# **Declarative Dynamic Object Reclassification**

Riccardo Sieve 🖾 🗈 University of Oslo, Norway

Eduard Kamburjan 🖂 💿 IT University of Copenhagen, Denmark

### **Ferruccio Damiani** ⊠ <sup>(D)</sup> University of Turin, Italy

Einar Broch Johnsen 🖂 💿

University of Oslo, Norway

#### — Abstract

In object-oriented languages, dynamic object reclassification is a technique to change the class binding of an object at runtime. Current approaches express when and how to reclassify inside the program's business code, while maintaining internal consistency. These approaches are less suited for programs that need to be consistent with an external context, such as autonomous systems interacting with a knowledge base. This paper proposes declarative dynamic object reclassification, a novel technique that provides a separation of concerns between a program's business code and its adaptation logic for reclassification, expressed via a knowledge base. We present Featherweight Semantically Reflected Java, a minimal calculus for declarative dynamic object reclassification that enables the programmer to define consistency both internally (using a type system) and externally (using declarative classification queries). We use this calculus to study how internal and external consistency interact for declarative dynamic object reclassification. We further implement the technique by extending SMOL, a language for reflective programming via external knowledge bases.

**2012 ACM Subject Classification** Software and its engineering  $\rightarrow$  Abstraction, modeling and modularity; Software and its engineering  $\rightarrow$  Object oriented languages

Keywords and phrases Dynamic Object Reclassification, Dynamic Software Updates, Featherweight Java, Knowledge Bases, Semantic Reflection, Type Soundness

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2025.2

Supplementary Material Software (Artifact): https://doi.org/10.5281/zenodo.15275245

# 1 Introduction

In traditional object-oriented programming [54, 56], the behaviour of an object instantiated from a class is defined by the methods and properties of that class. However, in many realworld applications, both the object's state and the context in which it operates can change over time. Dynamic object reclassification [15, 16] and typestate-oriented programming [2, 20] are techniques within class-based object-oriented programming that enable *dynamic* class binding of objects at runtime, allowing objects to change their class while retaining their identity. These powerful programming abstractions are particularly useful when an object's context changes over time, and the object needs to adapt to the new context. When the class-binding of objects is dynamic, affected objects need to remain *consistent* with their expected capabilities. For example, consistency becomes critical for dynamic software updates in asynchronous, IoT-driven ecosystems where the local processing of class re-binding may get delayed [34].

Previous approaches to such programming abstractions (e.g., [2,8,13–16,20,26]) handle the requirements for the reclassification process internally: programs describe classes and reclassification within the same language. This results in (1) a notion of consistency that does not capture the program's context, (2) an operational and low-level view of the conditions for



<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:1–2:30

Leibniz International Proceedings in Informatics

reclassification, and (3) a lack of separation of concerns between describing object behaviour and object reclassification. In contrast, autonomous and self-adaptive systems in, e.g., robotics and digital twins, typically need to adapt to a context represented by external knowledge bases containing information about both the environment and the system itself [7,42].

This paper proposes a novel approach to dynamic object reclassification, which shifts the reclassification process from ensuring internal consistency, with an operational view of the adaptation logic, to ensuring external consistency, with a declarative view of the adaptation logic. The approach combines three main elements:

- **Knowledge Base:** The context of the program is modelled as an external knowledge base, i.e., a logical representation of facts. This knowledge base evolves independently of the program and can be queried for Boolean results (*"is a certain formula implied by the knowledge base?"*) or retrieval (*"which values satisfy a given formula?"*).
- **Semantic Reflection** expresses that a program can query its own state in context of the knowledge base; this is enabled by a process of lifting the runtime state of a program into the knowledge base [38].
- **Declarative Object Reclassification** is defined via reclassification queries to the knowledge base. The first query defines when an object is *consistent* with a particular class; i.e., it is a *membership* or *classification* query. The second query that defines how to instantiate an object's fields when reclassifying an object into a particular class; i.e., it is a state *retrieval* query. These queries make use of semantic reflection, as they uniformly access both the lifted program state and the context in the knowledge base.

