the eKG without losing the lightweight representation of plain JSON.

## 4 Conclusions

In this paper, we described eight desiderata for the representation of multimodal interactions in collaborative contexts. We argued that existing



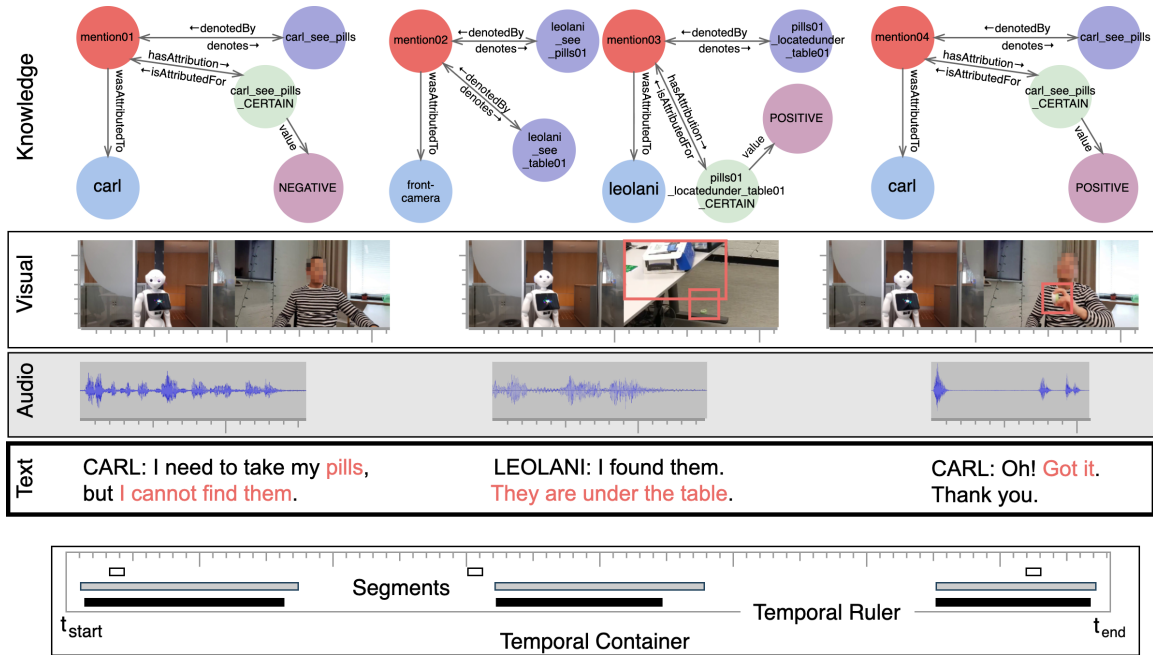


Figure 2: Visualization of four modalities (text, audio, visual, and knowledge) from the CarLani scenario. Signals are grounded in a temporal container on the horizontal axis, with bars marking alignments through the temporal ruler. Red boxes mark *segments* annotated as *mentions* of objects (pills and the table). Text *segments* highlighted in red are annotated as *mentions* of triples. The upper graphs represent corresponding triples from eKG generated from the annotated source modalities along the temporal sequence. The visual modality shows two different camera viewpoints (left is what Carl sees and right is what Leolani sees) concatenated side by side.

representations do not satisfy all these desiderata and therefore presented EMISSOR for referential interpretations of multimodal interactions to yield a Knowledge Graph as an explicit episodic memory of the experiences (eKG). EMISSOR combines light-weight JSON-LD representations for sequential media with semantic web-based RDF models of interpreted worlds. Through this we model cumulative growth of knowledge and information in the eKG as a result of processing multimedia streams over time. EMISSOR is designed to address all eight desiderata. It enables to create and compare recordings, annotations and interpretations of interactions in real-world contexts. This allows researchers to more easily share experiments and compare the interactions across different experiments, regardless of the specifics of agent systems or humans that participate in the experiment. Our model, software and converted data sets are available<sup>18</sup> according to the Apache open source license. Our release includes an annotation tool to create scenarios manually or inspect and annotate any recording of data.

As next steps, we develop more tooling and automatic linking of multimodal segments to identities.

<sup>18</sup><https://github.com/cltl/EMISSOR>

We will also provide more data either by converting existing public data to our framework or by rendering data through our robot platform. So far we focused on grounding segments in temporal containers but not yet in spatial containers. Our interactions do not include motion and navigation. In future work, we hope to include spatial grounding and reasoning. Finally, we will include an evaluation framework for analysing system performance in relation to 1) qualitative properties of the interaction, 2) goals and intentions following a reinforcement learning approach and 3) by evaluating the resulting eKG.

