

# Translating Universal Scene Descriptions into Knowledge Graphs for Robotic Environment

Giang Hoang Nguyen<sup>1</sup>, Daniel Beßler<sup>1</sup>, Simon Stelter<sup>1</sup>, Mihai Pomarlan<sup>2</sup> and Michael Beetz<sup>1</sup>

**Abstract**—Robots performing human-scale manipulation tasks require an extensive amount of knowledge about their surroundings in order to perform their actions competently and human-like. In this work, we investigate the use of virtual reality technology as an implementation for robot environment modeling, and present a technique for translating scene graphs into knowledge bases. To this end, we take advantage of the Universal Scene Description (USD) format which is an emerging standard for the authoring, visualization and simulation of complex environments. We investigate the conversion of USD-based environment models into Knowledge Graph (KG) representations that facilitate semantic querying and integration with additional knowledge sources. The contributions of the paper are validated through an application scenario in the service robotics domain.

## I. INTRODUCTION

The vision of a world with accurate digital twins for every environment holds promise for advancing the field of service robotics. With access to detailed environment models, robots would be able to interact with their surroundings in a more sophisticated manner. Advanced game engines with physics simulation capabilities are a significant step towards achieving the vision of highly detailed environment models. These models use scene graph data structures and computational techniques for real-time maintenance. This work aims to explore the assumption that environment models used in modern games and the associated infrastructure is appropriate for robot environment modeling.

Scene graphs enable the representation of the compositional structure of environments, where each component and its configurations are modeled as sub-graphs within the scene graph. This allows selective consideration of only relevant parts of the graph based on the agent’s viewpoint. A robot can start with a coarse object model, refine it with additional knowledge, and insert it at different positions in the graph as the object moves. Nonetheless, the use of detailed environment models presents challenges in data management and must be addressed to fully utilize them.

The *Universal Scene Description* (USD) format is an attempt to address these challenges [5]. USD is an open-source, industry-standard format for describing and exchanging 3D scenes and assets. It enables artists and developers to work

The research reported in this paper has been (partially) supported by the European Union’s Horizon Europe Framework Programme, as part of euROBIN and FAME, and the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 Project-ID 329551904 “EASE - Everyday Activity Science and Engineering”, University of Bremen (<http://www.ease-crc.org/>)

<sup>1</sup>Institute for Artificial Intelligence, University of Bremen, Germany

<sup>2</sup>Applicable Linguistics, University of Bremen, Germany

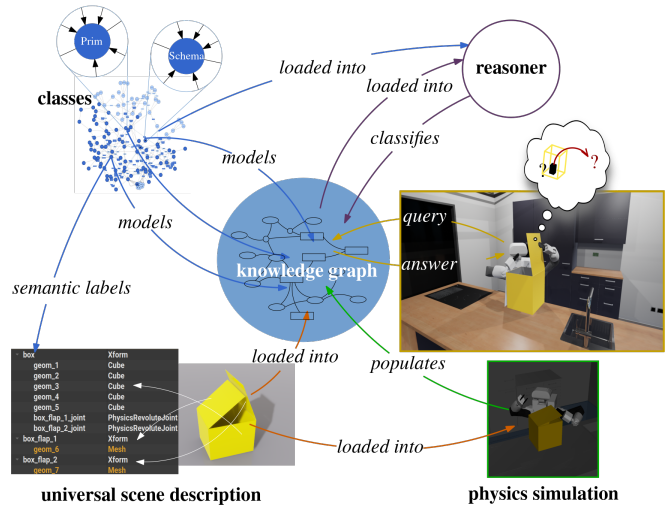


Fig. 1. Interactions of a Knowledge Graph with different components: Classes, Scene Description, Physics Simulation and Reasoner

with complex scenes and assets across multiple applications and workflows. It further supports the simulation of scenes, allowing artists and developers to test and refine their virtual environments by adjusting elements such as lighting and physics without the need to export or import data between different software applications. This streamlined workflow saves time and improves the quality of the final product.

To be practically useful for a robot, environment models need to be connected to the task of the robot and its body. Our approach is to represent such information as a linked KG, which is composed of knowledge covering robot, environment and task. Here, we focus on extracting environment knowledge from USD files to provide additional knowledge required for task execution.

Accordingly, the main research question motivating this work is whether standardized scene graphs represented in the USD format can be transformed into a representation that enables robots to model their environment while also rendering the model interpretable through semantic queries. Thus, a mapping from the USD format to a KG representation is proposed, which is fixed by an ontology that also supports the translation process via automated reasoning. Definitions in the ontology are further aligned with a common ontology, enabling linking to other knowledge sources. The resulting KG incorporates data mapped from the USD file, inferences made by the reasoner, and dynamical data computed by a simulator in order to build a more comprehensive relational

model of the environment. This process is depicted in Figure 1. The resulting infrastructure can be easily linked to other information sources, and provides a lowcost way of representing, manipulating and testing environments. In summary, this work presents the following contributions:

- a formalization of the USD format with extensions for semantic tagging of its entities,
- an ontological model of USD that fixes a KG and supports the translation process via inference, and
- the design and implementation of a method for automated KG construction from USD files.