In combination, restrictions on queries and on class inheritance ensure that reclassification preserves type safety. In short, the main contributions of this paper are as follows:

- We propose a declarative dynamic object reclassification technique that uses (1) reflection of the program state into a knowledge base, and (2) the additional context of an external knowledge base, to provide an interface for reclassification queries that encapsulates the adaptation logic from the programmer.
- We formalise our technique in Featherweight Semantically Reflected Java (FSRJ), a novel minimal core calculus in the spirit of Featherweight Java [31], which supports semantic reflection into knowledge bases.
- We prove type soundness for FSRJ and give a precise characterisation of *program coherence*, describing the conditions under which reclassification queries ensure type safety.
- We provide a prototype implementation of declarative dynamic object reclassification as an extension of SMOL [38], a language for self-adaptive digital twins [40] based on Knowledge Graphs [29], including a static check for program coherence.
- We evaluate our prototype implementation by application to GreenhouseDT, an opensource digital twin of a mini-greenhouse [41].

The remainder of this paper is structured as follows: Section 2 illustrates declarative dynamic object reclassification with an example before Section 3 describes Featherweight Semantic Reflected Java, the minimal calculus for reflection into a knowledge base and subsequent object reclassification, as well as the type system. Section 4 describes the implementation in the SMOL language and Section 5 the evaluation on a self-adaptative digital twin. Finally, Section 6 discusses related work and Section 7 concludes with a discussion.

# 2 Overview

This section introduces the programming challenges that we address with declarative dynamic object reclassification, and motivates different technical aspects of our solution. To ease

© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:2–2:30

Leibniz International Proceedings in Informatics

```
Java
1 class Plant { int id; String species; }
2 abstract class Pump {
    int id; int gpioPin; Plant plant;
3
    void pump(){ ... }; /* uses gpioPin and waters the plant */
4
5}
6 class NormalPump extends Pump { ... /* methods */ }
7 class OverheatingPump extends Pump { int maximal; ... /* methods */ }
8 class Main() {
    Plant pl = new Plant(1, "Ocimum basilicum");
9
10
    Pump pu = new NormalPump(2, 7, pl);
    void loop() { while (true) { pu.pump(); System.wait(1); } }
11
    public static void main(String[] args) { new Main().loop(); }
12
13 }
```

**Figure 1** Greenhouse digital twin, without domain knowledge and dynamic object reclassification.

readability in the examples in this section and in the formalisation of Section 3, programs are given with a Java-like syntax and knowledge bases in first-order logic (FOL).

#### 2.1 Programming Challenges

**Motivating Scenario.** Let us consider as a running example a digital twin of a greenhouse which contains plants and water pumps. These are monitored through sensors measuring their physical *properties*: the height, oxygen level and soil moisture for each plant and the temperature for each pump. Each plant has an associated pump that waters the plant according to a watering profile. The concrete watering profile depends on the stage of the plant (e.g., seedling or mature plant) and the level of functionality of the pump. We first consider a simple configuration of the example in Sections 2 and 3, with one pump, two plants and a temperature monitor for the pump (see Figure 1); we later consider a larger configuration in Section 5, with three pumps, three plants and monitoring more properties.

In the digital twin, each plant is modelled by a Plant object and each pump by a Pump object. The Plant class contains the plant's identifier and scientific species name, so further information about the plant can be retrieved from a *domain knowledge base*. The Pump class contains the pump's identifier, a reference pl to the plant it is watering and a reference gpioPin to the output pin needed to operate the physical pump. The pump method implements the watering profile for the pump, which depends on the stage of the plant and whether the pump is fully operational or not: A pump is either operating normally (NormalPump) or overheated and in need of maintenance (OverheatingPump). To avoid overheating, an OverheatingPump should run at reduced power. Let us further assume that sensors (not specified here) update a synchronisation knowledge base containing the current state of the greenhouse. Thus, our program can retrieve information about the greenhouse from the knowledge base as needed.