## 5 Acknowledgements

This research was funded by the Vrije Universiteit Amsterdam and the Netherlands Organisation for Scientific Research (NWO) via the *Spinoza* grant (SPI 63-260) awarded to Piek Vossen, the *Hybrid Intelligence Centre* via the *Zwaartekracht* grant (024.004.022) and the research project (9406.D1.19.005) *Communicating with and Relating to Social Robots: Alice Meets Leolani*.

## References

- Huda Alamri, Vincent Cartillier, Abhishek Das, Jue Wang, Anoop Cherian, Irfan Essa, Dhruv Batra, Tim K. Marks, Chiori Hori, Peter Anderson, Stefan Lee, and Devi Parikh. 2019. Audio-visual scene-aware dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Michael Beetz, Moritz Tenorth, and Jan Winkler. 2015. Open-ease. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1983–1990. IEEE.
- Dan Bohus, Sean Andrist, and Mihai Jalobeanu. 2017. Rapid development of multimodal interactive systems: a demonstration of platform for situated intelligence. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pages 493–494.
- Hennie Brugman, Albert Russel, and Xd Nijmegen. 2004. Annotating multi-media/multi-modal resources with elan. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*.
- Harry Bunt, Michael Kipp, and Volha Petukhova. 2012. Using diaml and anvil for multimodal dialogue annotations. In *Eight International Conference on Language Resources and Evaluation (LREC'12)*, pages 1301–1308.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42(4):335–359.
- Jinho D. Choi and Henry Y. Chen. 2018. [SemEval 2018 task 4: Character identification on multiparty dialogues](#). In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 57–64, New Orleans, Louisiana. Association for Computational Linguistics.
- Paul A Crook, Shivani Poddar, Ankita De, Semir Shafi, David Whitney, Alborz Geramifard, and Rajen Subba. 2019. Simmc: Situated interactive multimodal conversational data collection and evaluation platform. *arXiv e-prints*, pages arXiv–1911.
- Kate Dupuis and M Kathleen Pichora-Fuller. 2010. Toronto emotional speech set (tess). Toronto: University of Toronto, Psychology Department.
- Antske Fokkens, Piek Vossen, Marco Rospocher, Rinke Hoekstra, Willem R van Hage, and Fondazione Bruno Kessler. 2017. Grasp: grounded representation and source perspective. *Proceedings of Knowledge Resources for the Socio-Economic Sciences and Humanities associated with RANLP*, 17:19–25.
- Deepanway Ghosal, Navonil Majumder, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. Cosmic: Commonsense knowledge for emotion identification in conversations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 2470–2481.
- S. Haq and P.J.B. Jackson. 2010. *Machine Audition: Principles, Algorithms and Systems: Principles, Algorithms and Systems*, chapter Multimodal Emotion Recognition. IGI Global.
- Hengyuan Hu, Denis Yarats, Qucheng Gong, Yuan-dong Tian, and Mike Lewis. 2019. Hierarchical decision making by generating and following natural language instructions. In *Advances in neural information processing systems*, pages 10025–10034.
- Nancy Ide and Laurent Romary. 2007. Towards international standards for language resources nancy ide and laurent romary. In *Evaluation of text and speech systems*, pages 263–284. Springer.
- Michael Kipp. 2001. Anvil-a generic annotation tool for multimodal dialogue. In *Seventh European Conference on Speech Communication and Technology*.
- Dan Kondratyuk and Casey Kennington. 2017. Towards a dialogue system with long-term, episodic memory. *SEMDIAL 2017 SaarDial*, page 160.
- Alan M Leslie. 1987. Pretense and representation: The origins of "theory of mind.". *Psychological review*, 94(4):412.
- Steven R Livingstone and Frank A Russo. 2018. The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english. *PLoS one*, 13(5):e0196391.
- Alexander Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. Parlai: A dialog research software platform. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 79–84.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. Meld: A multimodal multi-party dataset for emotion recognition in conversations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 527–536.
- Soujanya Poria, Navonil Majumder, Devamanyu Hazarika, Deepanway Ghosal, Rishabh Bhardwaj, Samson Yu Bai Jian, Romila Ghosh, Niyati Chhaya, Alexander Gelbukh, and Rada Mihalcea. 2020. Recognizing emotion cause in conversations. *arXiv preprint arXiv:2012.11820*.
- David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4):515–526.