The contributions will be validated experimentally by constructing a KG from a USD file, using a reasoner to finalize the KG and executing a robotics task in a physics simulation to update the KG.

## II. RELATED WORK

Developing robotic agents that can perform various indoor tasks is a complex and challenging problem that requires a comprehensive understanding of the environment. In this section, we review the related work that emphasizes the importance of using the USD in generating 3D objects and semantic models for robotics reasoning.

Mapping 3D objects and scene graphs using sensor data is an approach to understanding the physical properties of the environment. Rusu et al. [16] and Antoni et al. [15] presented systems for acquiring accurate 3D object maps using depth cameras and inertial measurement units. Beetz et al. [3] focused on mapping objects in retail environments. However, these models often suffer from limited transferability, making it challenging to reuse them in different contexts. Moreover, the models' representation in XML-based formats or point-clouds can be difficult to import into different simulation tools and prone to errors, as noted by Ivanou et al. [9] and Tola et al. [20]. The USD addresses these challenges by describing a scene graph with capability to handle large-scale [10] and complex datasets [21].

Creating accurate 3D object maps is insufficient for intelligent robotic agents. They also need to reason about potential uses of objects in their environment. Beßler et al. [4] proposed a conceptualization of affordances based on physical properties of objects and their potential uses. Similarly, Crespo et al. [6] presented a semantic relational model for robots to manage and query their environment. Semantic mapping is essential for creating a representation of an environment that includes both geometric and semantic properties. USD's ability to represent complex scenes, objects, and their relationships can support semantic mapping and thus improve the reasoning capabilities of robotic agents.

Recent research has introduced the application of semantic models and scene graphs in robotics reasoning ([22], [12], [1]). To enable comprehensive robotics reasoning with semantic models, robust knowledge processing is essential. In this regard, Haidu et al. [8] introduced KnowRobSIM, an extension of the KnowRob architecture ([2,18]) that encompasses subsymbolic world modeling, action simulation based on a physics engine, and scene rendering. However,

this approach relies solely on scene description from Unreal Engine and its physics. Thus, the inclusion of semantic information into the 3D data represented in USD offers a promising opportunity to advance robotics reasoning in a more comprehensive manner.

## III. THE UNIVERSAL SCENE DESCRIPTION FORMAT

The USD file format stores 3D scenes as hierarchical scene graphs (or *stages*). In USD, data is arranged hierarchically into namespaces of *prims* (primitives). Each prim can hold child prims, as well as *attributes* and *relationships*, referred to as *properties*. Attributes have typed values that can change over time, while relationships are pointers to other objects in the hierarchy. All of these elements, including prims and their contents, are stored in a *layer*, which is a scene description container for USD. This hierarchical structure allows for easy organization of the scene, and enables transformations and properties to be inherited from parent prims to their children. Listing 1 shows an example of how a cardboard box with two flaps can be represented using a hierarchy of prims in USD.

---

```
def Xform "world" () {
  def Xform "box" (
    prepend apiSchemas = ["PhysicsMassAPI", ...]) {
    matrix4d xformOp:transform = ( ... )
    token[] xformOpOrder = ["xformOp:transform"]
    float physics:mass = 2.79
    def Cube "geom_1" ( ... ) { ... }
    def PhysicsRevoluteJoint "box_flap_1_joint" {
      rel physics:body0 = </world/box>
      rel physics:body1 = </world/box_flap_1>
    }
    def PhysicsRevoluteJoint "box_flap_2_joint" { ... }
  }
  def Xform "box_flap_1" ( ... ) { ... }
  def Xform "box_flap_2" ( ... ) { ... }
}
```

---

Listing 1. Example of a box with two flaps in USD format

In the remainder of this section, we will provide a more detailed definition of the characteristics of USD that are relevant for the translation method.

### A. Composed prims

We describe a scene graph as a set of prims which we denote as  $\mathcal{P}$ . Each prim  $p \in \mathcal{P}$  is associated with a set of child prims. In Listing 1, child prims are used to represent different parts of a box, including the main part named *box*, the geometry that form the box *geom\_1* and the flaps named *box\_flap\_1* and *box\_flap\_2*. Each prim further has a specifier, a typed schema, an associated path, a set of API schemas and a set of properties. Accordingly, we say that each prim  $p$  has the form  $(h, t, n, P, M, E)$  where  $h \in \mathcal{S}$  is the specifier,  $t \in \mathcal{M}_T$  is the type schema,  $n \in \mathcal{N}_P$  is the path,  $P \subset \mathcal{P}$  are the children,  $M \subset \mathcal{M}_A$  are the (applied) API schemas, and  $E \subset \mathcal{E}$  are the (inherited) properties of the prim. Note that only a subset of properties from built-in schemas of the USD language is considered. Schemas used to describe shaders and rendering techniques are ignored. Here we focus on components of USD that are used in our example domain. Table I shows a selection of considered properties.