The physical entities in the greenhouse evolve over time (e.g., a pump may threaten to overheat). The objects in the digital twin need to adapt their behaviour in accordance with the evolution of the physical entities they are twinning. To address this kind of problem, we propose a dynamic object reclassification technique that is *declarative*; i.e., the adaptation logic for reclassification is derived in the knowledge base. We now introduce a knowledge base for our example, and extend our program with constructs to interact with this knowledge base.



<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:3-2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Knowledge Bases and Consistency. A knowledge base is used to organise and make contextual information accessible via an API. To keep the presentation abstract, we here represent knowledge in terms of first-order logic formulas and structure the knowledge base for our digital twin as follows: (1) domain knowledge about plants and pumps, see Figure 2a, including the description of the physical pumps and plants in the greenhouse (the so-called domain objects); and (2) synchronisation knowledge about the current value of the properties of these domain objects, see Figure 2b. Observe how consistency here comes into play: a Pump object in the digital twin should be classified in the subclass of Pump implementing the watering policy that is appropriate for the current stage of the pump domain object it twins and controls.

We use patterns from ontologies and knowledge representation to model domain knowledge, represent domain objects and domain classes as constants (nullary predicates), and prefix all predicates concerning the domain by ctx\_. The *membership* of a domain object (e.g., ctx\_o) in a domain class (e.g., ctx\_C) is expressed by a binary predicate ctx\_in(ctx\_o, ctx\_C). the value x of a property  $\mathsf{ctx\_prop}$  for a domain object y by the predicate  $\mathsf{ctx\_prop}(y, x)$ . For instance, the temperature x of a pump y is expressed by the predicate  $\mathtt{ctx\_temp}(y, x)$ .

The formulas E1–E4 in Figure 2a express general knowledge about pumps; the remaining formulas specific knowledge about domain objects (here, one plant and one pump). Formula E1 states that a pump (ctx\_Pump) is operating normally (ctx\_NormalPump) if its temperature is  $\leq 50^{\circ}$ C, and **E2** that a pump is in danger of overheating (ctx\_OverheatingPump) if its temperature is  $> 50^{\circ}$ C. Formulas E3 and E4 express that the property ctx\_id uniquely identifies a domain object and that each pump can have only one temperature. The remaining formulas describe the plant and pump in our greenhouse, represented by constants ctx plant and ctx\_pump; ctx\_Plant is the domain class of all plants. The last formula states that the pump ctx\_pump has some temperature. The synchronisation knowledge base, in Figure 2b, contains, for each non-constant property ctx\_prop of the system, an axiom that states the current value of ctx\_prop. It contains a formula that states the actual temperature of ctx\_pump.

To express consistency, we need to relate the Pump program objects to the current state of the actual pumps in the greenhouse, as represented by the ctx\_Pump domain objects. To this aim, we represent the program code and its runtime configuration as knowledge bases. Figure 2c illustrates how the code in Figure 1 can be represented by a *code knowledge base*: predicate isCls(x) express whether x is a program program class; and subclass(x, y) express whether x is a subclass of y.

The representation of a program configuration in a knowledge base is known as *semantic* lifting [38], and uses a function defined in terms of the language to serialise runtime configurations into a *lifted-heap* knowledge base, where: isObj(x) express whether x is a program object (reference); instOf(x, y) whether x is an object instance of class y; in(x, y) whether x is an object instance of a subclass of y; and  $C_f(x, y)$  express whether x is an object in class a C and its field f has value y. Figure 2d shows the lifted runtime configuration of the program in Figure 1 after execution of the expression new Main() in Line 14 and field initialisation in Lines 11 and 12: here, the program object  $\iota_1$ , which is stored in pl, is an instance of program class Plant and  $\iota_2$  is an instance of program class NormalPump; and Pump\_plant  $(\iota_2, \iota_1)$  states that the field plant in  $\iota_2$  of program class Plant has value  $\iota_1$ .