- James Pustejovsky and Nikhil Krishnaswamy. 2016. Voxml: A visualization modeling language. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 4606–4613.
- James Pustejovsky and Amber Stubbs. 2011. Increasing informativeness in temporal annotation. In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 152–160.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8689–8696.
- Roxane Segers, Tommaso Caselli, and Piek Vossen. 2018. The circumstantial event ontology (ceo) and ecb+/ceo; an ontology and corpus for implicit causal relations between events. In *Proceedings of the 11th edition of the Language Resources and Evaluation Conference (LREC2018), Miyazaki, Japan, May 7-12*.
- Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2018. A survey of available corpora for building data-driven dialogue systems: The journal version. *Dialogue & Discourse*, 9(1):1–49.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Willem Robert Van Hage, Véronique Malaisé, Roxane Segers, Laura Hollink, and Guus Schreiber. 2011. Design and use of the simple event model (sem). *Journal of Web Semantics*, 9(2):128–136.
- Piek Vossen, Rodrigo Agerri, Itziar Aldabe, Agata Cybulska, Marieke van Erp, Antske Fokkens, Egoitz Laparra, Anne-Lyse Minard, Alessio Palmero Aprosio, German Rigau, et al. 2016. Newsreader: Using knowledge resources in a cross-lingual reading machine to generate more knowledge from massive streams of news. *Knowledge-Based Systems*, 110:60–85.
- Piek Vossen, Selene Baez, Lenka Bajčetić, and Bram Kraaijeveld. 2018. Leolani: a reference machine with a theory of mind for social communication. In *International conference on text, speech, and dialogue*, pages 15–25. Springer.
- Piek Vossen, Selene Baez, Lenka Bajčetić, Suzana Bašić, and Bram Kraaijeveld. 2019a. Leolani: A robot that communicates and learns about the shared world. In *2019 ISWC Satellite Tracks (Posters and Demonstrations, Industry, and Outrageous Ideas), ISWC 2019-Satellites*, pages 181–184. CEUR-WS.
- Piek Vossen, Lenka Bajčetić, Selene Báez Santamaria, Suzana Bašić, and Bram Kraaijeveld. 2019b. Modelling context awareness for a situated semantic agent. In *Proceedings of 11th International and Interdisciplinary Conference on Modeling and Using Context, CONTEXT 2019*.
- Zhengzhe Yang and Jinho D. Choi. 2019. FriendsQA: Open-domain question answering on TV show transcripts. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 188–197, Stockholm, Sweden. Association for Computational Linguistics.

## Appendix A: Annotation tool and example data sets

### A Example Data Sets

Any multimodal interaction data can be represented and annotated in the EMISSOR annotation format with minimal effort. We released EMISSOR representations and annotations of some popular public data sets (e.g. MELD and IEMOCAP), together with the scripts to convert them from their original formats. In the near future, we will add scripts for other popular data sets. In addition to the conversion scripts, we also created scripts for segmentation of modalities, such as bounding boxes for objects, faces and text tokenization with named entity detection. An additional baseline script resolves the identities of faces and entities against an eKG by selecting the first matching name. These scripts prepare any video recording or collection of multimedia data for annotation in the annotation tool described in Section B. In the future, we replace the baseline scripts with SOTA modules for resolving referential ambiguity.

We also created our own data set called CarLani directly rendered from interacting with our robot platform. Figure 2 shows an example dialogue of three utterances from this data set between the human (Carl) and the robot (Leolani), assuming the context of a care taking robot in an elderly home.

By running interactions through our robot platform with humans in a physically perceived world, the multimodal data is automatically grounded in the knowledge graph according to the EMISSOR framework. This will automatically generate rich referential relations between mentions and perceptions with identities, within a functional communicative contexts. These can be analysed, evaluated and adapted to gold annotations for training and testing.

### B Annotation Tool

Along with the proposed data representation, we are developing a GUI tool capable of reading EMISSOR data representations (with or without annotations). The purpose of the tool is a first inspection of data sets by providing a comprehensible visualisation of the signals in different modalities, their grounding to the temporal (and spatial) containers, as well as their interpretations, including segment alignments, situated references and explicit semantic representations. Second, it allows modification of the aforementioned properties for a given data set, e.g. to add gold annotations, perform corrections or add additional interpretations. Third, gold scenarios for a given task or problem can be created manually from scratch without the need for an actual agent implementation.

Besides conversion issues to other data representations<sup>19</sup>, existing tools like e.g. Anvil<sup>20</sup> (Kipp, 2001) or Elan<sup>21</sup> (Brugman et al., 2004) only ground the conversation to speakers, audio, faces, gestures but do not ground referential expressions to the situation. Our tool focuses on segmentation and grounding to mediate between the media data and the identities in the Knowledge Graph.