TABLE I  
A SELECTION OF PROPERTIES CONSIDERED IN THIS WORK.

	Schema	Property	Data type
Typed schema	Xform	<i>xformOpOrder</i>	token[]
		<i>xformOp:transform</i>	matrix4d
	Gprim	<i>primvars:displayColor</i>	color3f[]
		<i>primvars:displayOpacity</i>	float[]
	Cube	<i>size</i>	double
	PhysicsJoint	<i>physics:body0</i>	rel
		<i>physics:body1</i>	rel
API schema	PhysicsMassAPI	<i>physics:mass</i>	float
		<i>physics:centerOfMass</i>	point3f
	RdfAPI	<i>rdf:namespace</i>	string
		<i>rdf:conceptName</i>	string
	SemanticTagAPI	<i>semanticTag:semanticLabel</i>	rel

a) *Specifier*: Each prim is associated to a so called *specifier* that is used to distinguish concrete and abstract prim declarations. We denote the set of prim specifiers as  $\mathcal{S} = \{def, over, cls\}$ . These specifiers are used for prims that describe concrete objects in the scene, to overwrite certain properties, and to define parent classes of prims respectively. In the listed example, only the *def* specifier is used as each prim corresponds to a concrete object.

b) *Paths*: USD uses paths to locate prims within a stage, starting with the root path `"/`. Each prim has a name string, and child prims are located using paths prefixed by their parent's path. In the following, we denote the set of all possible prim names as  $\mathcal{N}_N$ , and assume that each prim name is prefixed by a slash. Hence, we can define the set of paths that can be constructed over these names as  $\mathcal{N}_P = \mathcal{N}_N^*$ .

c) *Schemas*: In USD, *schemas* set out conventions and rules for structuring, organizing, and accessing data. Two types of schemas are distinguished: *typed schemas* and *API schemas*. Typed schemas categorize the general types of prims, whereas API schemas specify certain attributes and relationships to describe object attributes. A prim can have only one typed schema but multiple API schemas. We denote the set of typed schemas as  $\mathcal{M}_T$ . The set contains elements for all built-in and user-defined typed schemas and an additional constant  $\lambda \in \mathcal{M}_T$  indicating that a prim has no typed schema. Similarly, we denote the set of API schemas as  $\mathcal{M}_A$ . The form of schemas is only relevant for resolving property values based on inheritance relations.

d) *Properties*: Properties of prims define specific aspects of the associated object, such as its transformation or material. USD distinguishes between two types of properties: attributes and relationships. Accordingly, we see the set of properties  $\mathcal{E}$  as the union of attributes and relations, i.e.,  $\mathcal{E} = \mathcal{E}_A \cup \mathcal{E}_R$  where  $\mathcal{E}_A$  is the set of attributes and  $\mathcal{E}_R$  the set of relations. Attributes are seen as tuples of attribute name and value. Let  $\mathcal{N}_A$  and  $\mathcal{V}_A$  be the set of allowed attribute names and values respectively. Then each attribute  $a \in \mathcal{E}_A$  can be written as a tuple  $(n, v)$  where  $n \in \mathcal{N}_A$  and  $v \in \mathcal{V}_A$ . In the example displayed in Listing 1, the prim named *box* specifies four attributes. Their names often indicate in which schema they are defined, e.g. the *xformOp* prefix indicates that attributes are derived from the *Xform* schema. Prims can

also have relationships with other prims. These are indicated by the keyword `"rel"` in the USD format. The value of a relation is given by a set of USD paths, i.e. one relation can link to many paths. Hence, we can define the set of relations as  $\mathcal{E}_R = \mathcal{N}_R \times \mathbb{P}(\mathcal{N}_P)$  where  $\mathcal{N}_R$  is the set of allowed relation names, and  $\mathbb{P}(\mathcal{N}_P)$  the powerset of paths.

### B. Prim specifications

Prim specifications can be seen as prims in their uncomposed form, and as part of a layer without the context in which the layer is embedded. Composed prims are created through a composition method where many prim specifications may contribute to the composed description. Such a composed prim includes properties inherited from prim specifications and the schemas they refer to. Prim specifications, on the other hand, only include properties and schemas that are referred to in the scope of the specification itself, but they additionally specify *composition arcs* which are used to compute the inheritance relation when prim compositions are created. Of the six composition arc types in USD, this paper focuses on the *sublayer* and *inherits* compositions. The sublayer composition integrates entire USD files into a main USD file, allowing the handling of complex scenes. Meanwhile, the inherits composition facilitates property inheritance among prims, aiding in creating variations and similar resources.