Observe that the fact D1 ctx\_in(ctx\_pump, ctx\_OverheatingPump) can be derived in the knowledge base in Figure 2 (by instantiating the universal quantifier in the axioms E2and E3 with ctx\_plant), which shows that the lifted runtime configuration in Figure 2d (stating that  $\iota_2$  is an instance of NormalPump) is not consistent with the context. To enable declarative dynamic object reclassification, we need to connect program objects (e.g.,  $\iota_1$  and  $\iota_2$ ) and classes (e.g., Pump, NormalPump and OverheatingPump) to corresponding domain



licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:4-2:30 Leibniz International Proceedings in Informatics

- (E1)  $\forall x. \operatorname{ctx_in}(x, \operatorname{ctx_NormalPump}) \Leftrightarrow (\operatorname{ctx_in}(x, \operatorname{ctx_Pump}) \land \exists y. \operatorname{ctx_temp}(x, y) \land y \leq 50),$
- (E2)  $\forall x. \mathsf{ctx\_in}(x, \mathsf{ctx\_OverheatingPump}) \Leftrightarrow (\mathsf{ctx\_in}(x, \mathsf{ctx\_Pump}) \land \exists y. \mathsf{ctx\_temp}(x, y) \land y > 50),$

```
(E3) \forall x, y, z. (\mathtt{ctx\_id}(x, z) \land \mathtt{ctx\_id}(y, z)) \Rightarrow x \doteq y,
```

(E4)  $\forall x, y, z. (\mathsf{ctx\_temp}(x, y) \land \mathsf{ctx\_temp}(x, z)) \Rightarrow y \doteq z,$ 

- ctx\_in(ctx\_plant, ctx\_Plant), ctx\_in(ctx\_pump, ctx\_Pump),
- $ctx_id(ctx_plant, 1), ctx_id(ctx_pump, 2), \exists x.ctx_temp(ctx_pump, x)$
- (a) The *domain knowledge base* for the greenhouse. ctx\_temp(ctx\_pump, 52)
- (b) A synchronisation knowledge base describing the temperature of the pump. isCls(Plant), isCls(Pump), isCls(NormalPump), isCls(OverheatingPump), isCls(Main), subclass(NormalPump, Pump), subclass(OveheatingPump, Pump)
- (c) A snippet of the code knowledge base describing the code in Figure 1.  $isObj(\iota_1)$ ,  $isObj(\iota_2)$ ,  $isObj(\iota_3)$ ,  $instOf(\iota_1, Plant)$ ,  $instOf(\iota_2, NormalPump)$ , instOf( $\iota_3$ , Main), Plant\_id( $\iota_1$ , 1), Plant\_name( $\iota_1$ , "Ocimum basilicum"), Pump\_plant( $\iota_2, \iota_1$ ), Pump\_id( $\iota_2, 2$ ), Pump\_gpioPin( $\iota_2, 7$ )
- (d) The *lifted-heap knowledge base*, for the program's heap after the Main object has been initialised.

**Figure 2** Knowledge bases for the digital twin: domain, synchronisation, code and lifted heap.

objects (e.g., ctx\_plant and ctx\_pump) and classes (e.g., ctx\_Pump, ctx\_NormalPump and ctx\_OverheatingPump) in the knowledge base. In particular, we face the following challenges:

- **Ch1:** How can we relate program objects with the external knowledge base, and define external consistency, i.e., consistency between program and external knowledge base?
- **Ch2:** How can we program reactions to changes in the consistency of the relation between program and knowledge base?
- **Ch3**: How can we ensure that establishing external consistency does not break internal consistency, i.e., the typing of the program?