The current version of the tool supports image, audio, and text as modalities in a scenario, allows to add and remove signals to them and to position (ground) the signals on (to) the timeline of the scenario. In any situation, it is possible to create segments and annotations automatically or manually. On image signals, rectangular segments (bounding boxes) can be defined manually and annotated. Alternatively, during data preparation, boxing scripts can be used to automatically generate bounding boxes beforehand. Text signals can be automatically tokenized using scripts beforehand as well. The tool then allows for token selections to be annotated. The tool also provide choices for reference linking to known (listed) entities as annotation values. These entities can be taken from the eKG or from any other registration. In addition, you can create new identities directly in the tool through annotations, as well create triples that express properties or relations between entities. The URIs and triples created during the annotation can then be added as gold knowledge to the eKG a posteriori.

Finally, the annotation tool can be used to create scenarios manually in a very controlled way. Researchers can store images and conversations in the corresponding media folders manually and next use

---

<sup>19</sup>Anvil supports DIAML and stores conversational units in sequences with name identifiers and time stamps for an associated video file. ELAN stores conversations in EAF (Eudico Annotation Format), which is a propriety XML format.

<sup>20</sup>[www.anvil-software.org](http://www.anvil-software.org)

<sup>21</sup>[www.mpi.nl/corpus/html/elan/](http://www.mpi.nl/corpus/html/elan/)




Figure 3: Tool for visualizing and annotating EMISSOR

the tool to place them in the proper order. In the near future, we will add a function to play such a scenario as well following the temporal specification. Currently, the user can play it by moving forward manually.

## Appendix B: CarlLani Example Data from redacted source

In the following we include an excerpt of the CarlLani data set as the original source was redacted from the submission for anonymization purposes. For space and readability reasons *text.json* and *image.json* metadata files are shortened by removing part of the signals and/or mentions. Also the context referenced in the JSON-LD *@context* element is included.

### Scenario structure

```

|- carl-robot/
  |- audio/
    |- carl-robot-000_frame0_0.wav
    |- carl-robot-000_frame30_1000.wav
    |- carl-robot-000_frame60_2000.wav
    |- ....
  |- image/
    |- carl-robot-000_frame0_0.jpg
    |- carl-robot-000_frame30_1000.jpg
    |- carl-robot-000_frame60_2000.jpg
    |- ....
  |- rdf/
    |- episodic_memory.trig
    |- statement1.trig
    |- objectdetection1.trig
    |- statement2.trig
    |- statement3.trig
  |- text/
    |- carl-robot.csv
  |- video/
    |- carl-robot-000_frame0_0.mp4
    |- carl-robot-000_frame30_1000.mp4
    |- carl-robot-000_frame60_2000.mp4

```

```
|- ....
|- carl-robot.json
|- audio.json
|- image.json
|- text.json
|- video.json
```

### carl-robot.csv

```
speaker,utterance,time
Carl,"I need to take my pills, but I cannot find them.",0
Leolani,"I found them. They are under the table.",3933
Carl,"Oh! Got it. Thank you.",7133
```

### carl-robot.json

```
1 {
2   "@context" : "http://emissor.org/jsonldcontext.jsonld",
3   "type": "Scenario",
4   "id": "carl-robot",
5   "context": {
6     "agent": "robot_agent",
7     "objects": [],
8     "persons": [],
9     "speaker": {
10      "@context" : "http://schema.org/docs/jsonldcontext.jsonld",
11      "id": "bc913d64-a597-4876-a3fe-fe47472cd274",
12      "type": "Person",
13      "birthdate": "1995-04-09T20:00:00Z",
14      "gender": "Male",
15      "name": "Carl"
16    }
17  },
18  "ruler": {
19    "type": "TemporalRuler",
20    "container_id": "carl-robot",
21    "end": 11133,
22    "start": 0
23  },
24  "signals": {
25    "image": "./image.json",
26    "text": "./text.json"
27  }
28 }
```