The representation of prims and their specifications is very similar. For our purpose here, the notion of prims only needs to be extended to cover the inheritance relation. Accordingly, we denote the set of prim specifications in a layer as  $\mathcal{P}_S$ , and say that each prim specification  $s \in \mathcal{P}_S$  has the form  $(h, t, n, P, I, E)$  where  $h \in \mathcal{S}$  is the specifier,  $t \in \mathcal{M}_T$  is the type schema,  $n \in \mathcal{N}_P$  is the path,  $P \subset \mathcal{P}_S$  are the children,  $I \in \mathbb{P}(\mathcal{M}_A) \times \mathbb{P}(\mathcal{N})$  is a tuple of API schemas and inheritance relations, and  $E \subset \mathcal{E}$  are the properties of  $s$ .

## IV. ONTOLOGICAL MODELING OF USD ENVIRONMENTS

The structure of KGs can be fixed through an ontological characterization of the vocabulary which is used to label nodes and edges. Here, we explore the use of an ontology language in the *Description Logic* fragment of the *Web Ontology Language* to support the translation process from USD format to a KG representation via logical axioms. The highest level of the ontology hierarchy we employ are externally defined ontologies that are widely applicable. The lower levels are aligned to the upper level, and thus must commit to the design decisions established at the upper level.

*Upper-Level Ontology*: For the most general definitions, we import the *DOLCE+DnS Ultra-lite* (DUL) foundational framework [11]. It defines concepts such as `PhysicalObject` which is the class of all things that occupy space, and properties that can characterize objects further such as *hasComponent* which is used to represent component hierarchies. DUL further defines a general model for object qualities around the notion of `Quality` that allow to form statements about individual aspects of entities.

### A. The USD Ontology

The purpose of the USD ontology is to establish a graph model of the USD language, and to support translation from a USD to a KG scene representation. To this end, we define the USD vocabulary in terms of an ontology. In addition, a set of re-usable built-in schemas within the USD language is used to describe prims. In order to facilitate translation, the ontological model attempts to replicate the structure of USD as closely as possible while in addition entailing semantics of USD entities. In the following, some central axioms of the ontology will be highlighted. However, due to space restrictions only a few will be listed in this paper.

In its current version, the USD ontology consists of 35 concepts, 42 properties, and 213 logical axioms. The ontology characterizes the composed view on prims. The layered view is not considered as it is mostly relevant for the authoring of environment models. The ontology further covers a set of built-ins of the USD language. These are 3 API schemas, and 11 concepts representing different typed schemas covering the requirements of the usecase considered in this work. However, many other built-in schemas can be characterized analogously.

1) *Representation of prims*: The basic notion in the USD ontology is the `Prim` concept. It represents concretely defined prims (i.e., using the `def` keyword). Prims with other specifiers (i.e., `over` and `class`) that only contribute to the uncomposed view are not considered in this work. In terms of the toplevel ontology, we define that  $\text{Prim} \sqsubseteq \text{Object}$ , where the notion of `Object` also includes objects without a proper space region. Parthood of prims is represented using the transitive and reflexive *hasPart* relationship of the toplevel ontology.

Each prim has a set of data attributes and relationships. Relations trivially map to edges between prims in the KG. Data attributes are used in the USD language to quantify characteristics of prims. The attribute *xformOp:scale*, for example, is used to quantify how a prim is scaled along three dimensions. It thus contributes to the shape of the prim. Following the toplevel ontology, we represent such characteristics using the `Quality` concept, e.g.,  $\text{Shape} \sqsubseteq \text{Quality}$  represents the quality of having a shape. Each quality concept further defines a property used to connect a prim to its quality:

$$\text{hasShape} \sqsubseteq \text{dul:hasQuality} \quad (1)$$

$$\text{hasShape}(x, y) \rightarrow \text{PO}(x) \wedge \text{Shape}(y) \quad (2)$$

$$\text{Prim} \sqsubseteq (\leq 1)\text{dul:hasQuality.Shape} \quad (3)$$

Axiom (2) defines range and domain of the property *hasShape*. Namely it asserts that having a shape implies that a prim must be of type `dul:PhysicalObject` (short `PO`). Axiom (3) asserts that a prim may have at max one shape.

We further define a set of data types of the USD language in the ontology such that they can be used to represent typed literals in the KG. This includes data types such as *float3* and *float4* that represent numeric vectors. The datatypes and quality concepts are then used to define built-in attributes of

the USD language. For example, *xformOp:scale* is defined as an attribute of the `Shape` quality of a prim where the value of the attribute has the data type `float3`.