Knowledge Bases as Interoperability Layers. While our approach is general to dynamic reclassification, it is worth commenting on the role of the knowledge base, which is motivated by digital twins. In digital twins, the focus is not on highly complex low-level control, but on modularity, composition, and semantic interoperability of different components, where knowledge bases are widely used [42,60]. Often, all communication between the program and the physical component is handled via the knowledge base. The program is acting as a controller and a coordinating component that acts on data that has been already ingested into the knowledge base. This approach, which lends to blackboard architectures [25], also motivates our terminology of *external consistency*: The program is consistent with an external component (namely the physical twin), as described by the synchronization knowledge base. From a model-based perspective, this corresponds to intermodel consistency (see, e.g., Feichtinger et al. [19]), in contrast to internal, or intramodel, consistency, which corresponds to type soundness.

#### 2.2 **Declarative Dynamic Object Reclassification**

Based on the knowledge base and scenario above, we now illustrate the mechanisms in the programming language to connect the program with its context at runtime, shown in Figure 3. The formalisation and properties for these mechanisms will be explored in Section 3.



```
Extended Java code
1 class Plant { int id; String species; }
2 abstract class Pump classifies \lambda self. ctx_in(self, ctx_Pump) {
     int id; int gpioPin; Plant plant;
3
     void pump(){ ... }; /* uses gpioPin and waters the plant */
4
5}
6 class NormalPump extends Pump
     links \lambda self. ctx_in(self, ctx_Pump) \land \forall x. Pump_id(self, x) \Rightarrow ctx_id(self, x)
     classifies \lambda self. ctx_in(self, ctx_NormalPump) { ... /* methods */ }
8
9 class OverheatingPump extends Pump
     links \lambdaself. ctx_in(self, ctx_Pump) \wedge \forall x. Pump_id(self, x) \Rightarrow ctx_id(self, x)
10
     classifies \lambda self. ctx_in(self, ctx_OverheatingPump)
11
     retrieves \lambda self, maximal. \exists x. \ \texttt{ctx\_profile}(self, x) \land \texttt{ctx\_maximalPower}(x, maximal)
12
13
     { int maximal; ... /* methods */ }
14 class Main {
    Plant pl = new Plant(1, "Ocimum basilicum");
15
     Pump pu = new NormalPump(2, 7, pl);
16
     void loop() { while (true) { adapt(pu); pu.pump(); System.wait(1); } }
17
     public static void main(String[] args) { new Main().loop(); }
18
19 }
```

(a) The greenhouse digital twin programmatic part with links, classifies, retrieves and adapt.

 $\forall y. \text{ instOf}(y, \text{NormalPump}) \Rightarrow (\text{ctx}_in(y, \text{ctx}_Pump) \land \forall x. \text{Pump}_id(y, x) \Rightarrow \text{ctx}_id(y, x)), \\\forall y. \text{ instOf}(y, \text{OverheatingPump}) \Rightarrow (\text{ctx}_in(y, \text{ctx}_Pump) \land \forall x. \text{Pump}_id(y, x) \Rightarrow \text{ctx}_id(y, x)) \\(\textbf{b}) \text{ Linkage knowledge base. Tautologies from the default links clauses of Plant and Main are omitted.}$ 