## image.json (excerpt)

```
1 [{
2   "@context" : "http://emissor.org/jsonldcontext.jsonld",
3   "type": "ImageSignal",
4   "id": "21830691-4410-45f2-b611-f61cb4dbc0de",
5   "files": [
6     "image/car1-robot-000_frame0_0.jpg"
7   ],
8   "modality": "image",
9   "time": {
10    "type": "TimeSegment",
11    "container_id": "car1-robot",
12    "start": 0,
13    "end": 33
14  },
15  "ruler": {
16    "type": "MultiIndex",
17    "container_id": "21830691-4410-45f2-b611-f61cb4dbc0de",
18    "bounds": [0, 0, 3840, 1080]
19  },
20  "mentions": [
21    {
22      "type": "Mention",
23      "id": "54920da9-41d4-421e-b3f4-7955e71f053a",
24      "annotations": [
25        {
26          "type": "Annotation",
27          "source": "machine",
28          "timestamp": 0,
29          "type": "person",
30          "value": {
31            "type": "Face",
32            "instance": {
33              "@context" : "http://schema.org/docs/jsonldcontext.
34                jsonld",
35              "id": "bc913d64-a597-4876-a3fe-fe47472cd274",
36              "type": "Person",
37              "birthDate": "1995-04-09T20:00:00Z",
38              "gender": "Male",
39              "name": "Speaker"
40            },
41            "age": 23,
42            "gender": "male",
43            "faceprob": 1.0
44          }
45        }
46      ],
47      "segment": [
48        {
49          "type": "BoundingBox",
50          "container_id": "21830691-4410-45f2-b611-f61cb4dbc0de",
```

```

50     "bounds": [2830, 241, 3034, 521]
51     }
52   ]
53 }
54 ]},
55
56
57
58
59
60 {
61   "@context" : "http://emissor.org/jsonldcontext.jsonld",
62   "type": "ImageSignal",
63   "id": "88a31791-4410-45f2-b611-f61cb4d321ff",
64   "files": [
65     "image/carl-robot-000_frame30_1000.jpg"
66   ],
67   "modality": "image",
68   "time": {
69     "type": "TimeSegment",
70     "container_id": "carl-robot",
71     "start": 1000,
72     "end": 1033
73   },
74   "ruler": {
75     "type": "MultiIndex",
76     "container_id": "88a31791-4410-45f2-b611-f61cb4d321ff",
77     "bounds": [0, 0, 3840, 1080]
78   },
79   "mentions": [
80     {
81       "type": "Mention",
82       "id": "92af1ea9-41d4-421e-b3f4-7955e71a1a97",
83       "annotations": [
84         {
85           "type": "Annotation",
86           "source": "machine",
87           "timestamp": 1000,
88           "type": "person",
89           "value": {
90             "type": "Face",
91             "instance": {
92               "@context" : "http://schema.org/docs/jsonldcontext.
93                 jsonld",
94               "@id": "bc913d64-a597-4876-a3fe-fe47472cd274",
95               "type": "Person",
96               "birthDate": "1995-04-09T20:00:00Z",
97               "gender": "Male",
98               "name": "Speaker"
99             },

```



```
100         "gender": "male",
101         "faceprob": 1.0
102     }
103 }
104 ],
105 "segment": [
106     {
107         "type": "BoundingBox",
108         "container_id": "88a31791-4410-45f2-b611-f61cb4d321ff",
109         "bounds": [2831, 235, 3036, 514]
110     }
111 ]
112 }]]}, .....]
```