2) *Representation of schemas*: Most of the built-in schemas of the USD language are used to describe prims, but a few also exist that describe properties. We characterize schemas ontologically using the notion of `Description`. Each schema entails that a certain set of properties is used for describing some phenomena such as properties used to quantify the mass of an object. As such schemas describe objects where certain qualities of the object are quantified in a specific way. They may further imply type constraints for objects they describe.

The API schema and typed schema (also called `IsA` schema) in the USD language is represented via the concepts `APISchema` and `TypedSchema` respectively. Both follow the same pattern how they are characterized: they may restrict the type of prims described by them, and further assert a particular quantification for a quality of the prim. For example, we define the `Xformable` schema as follows:

$$\text{Xformable} \sqsubseteq \text{TypedSchema} \quad (4)$$

$$\text{Xformable} \sqsubseteq \text{VisSchemaOf.PhysicalObject} \quad (5)$$

$$\text{Xformable} \sqsubseteq \text{VisSchemaOf.WithXform} \quad (6)$$

$$\text{WithXform} \equiv \exists \text{hasShape} . (\exists \text{xformOp:scale} \sqcap \dots) \quad (7)$$

Hence, only prims with type `PhysicalObject` can have this schema (5), and it establishes the use of data properties for the quantification of the shape (6, 7). Further note that *isSchemaOf* is aligned to the toplevel via its parent property *describes*.

Qualities of prims can also be classified via schemas. For instance,  $\text{CubeShape} \sqsubseteq \text{Shape}$  represents the quality of having a cube-like shape. We define `CubeSchema` as a schema that must have a `CubeShape`. Hence, any shape of a `CubeSchema` can be classified as a `CubeShape` which enables the classification via reasoning.

3) *Connectedness of Prims*: Prims are thought to be directly connected if there exists a joint that connects them. Joints are represented as prims that are described by a `PhysicsJointSchema`. The schema establishes the use of the relations *physics:body0*  $\sqsubseteq$  *hasConnection* and *physics:body1*  $\sqsubseteq$  *hasConnection* to link joint with prims:

$$\text{hasConnection}(x, y) \rightarrow \text{hasConnection}(y, x) \quad (8)$$

$$\text{physics:body0} \sqsubseteq \text{hasConnection} \quad (9)$$

$$\text{physics:body1} \sqsubseteq \text{hasConnection} \quad (10)$$

Hence, the (direct) connection between a joint and the two prims it controls is bidirectional (8).

However, prims of a joint are also connected with each other via the joint. This indirect connectedness can be inferred through the introduction of  $HC^T$  which is a transitive variant of the property:

$$\text{hasConnection} \sqsubseteq HC^T \quad (11)$$

$$HC^T(x, y) \wedge HC^T(y, z) \rightarrow HC^T(x, z) \quad (12)$$

## B. Logical Theories for Robot Decision Making

Formal ontologies enable the use of logical reasoning as they define a theory around notions they consider. The USD ontology encodes such a theory which is designed for reasoning about the mapping between USD and KG representations. Additional theories can be linked to the USD ontology to support robot decision making. As a proof of concept, in the scope of this work, we define an object type `Box` with two flaps of type `Flap` can be classified as either *opened* or *closed* based on a fixed threshold of the joint value connecting the box with the flaps. Accordingly, we assert that  $\text{hasQuality}(\exists \text{hasJointValue}(\geq 0.1))$  is a sufficient condition for a joint to be classified as *opened*, and that a `Box` can be classified as *opened* whenever its joints are opened. Axioms for being closed are defined analogously.

## V. KNOWLEDGE GRAPH CONSTRUCTION

The first step in our proposed method for generating a KG representation of a USD scene is to establish a USD layer that includes a collection of class prims representing the TBox ontology. Another USD layer representing the scene graph imports the TBox layer, and uses a custom API for tagging prims with ontological concepts. The semantically tagged prims are then translated into the KG.

### A. USD sublayer for semantic tagging

In order to link prims into the KG, we introduce the *SemanticTagAPI* schema. It views semantic tags as relationships in the USD framework, and uses the *semanticLabel* property to link prims to class prims corresponding to ontological concepts. The class prims utilize the *RdfAPI* schema to represent identifier of entities in the KG using the properties *rdf:namespace* and *rdf:conceptName*.

The USD layer for semantic tagging can then easily be generated given an ontology. A formalization of this procedure is omitted here due to space restrictions. In principle, it converts a set of concept names into a USD layer consisting prim specifications. For each concept in the ontology, a class prim is created according to the class hierarchy in the ontology, the concept name, and its associated namespace. The main USD layer that requires semantic tagging may then make use of class prims from the generated sublayer to classify prims.