**Figure 3** The greenhouse digital twin programmatic part and the linkage knowledge base.

**Linkage.** To address the first part of challenge **Ch1**, we propose to dynamically add knowledge about the dynamically created objects of the program's runtime state to the knowledge base. Technically, we add to each non-abstract class C a links clause with a singleparameter  $\lambda$ -predicate  $\lambda$ self. $\phi$  to express that  $\forall x.instOf(x, C) \Rightarrow \phi[self := x]$  is added to the knowledge base (the  $\lambda$ -predicate  $\lambda$  self.true here acts as default). At runtime, dynamically created objects are lifted into the lifted-heap knowledge base by adding a fact about their class membership, e.g.,  $instOf(\iota, C)$  for some object  $\iota$  (cf. Section 2.1). In Figure 3 (Section 2.1), the program class NormalPump declares its objects to be members of domain class ctx\_NormalPump and the domain id (ctx\_id) and program id (Pump\_id) to be the same. Similarly for OverheatingPump (Section 2.1). Linkage connects the objects of C to the context, using domain terminology (i.e., predicate symbols). Figure 3b shows the linkage knowledge base containing the added axioms. The first states that every program object that is an instance of NormalPump, is also a pump (i.e., a ctx\_Pump in the domain context), and that the ids are the same. The same states the same for OverheatingPump. This links the program object to the domain knowledge, as one can use *domain* axioms (such as E1 and E2) to reason about the program objects. This is done as follows. From the lifted heap in Figure 2d, we derive that  $\iota_2$  is member of ctx\_OverheatingPump as follows: By lifting, we have instOf( $\iota_2$ , NormalPump) and by linkage,  $\mathtt{ctx\_id}(\iota_2, 2)$ . Since  $\mathtt{ctx\_id}$  uniquely identifies an object (cf. E3),  $\mathtt{ctx\_pump} \doteq \iota_2$ . Applying this equality in the formula ctx\_in(ctx\_pump, ctx\_OverheatingPump) (D1, see above), we obtain  $\mathtt{ctx\_in}(\iota_2, \mathtt{ctx\_OverheatingPump})$ .

**Consistency Declaration.** To address the second part of challenge Ch1, we propose to characterise in the knowledge base when an object should belong to a particular class. To this aim, we let objects change between *reclassifiable* subtypes of an *adaptable* (abstract) class. Technically, we add a **classifies** clause with a single-parameter  $\lambda$ -predicate to these



© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:6–2:30 Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



**Figure 4** The structure of a class hierarchy for declarative dynamic object classification.

classes. If this  $\lambda$ -predicate holds when instantiated for a program object of the given program class, the program object is considered consistent with its context. In our example (see Figure 2),  $\iota_2$  is *not* consistent. The only classes with a **classifies** clause are

- 1. *adaptable* classes, i.e, abstract classes that extend classes with no classifies clause; and
- 2. *reclassifiable* classes, i.e., all the subclasses of the adaptable classes (introduced in point 1), that must be non-abstract and have no subclasses.

Remark that, as illustrated in Figure 4, each reclassifiable class must (1) directly inherit from an adaptable class; (2) not be inherited from; and (3) have only reclassifiable sibling classes.

**Reclassification Queries.** Reclassification is the process of changing the program class of an inconsistent program object to a program class that makes the program object consistent. To address challenge Ch2, we provide language support for programming such reclassification. The expression adapt(e) (Section 2.1) takes a program object that must be an instance of a reclassifiable class as argument, and checks whether this program object is consistent with the context; if this is not the case, it is reclassified. This is done by searching the knowledge base for a sibling to its current program class in the class hierarchy, that would make the object consistent. For the running example (see Figure 2), this would be the OverheatingPump class. Remark that the adapt expression is an explicit mechanism to ensure external consistency, but far from the only one. External consistency can also be established implicitly, either lazily, whenever an object is accessed, or periodically in the background, akin to a garbage collector.

To support state transfer for the reclassified object to the fields of its new program class, the programmer adds a **retrieves** clause to each reclassifiable class. This clause provides a  $\lambda$ -predicate with one parameter for the object in question and one for each field declared in the class. We can then infer values from the knowledge base that make the  $\lambda$ -predicate hold. Consider the OverheatingPump in Figure 3a: the reclassification query (Section 2.1) first retrieves the profile of the given pump, and then the maximal power it can use when in this stage.