## text.json (excerpt)

```
1  {{
2    "@context": "http://emissor.org/jsonldcontext.jsonld",
3    "files": ["text/carl-robot.csv#0"],
4    "id": "85c27957-9b18-497e-9557-761b02bdbc21",
5    "mentions": [
6      {
7        "type": "Mention",
8        "id": "0d830564-ab25-4aac-82f6-f34fc61b0481",
9        "annotations": [
10       {
11         "source": "annotation_tool",
12         "timestamp": 1616442473,
13         "type": "token",
14         "value": {
15           "id": "b1ec4a11-cd35-4c10-be47-244147da1086",
16           "ruler": {
17             "container_id": "b1ec4a11-cd35-4c10-be47-244147da1086",
18             "type": "AtomicRuler"
19           },
20           "type": "Token",
21           "value": "I"
22         }
23       }
24     ],
25     "segment": [
26       {
27         "container_id": "85c27957-9b18-497e-9557-761b02bdbc21",
28         "start": 0,
29         "stop": 1,
30         "type": "Index"
31       }
32     ]
33   },
34   ...
35   {
36     "type": "Mention",
37     "id": "a930c234-f3f2-4932-a32d-bde0acc2aafd",
38     "annotations": [
39       {
40         "source": "annotation_tool",
41         "timestamp": 1616442473,
42         "type": "token",
43         "value": {
44           "id": "13d77c30-4f10-481a-b0c4-3b80532b038f",
45           "ruler": {
46             "container_id": "13d77c30-4f10-481a-b0c4-3b80532b038f",
47             "type": "AtomicRuler"
48           },
49           "type": "Token",
50           "value": "."
51         }
52       }
53     ],
54     "segment": [
55       {
56         "container_id": "85c27957-9b18-497e-9557-761b02bdbc21",
57         "start": 47,
58         "stop": 48,
59         "type": "Index"
60       }
61     ]
62   }
63 ],
64 "modality": "text",
65 "ruler": {
66   "container_id": "85c27957-9b18-497e-9557-761b02bdbc21",
67   "start": 0,
68   "stop": 48,
69   "type": "Index"
70 },
71 "seq": ["I", " ", "n", "e", "e", "d", " ", "t", "o", " ", "t", "a", "k", "e", " ", "m", "y", " ", "p", "i", "l", "l", "s", " ", " ",
72 " ", "b", "u", "t", " ", "I", " ", "c", "a", "n", "n", "o", "t", " ", "f", "i", "n", "d", " ", "t", "h", "e", "m", "." ],
73 "time": {
74   "container_id": "carl-robot",
75   "end": 0,
76   "start": 0,
77   "type": "TemporalRuler"
78 },
79 "type": "TextSignal"
80 },
81 ...
82
83
84
85
```

```

86
87
88
89
90 {
91   "@context" : "http://emissor.org/jsonldcontext.jsonld",
92   "files": [
93     "text/carl-robot.csv#2"
94   ],
95   "id": "2142b6d8-4cda-481b-a056-1b6d874da648",
96   "mentions": [
97     {
98       "type": "Mention",
99       "id": "c851ca48-81b6-44fe-a772-9f62840ca2f6",
100      "annotations": [
101        {
102          "source": "annotation_tool",
103          "timestamp": 1616442473,
104          "type": "token",
105          "value": {
106            "id": "d7770947-0be5-413f-9c1e-4e9d130e6a41",
107            "ruler": {
108              "container_id": "d7770947-0be5-413f-9c1e-4e9d130e6a41",
109              "type": "AtomicRuler"
110            },
111            "type": "Token",
112            "value": "Oh"
113          }
114        }
115      ],
116      "segment": [
117        {
118          "container_id": "2142b6d8-4cda-481b-a056-1b6d874da648",
119          "start": 0,
120          "stop": 2,
121          "type": "Index"
122        }
123      ]
124    },
125    ...
126    {
127      "type": "Mention",
128      "id": "e62ae54b-bbb4-4464-8796-fe1a5ce22fac",
129      "annotations": [
130        {
131          "source": "annotation_tool",
132          "timestamp": 1616442473,
133          "type": "token",
134          "value": {
135            "id": "fb7a3f36-11c4-486c-bd60-aedd4377bb7",
136            "ruler": {
137              "container_id": "fb7a3f36-11c4-486c-bd60-aedd4377bb7",
138              "type": "AtomicRuler"
139            },
140            "type": "Token",
141            "value": "."
142          }
143        }
144      ],
145      "segment": [
146        {
147          "container_id": "2142b6d8-4cda-481b-a056-1b6d874da648",
148          "start": 21,
149          "stop": 22,
150          "type": "Index"
151        }
152      ]
153    }
154  ],
155  "modality": "text",
156  "ruler": {
157    "container_id": "2142b6d8-4cda-481b-a056-1b6d874da648",
158    "start": 0,
159    "stop": 22,
160    "type": "Index"
161  },
162  "seq": ["O", "h", "!", " ", "G", "o", "t", " ", " ", "i", "t", ".", " ", " ", "T", "h", "a", "n", "k", " ", " ", "y", "o", "u", "." ],
163  "time": {
164    "container_id": "carl-robot",
165    "end": 7133,
166    "start": 10976,
167    "type": "TemporalRuler"
168  },
169  "type": "TextSignal"
170 }
171 ]

```

## JSON-LD context (<http://emissor.org/jsonldcontext.jsonld>)

```
1 {
2   "@context" : {
3     "@base" : "http://experiment.my/",
4     "@vocab" : "https://emissor.org/emissor#",
5     "type" : "@type",
6     "id" : "@id",
7     "emissor" : "http://emissor.org/emissor#",
8     "grasp" : "http://groundedannotationframework.org/grasp#",
9     "container_id" : {"@type" : "@id"},
10    "signal" : "@nest",
11    "Mention" : "grasp:Mention"
12  }
13 }
```