### B. USD to KG translation

This subsection details the transformation of a USD layer into a KG. Algorithm 1 formalizes this process for a set of prims  $\mathcal{P}$ . To this end, an ABox ontology  $\mathcal{A}$  is generated that contains facts representing the prims in  $\mathcal{P}$ . For simplicity, we assume that  $\mathcal{P}$  also includes class prims (not only defined prims). The ABox is initially empty, and successively populated during an iteration of prims in  $\mathcal{P}$ . In the following, several assertions about each prim in  $\mathcal{P}$  are added to the ABox  $\mathcal{A}$ . These are statements involving the prim identifier  $i$  as a subject which identifies the prim at path  $n$  in the KG. The path  $n$  is mapped via the function *iri* to an identifier of the prim in the KG. The resulting ABox  $\mathcal{A}$  contains named individuals representing prims in  $\mathcal{P}$  and their qualities.

## Algorithm 1: USD to KG translation

---

**Input** : A set of prims  $\mathcal{P}$  forming a USD stage.  
**Output**: An ABox that describes the prims in  $\mathcal{P}$ .

```

1  $\mathcal{A} \leftarrow \emptyset$  // an empty ABox
2 for  $(\text{def}, t, n, P, M, E) \in \mathcal{P} \wedge i = \text{iri}(n)$  do
3   if  $t \neq \lambda \wedge c = \text{iri}_S(t)$  then // assert  $t$  as type of  $i$ 
4      $\mathcal{A} \leftarrow \mathcal{A} \cup \{(\exists \text{hasTypedSchema}.c)(i)\}$ 
5   endif
6   for  $(\text{def}, \rightarrow, n_c, \rightarrow, \rightarrow, -) \in P$  do // assert parts of  $i$ 
7      $\mathcal{A} \leftarrow \mathcal{A} \cup \{\text{hasPart}(i, j)\}$  where  $j = \text{iri}(n_c)$ 
8   endfor
9   for  $m \in M$  do // assert API schema  $m$  of  $i$ 
10     $\mathcal{A} \leftarrow \mathcal{A} \cup \{(\exists \text{hasAPI}.c)(i)\}$  where  $c = \text{iri}_S(m)$ 
11  endfor
12  for  $(n_E, v_E) \in E \wedge i_E = \text{iri}_P(n_E)$  do
13    if  $n_E = \text{semanticTag}:\text{semanticLabel}$  then
14      // assert semantic tags as types of  $i$ 
15      for  $n_r \in v_E \wedge (\text{cls}, \rightarrow, n_r, \rightarrow, \rightarrow, F) \in \mathcal{P}$  do
16         $i_t \leftarrow b \oplus g$ , where
17         $(\text{rdf:ns}, b) \in F, (\text{rdf:name}, g) \in F$ 
18         $\mathcal{A} \leftarrow \mathcal{A} \cup \{i_t(i)\}$ 
19      endfor
20    else if  $n_E \in \mathcal{N}_R$  then // assert a relation of  $i$ 
21      for  $n_r \in v_E$  do // to a prim at path  $n_r$ 
22         $\mathcal{A} \leftarrow \mathcal{A} \cup \{i_E(i, j)\}$  where  $j = \text{iri}(n_r)$ 
23      endfor
24    else
25       $\mathcal{A} \leftarrow \mathcal{A} \cup \{\text{hasQuality}(i, i_Q), i_E(i_Q, v_E)\}$ 
26      where  $i_Q = \text{iri}(\text{Quality})$ 
27    endif
28  endfor
29 return  $\mathcal{A}$ 

```

---

## VI. EVALUATION

This section demonstrates how USD supports robotics agents to perform human-scale tasks. The evaluation utilizes the robots simulation from MuJoCo [19], which is controlled by Giskard [17], a constraint-based motion planning using ROS communication and visualized by Unreal Engine. It also highlights USD's role in parsing scene descriptions between MuJoCo, ROS, Unreal Engine, and translating into the KG.

### A. Universal Scene Description Parser

Utilizing USD for scene description parsing offers consistency and minimal data loss. Analogous to the KG translation process, any new attributes specified in the scene description format, which USD doesn't originally recognize, can be incorporated into USD by introducing a new schema API. In this work, we only process commonly shared data, including transformations, geometry, and inertial properties of bodies and their associated joints. In essence, given a scene description in USD or any other format, this parser optimizes and translates it into USD, streamlining its conversion into KG or another format, as depicted in Figure 2.

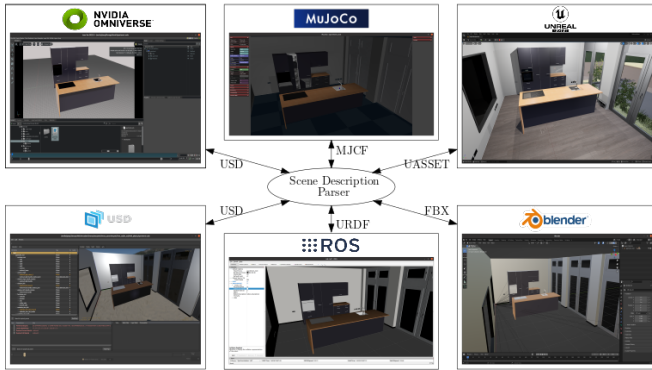


Fig. 2. Usage of USD as the translation medium for scene descriptions across various formats

### B. Giskard

Giskard [17] is a constraint-based motion planning and control framework. In contrast to many alternatives, it models not only the robot, but also the world kinematics[14]. This simplifies the expression of environment manipulation motions, by talking in terms of how the environment should change. Additionally, the state of Giskard’s world model can be synchronized by reading the joint state from the simulation environment or estimating it in reality.