**Ensuring Consistency and Type Soundness.** To address challenge Ch3, let us observe that it is crucial that the dynamic operations on the knowledge base maintain satisfiability of the knowledge base — otherwise, the program's interactions with the knowledge base would allow incorrect information to be inferred, potentially violating type soundness. Technically, we first ensure that reclassifiable classes are not used as types; e.g., a variable of type OperationalPump is not allowed, as it may cause an error when accessed after reclassification. Both this restriction and the structure illustrated in Figure 4 are enforced by the type system; they guarantee that no unsafe access is possible.

To ensure that the interaction between the program and the knowledge base does not introduce errors, we introduce a notion of *program coherence* roughly expressing that

1. every lifted object of a reclassifiable class can be classified as member of the adaptable superclass;



Prg ::=	$\mathcal{K} \ \overline{\mathtt{CD}} \ \mathtt{e}$	Program
$\mathcal{K} ::=$	$\{\overline{\phi}\}$	Knowledge base
CD ::=	class C [extends C] [Links] [Adapt] $\{\overline{FD} \ \overline{MD}\}$	Class
FD ::=	T f;	Field
T ::=	C   int	Type
MD ::=	MH {return e; }	Method
MH ::=	$T m(\overline{T} \overline{x})$	Method header
e ::=	$\texttt{x} \mid n \mid \texttt{e.f} \mid \texttt{e.m}(\overline{\texttt{e}}) \mid \texttt{new} \ \texttt{C}(\overline{\texttt{e}}) \mid \texttt{e.f} = \texttt{e} \mid \texttt{null} \mid \texttt{adapt}(\texttt{e})$	Expression
Links ::=	links $\lambda z.\phi$	Linkage
Adapt ::=	classifies $\lambda z.\phi~[{ t retrieves}~\lambda z\overline{z}.\phi]$ ;	Adaptation

**Figure 5** FSRJ syntax, where  $\phi$  ranges over first-order logic (FOL) formulas, C over class names, f over field names, T over types (i.e., class names and the basic type int), m over method names, x over local variables (i.e., formal parameters to methods) and z over term variables (so  $\lambda z.\phi$  is a unary  $\lambda$ -predicate and  $\lambda z\overline{z}.\phi$  a  $\lambda$ -predicate of arity  $\geq 1$ ).

- 2. every lifted object that is member of an adaptable class can also be classified as member of one of the reclassifiable subclasses; and
- 3. the retrieves  $\lambda$ -predicate associates a value of the correct type to each field.

We say that an object is *hot* to mean that a method invoked on the object has an activation record on the stack, and *cold* otherwise. In the next section we will see that program coherence ensures that reclassification of cold objects is always successful.

### **3** FSRJ: Featherweight Semantically Reflected Java

This section provides a formal account of programming with semantic reflection in terms of *Featherweight Semantically Reflected Java (FSRJ)*, a minimal core calculus for class-based object-oriented languages with semantic reflection. Our purpose with FSRJ is to formalise the interplay between program execution, semantic reflection and the querying of knowledge bases for dynamic object reclassification. FRSJ is a variant of Featherweight Java (FJ) [31].<sup>1</sup>

# 3.1 FSRJ Syntax

The syntax of FSRJ is given in Figure 5. We let  $\overline{\phantom{a}}$  represent (possibly empty) sequences and  $[\cdot]$  represent optional elements. Operations on sequences of pairs are abbreviated in a similar way; e.g.,  $\overline{C} \ \overline{f}$  is shorthand for  $C_1 \ f_1, \ldots, C_n \ f_n$ . The empty sequence is denoted by  $\bullet$  and the length of a sequence  $\overline{e}$  is denoted by  $\#(\overline{e})$ . The set of program variables includes the reserved name this, which cannot be used as a name of a field or a method's formal parameter.