## statements2.trig

```
1 @prefix robotContext: <http://emissor.org/robot/context/> .
2 @prefix xml1: <https://www.w3.org/TR/xmlschema-2/#> .
3 @prefix owl: <http://www.w3.org/2002/07/owl#> .
4 @prefix wdt: <http://www.wikidata.org/prop/direct/> .
5 @prefix ceo: <http://www.newsreader-project.eu/domain-ontology#> .
6 @prefix gaf: <http://groundedannotationframework.org/gaf#> .
7 @prefix ns1: <urn:x-rdflib:> .
8 @prefix wd: <http://www.wikidata.org/entity/> .
9 @prefix grasp: <http://groundedannotationframework.org/grasp#> .
10 @prefix xml: <http://www.w3.org/XML/1998/namespace> .
11 @prefix grasps: <http://groundedannotationframework.org/grasp/sentiment#> .
12 @prefix sem: <http://semanticweb.cs.vu.nl/2009/11/sem/> .
13 @prefix prov: <http://www.w3.org/ns/prov#> .
14 ...
15 @prefix foaf: <http://xmlns.com/foaf/0.1/> .
16 @prefix wgs: <http://www.w3.org/2003/01/geo/wgs84_pos#> .
17 @prefix graspf: <http://groundedannotationframework.org/grasp/factuality#> .
18 @prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
19 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
20 @prefix grasp: <http://groundedannotationframework.org/grasp#> .
21
22
23 robotWorld:Instances {
24   robotWorld:lani a gaf:Instance, robotMu:robot ;
25                   rdfs:label "lani" .
26   robotWorld:pills a gaf:Instance, robotMu:object ;
27                   rdfs:label "pills" ;
28                   gaf:denotedIn robotTalk:chat1_utterance2_char0-39 .
29   robotWorld:pills-277239 a gaf:Instance, robotMu:object, robotMu:pills ;
30                   rdfs:label "pills-277239" ;
31                   robotMu:id "277239"^^xml1:string ;
32                   gaf:denotedIn robotTalk:visual1_detection2_pixel0-3 ;
33                   eps:hasContext robotContext:context212127 .
34   robotWorld:table a gaf:Instance, robotMu:object ;
35                   rdfs:label "table" ;
36                   gaf:denotedIn robotTalk:chat1_utterance2_char0-39 .
37   robotWorld:table-208510 a gaf:Instance, robotMu:object, robotMu:table ;
38                   rdfs:label "table-208510" ;
39                   robotMu:id "208510"^^xml1:string ;
40                   gaf:denotedIn robotTalk:visual1_detection2_pixel0-3 ;
41                   eps:hasContext robotContext:context212127 .
42 }
43
44 robotTalk:Interactions {
45   robotWorld:Netherlands a robotMu:location, sem:Place, robotMu:country ;
46                   rdfs:label "Netherlands" .
47   robotWorld:Gelderland a robotMu:location, sem:Place, robotMu:region ;
48                   rdfs:label "Gelderland" .
49   robotWorld:Apeldoorn a robotMu:location, sem:Place, robotMu:city ;
50                   rdfs:label "Apeldoorn" .
51   robotTalk:chat1 a sem:Event, grasp:Chat ;
52                   rdfs:label "chat1" ;
53                   robotMu:id "1"^^xml1:string ;
54                   sem:hasSubEvent robotTalk:chat1_utterance2 .
55   robotTalk:visual1 a sem:Event, grasp:Visual ;
56                   rdfs:label "visual1" ;
57                   robotMu:id "1"^^xml1:string ;
58                   sem:hasSubEvent robotTalk:visual1_detection2 .
59   robotTalk:chat1_utterance2 a sem:Event, grasp:Utterance ;
60                   rdfs:label "chat1_utterance2" ;
61                   robotMu:id "2"^^xml1:string ;
62                   sem:hasActor robotFriends:lani .
63   robotTalk:visual1_detection2 a sem:Event, grasp:Detection ;
64                   rdfs:label "visual1_detection2" .
65                   robotMu:id "2"^^xml1:string ;
66                   sem:hasActor robotInputs:front-camera .
67   robotInputs:front-camera a gaf:Instance, grasp:Source, sem:Actor, robotMu:sensor ;
68                   rdfs:label "front-camera" .
69   robotFriends:lani a robotMu:person, gaf:Instance, grasp:Source, sem:Actor ;
70                   rdfs:label "lani" .
71   robotContext:home a robotMu:location, sem:Place ;
72                   rdfs:label "home" ;
73                   robotMu:id "251375"^^xml1:string ;
74                   robotMu:in robotWorld:Netherlands, robotWorld:Gelderland, robotWorld:Apeldoorn .
75   robotContext:context212127 a eps:Context ;
76                   rdfs:label "context212127" ;