### C. Task execution with USD

Robot’s comprehensive awareness of its environment is important for its performance in human-scale tasks. For tasks like tidying objects, robots might rely on online information or user cues, like “milk should be stored in the fridge”. Using USD and the built KG, robots execute tasks with a well-formed belief state, increasing their adaptability in human-centered tasks. In the context of our experiment, the robot is directed to open a box, identify its contents, and neatly place them in a proper location. The task execution pipeline is depicted in Figure 3.




Fig. 3. Task execution pipeline from knowledge to action

During the initialization phase, the scene graph of the apartment, consisting of 2247 prims (784 of type *Xform*, 658 of type *Gprim*, 45 of type *Joint*), is translated into the KG as an OWL file with 1487 prims, excluding 760 prims of type *Material*, *Shader*, etc. Subsequently, the KG with 6827 nodes and 24116 edges was subjected to a reasoning process to evaluate its consistency. As tasks are executed, the simulation updates the state to the KG, which is accessed by Konclude [7] for queries. Using Konclude’s *realization* operation, one

can determine an individual’s class, helping label items in a scene — for instance, determining if a flap is open or if an item is perishable. Another query identifies suitable storage locations for items, based on logic like “every spoon can be stored in any drawer”. While OWL-DL and Konclude have limitations in handling this, we use an ontology of item uses [13]. This combines OWL-DL’s object taxonomy with rules about how object classes can interact for tasks.

Giskard controls the robot using simulation data and constraints. Task decisions are derived from Konclude’s reasoning, which informs Giskard about the box’s status and object storage location. As the result, the robot successfully executes the task utilizing the scene graph-grounded KG. Figure 4 illustrates the task execution during each phase when Konclude updates Giskard’s belief state.

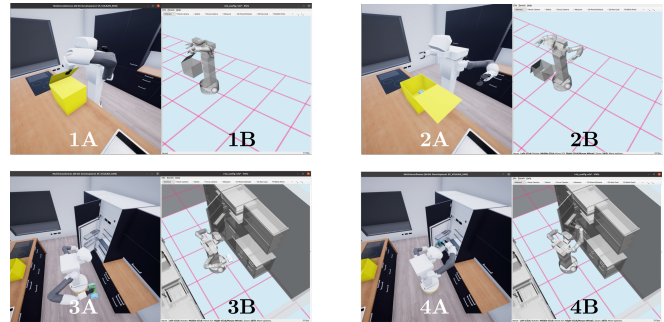


Fig. 4. Box unpacking scenario (Left: Unreal Engine, Right: RViz)

On the left, the visualization in Unreal Engine is depicted, whereas the right side represents the simulated robot’s belief state of Giskard in RViz. In the first phase, the box’s state is queried, disregarding the object inside it and the apartment. In the second phase, after the object *milk* is identified, the object’s storage container *fridge* is determined. In phases three and four, the robot picks the *milk*, opens the *fridge*, places it in the *tray*, and then closes the *fridge*.

## VII. CONCLUSION

This research presents a method for translating USD scene graphs into KGs, addressing three crucial aspects of the process. Firstly, the formalization and ontological modeling of the USD scene graph enables its representation in the KG and allows for the inference of additional knowledge. Secondly, a sublayer USD is generated, which represents the TBox ontology and enables semantic tagging using custom API schemas. The main USD layer is then linked to the knowledge base using these API schemas, and covers all individuals relevant for reasoning. Finally, the USD layer is converted into a KG, augmented with dynamical data from a physics simulator. By linking the USD into the KG, background knowledge for certain robotic decision-making tasks, as demonstrated in a box unpacking scenario. The resulting graph’s consistency was evaluated through reasoning, and its knowledge representation was verified through several relevant queries, thereby demonstrating the effectiveness of this approach.

## REFERENCES

- [1] Christopher Agia, Krishna Murthy Jatavallabhula, Mohamed Khodeir, Ondrej Miksik, Vibhav Vineet, Mustafa Mukadam, Liam Paull, and Florian Shkurti, *Taskography: Evaluating robot task planning over large 3d scene graphs*, Conference on Robot Learning, PMLR, 2022, pp. 46–58.
- [2] Michael Beetz, Daniel Beßler, Andrei Haidu, Mihai Pomarlan, Asil Kaan Bozcuoğlu, and Georg Bartels, *Know rob 2.0—a 2nd generation knowledge processing framework for cognition-enabled robotic agents*, 2018 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2018, pp. 512–519.
- [3] Michael Beetz, Simon Stelter, Daniel Beßler, Kaviya Dhanabalachandran, Michael Neumann, Patrick Mania, and Andrei Haidu, *Robots collecting data: Modelling stores*, pp. 41–64, Springer International Publishing, Cham, 2022.
- [4] Daniel Beßler, Robert Porzel, Pomarlan Mihai, Michael Beetz, Rainer Malaka, and John Bateman, *A formal model of affordances for flexible robotic task execution*, Proc. of the 24th European Conference on Artificial Intelligence (ECAI), 2020.
- [5] Alan Blevins and Mike Murray, *Zero to usd in 80 days: transitioning feature production to universal scene description at dreamworks*, ACM SIGGRAPH 2018 Talks, 2018, pp. 1–2.
- [6] Jonathan Crespo, Ramón Barber, and OM Mozos, *Relational model for robotic semantic navigation in indoor environments*, Journal of Intelligent & Robotic Systems **86** (2017), 617–639.
- [7] derivo GmbH, *Konclude reasoner*.
- [8] Andrei Haidu, Daniel Beßler, Asil Kaan Bozcuoğlu, and Michael Beetz, *Knowrobsim—game engine-enabled knowledge processing towards cognition-enabled robot control*, 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2018, pp. 4491–4498.
- [9] Mikhail Ivanou, Stanislav Mikhel, and Sergei Savin, *Robot description formats and approaches*, 2021 International Conference “Nonlinearity, Information and Robotics” (NIR), IEEE, 2021, pp. 1–5.
- [10] Krishna Murthy Jatavallabhula, Edward Smith, Jean-Francois Lafleche, Clement Fuji Tsang, Artem Rozantsev, Wenzheng Chen, Tommy Xiang, Rev Lebareddian, and Sanja Fidler, *Kaolin: A pytorch library for accelerating 3d deep learning research*, arXiv preprint arXiv:1911.05063 (2019).
- [11] Claudio Masolo, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari, *WonderWeb deliverable D18 ontology library (final)*, Tech. report, IST Project 2001-33052 WonderWeb: Ontology Infrastructure for the Semantic Web, 2003.
- [12] Runqing Miao, Qingxuan Jia, and Fuchun Sun, *Long-term robot manipulation task planning with scene graph and semantic knowledge*, Robotic Intelligence and Automation (2023).
- [13] Mihai Pomarlan and Robert Porzel, *Narrative objects*, Workshop on semantic techniques for narrative-based understanding at IJCAI-ECAI, 2022.
- [14] Adrian Röfer, Georg Bartels, Wolfram Burgard, Abhinav Valada, and Michael Beetz, *Kineverse: A symbolic articulation model framework for model-agnostic mobile manipulation*, IEEE Robotics and Automation Letters **7** (2022), no. 2, 3372–3379.
- [15] Antoni Rosinol, Arjun Gupta, Marcus Abate, Jingnan Shi, and Luca Carlone, *3d dynamic scene graphs: Actionable spatial perception with places, objects, and humans*, arXiv preprint arXiv:2002.06289 (2020).
- [16] Radu Bogdan Rusu, Zoltan Csaba Marton, Nico Blodow, Andreas Holzbach, and Michael Beetz, *Model-based and learned semantic object labeling in 3d point cloud maps of kitchen environments*, 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, 2009, pp. 3601–3608.
- [17] Simon Stelter, Georg Bartels, and Michael Beetz, *An open-source motion planning framework for mobile manipulators using constraint-based task space control with linear mpc*, 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022, pp. 1671–1678.
- [18] Moritz Tenorth and Michael Beetz, *Knowrob—knowledge processing for autonomous personal robots*, 2009 IEEE/RSJ international conference on intelligent robots and systems, IEEE, 2009, pp. 4261–4266.
- [19] Emanuel Todorov, Tom Erez, and Yuval Tassa, *Mujoco: A physics engine for model-based control*, 2012 IEEE/RSJ international conference on intelligent robots and systems, IEEE, 2012, pp. 5026–5033.
- [20] Daniella Tola and Peter Corke, *Understanding urdf: A survey based on user experience*, arXiv preprint arXiv:2302.13442 (2023).
- [21] Yizhou Zhao, Steven Gong, Xiaofeng Gao, Wensi Ai, and Song-Chun Zhu, *Vrkitchen2. 0-indoorkit: A tutorial for augmented indoor scene building in omniverse*, arXiv preprint arXiv:2206.11887 (2022).
- [22] Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, and Yuke Zhu, *Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs*, 2021 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2021, pp. 6541–6548.