77                   robotMu:id "212127"^^xml1:string ;
78                   eps:hasDetection robotWorld:pills-277239, robotWorld:table-208510 ;
79                   sem:hasBeginTimeStamp robotContext:2021-03-12 ;
80                   sem:hasEvent robotTalk:chat1, robotTalk:visual1 ;
81                   sem:hasPlace robotContext:home .
82   robotContext:2021-03-12 a sem:Time, time:DateTimeDescription ;
83                   rdfs:label "2021-03-12" ;
84                   time:day "12"^^xml1:gDay ;
```

```

85     time: month "3"^^xml1:gMonthDay ;
86     time: unitType time: unitDay ;
87     time: year "2021"^^xml1:gYear .
88 }
89
90
91
92 robotWorld: Claims {
93   robotWorld: lani_sense_front-camera a gaf: Assertion, sem: Event ;
94   rdfs: label "lani_sense_front-camera" .
95   robotWorld: lani_know_lani a gaf: Assertion, sem: Event ;
96   rdfs: label "lani_know_lani" ;
97   owl: sameAs robotWorld: lani .
98   robotWorld: pills_locatedunder_table a gaf: Assertion, sem: Event ;
99   rdfs: label "pills_locatedunder_table" ;
100   gaf: denotedBy robotTalk: chat1_utterance2_char0-39 .
101   robotWorld: lani_see_pills -277239 a gaf: Assertion, sem: Event ;
102   rdfs: label "lani_see_pills -277239" ;
103   gaf: denotedBy robotTalk: visual1_detection2_pixel0-3 ;
104   eps: hasContext robotContext: context212127 .
105   robotWorld: lani_see_table -208510 a gaf: Assertion, sem: Event ;
106   rdfs: label "lani_see_table -208510" ;
107   gaf: denotedBy robotTalk: visual1_detection2_pixel0-3 ;
108   eps: hasContext robotContext: context212127 .
109 }
110
111 robotTalk: Perspectives {
112   robotTalk: chat1_utterance2_char0-39 a gaf: Mention, grasp: Statement ;
113   rdfs: label "chat1_utterance2_char0-39"
114   rdf: value "I found them. They are under the table."^^xml1:string .
115   prov: wasDerivedFrom robotTalk: chat1_utterance2 ;
116   gaf: denotes robotWorld: pills_locatedunder_table ;
117   gaf: containsDenotation robotWorld: pills, robotWorld: table ;
118   grasp: wasAttributedTo robotFriends: lani ;
119   grasp: hasAttribution robotTalk: pills_locatedunder_table_CERTAIN -POSITIVE-NEUTRAL-NEUTRAL .
120   robotTalk: visual1_detection2_pixel0-3 a gaf: Mention, grasp: Experience ;
121   rdfs: label "visual1_detection2_pixel0-3" ;
122   prov: wasDerivedFrom robotTalk: visual1_detection2 .
123   gaf: denotes robotWorld: lani_see_pills -277239, robotWorld: lani_see_table -208510 ;
124   gaf: containsDenotation robotWorld: pills -277239, robotWorld: table -208510 ;
125   grasp: wasAttributedTo robotInputs: front-camera ;
126   grasp: hasAttribution robotTalk: pills_locatedunder_table_PROBABLE .
127   robotTalk: pills_locatedunder_table_CERTAIN -POSITIVE-NEUTRAL-NEUTRAL a grasp: Attribution ;
128   rdfs: label "pills_locatedunder_table_CERTAIN -POSITIVE-NEUTRAL-NEUTRAL" ;
129   rdf: value graspf: CERTAIN, graspf: POSITIVE, graspe: NEUTRAL, grasps: NEUTRAL ;
130   grasp: isAttributionFor robotTalk: chat1_utterance2_char0-39 .
131   robotTalk: pills_locatedunder_table_PROBABLE a grasp: Attribution ;
132   rdfs: label "pills_locatedunder_table_PROBABLE" ;
133   rdf: value graspf: PROBABLE ;
134   grasp: isAttributionFor robotTalk: visual1_detection2_pixel0-3 .
135   graspe: NEUTRAL a grasp: AttributionValue, graspe: EmotionValue .
136   grasps: NEUTRAL a grasp: AttributionValue, grasps: SentimentValue .
137   graspf: CERTAIN a grasp: AttributionValue, graspf: CertaintyValue .
138   graspf: POSITIVE a grasp: AttributionValue, graspf: PolarityValue .
139   graspf: PROBABLE a grasp: AttributionValue, graspf: CertaintyValue .
140 }
141
142 robotWorld: lani_know_lani {
143   robotWorld: lani robotMu: know robotFriends: lani .
144 }
145
146 robotWorld: lani_sense_front-camera {
147   robotWorld: lani robotMu: sense robotInputs: front-camera .
148 }
149
150 robotWorld: pills_locatedunder_table {
151   robotWorld: pills robotMu: locatedUnder robotWorld: table .
152 }
153
154 robotWorld: lani_see_pills -277239 {
155   robotWorld: lani robotMu: see robotWorld: pills -277239 .
156 }
157
158 robotWorld: lani_see_table -208510 {
159   robotWorld: lani robotMu: see robotWorld: table -208510 .
160 }

```