In FSRJ, a program Prg consists of a knowledge base  $\mathcal{K}$ , a list of class definitions  $\overline{CD}$  and a main expression e. The knowledge base  $\mathcal{K}$ , called the *domain knowledge base* of Prg, is set of FOL formulas. A class definition class C [extends D] [Links] [Adapt] {FD  $\overline{MD}$ } consists of a class name C, an optional clause extends D to express inheritance (so D is the immediate superclass of C), an optional linkage clause Links = links  $\lambda z.\phi$ , an optional adaptation clause Adapt = classifies  $\lambda z.\phi$  [retrieves  $\lambda z\overline{z}.\psi$ ] (where, in turn, the retrieves part is optional),

<sup>&</sup>lt;sup>1</sup> FSRJ does not technically extend FJ because cast-expressions are excluded; these are orthogonal to semantic reflection and object reclassification. Further, FSRJ is imperative (it features a field assignment expression), has no implicit root class (Object), and classes in FSRJ need not have a superclass and have an implicit constructor which initialises all the fields.



<sup>@</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:8–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

a list of field definitions  $\overline{\mathtt{PD}}$  and a list of method definitions  $\overline{\mathtt{MD}}$ . The knowledge base, formalising the context in which the program executes, is detailed in Section 3.2. The linkage clause expresses that every instance o of **C** is logically linked by  $\lambda z.\phi$ ; i.e., given the address  $\iota$  of o, the formula  $\phi[z := \iota]$  holds in the context of the program. The adaptation clause expresses that if **C** has the immediate superclass **D**, then every instance o of an immediate subclass of **D** is reclassifiable to **C** by  $\lambda z.\phi$  and  $\lambda z \overline{z}.\psi$ ; i.e., given the address  $\iota$  of o, if the formula  $\phi[z := \iota]$  holds in the context of the program then there exist values  $\overline{v}$  for the fields of **C** such that the formula  $\psi[z\overline{z} := \iota\overline{v}]$  holds in the context of the program. Types **T** include names of declared classes and the basic type int. Field and method declarations are standard; the fields declared in a class are added to the ones declared in its superclasses, and are assumed to have distinct names (i.e., no field shadowing is allowed). Method names are assumed to be unique in the class; i.e., no method overloading is allowed. It is, however, possible to override methods from the superclass. All fields and methods are public, and each class has an implicit constructor that takes a parameter for each field (inherited or defined) in the class and initialises all the fields.

We assume that the elements of lists of named entities (i.e., class definitions, field definitions, method definitions, method formal parameter declarations) have different names. To simplify the presentation, we use a program Prg as a mapping from class names C to the corresponding class definitions CD; and use a class definition CD as a mapping that maps field and method names to field and method definitions, may map the keyword **extends** to the name of its superclass D, may map each of the keywords **links** and **classifies** to a unary  $\lambda$ -predicate of arity one and the keyword **retrieves** to a  $\lambda$ -predicate of arity n + 1, where n is the number of fields (defined or inherited) in the class. We assume that programs Prg satisfy the following *well-formedness conditions:* 

1. For a class name C appearing anywhere in Prg, we have  $C \in dom(Prg)$ .

2. There are no cycles in the transitive closure of the immediate extends relation.

The *subclassing* relation <: is the reflexive and transitive closure of the immediate **extends** relation. The *subtyping* relation  $\leq$  extends <: by adding reflexivity on basic types, i.e., int  $\leq$  int.

In the sequel, we let the metavariable a range over names and the metavariable AD over definitions. The lookup of the definition of a field or method a in a class C is denoted by aDef(C)(a). For each class in dom(Prg), the function aDef(C) is defined as follows:

 $\texttt{aDef(C)(a)} = \begin{cases} \texttt{Prg}(\texttt{C})(\texttt{a}) & \text{ if } \texttt{a} \in \texttt{dom}(\texttt{Prg}(\texttt{C})) \\ \texttt{aDef(D)(a)} & \text{ if } \texttt{a} \notin \texttt{dom}(\texttt{Prg}(\texttt{C})) \text{ and } \texttt{Prg}(\texttt{C})(\texttt{extends}) = \texttt{D} \end{cases}$ 

Given a field definition FD = T f and a method definition  $MD = T m (\overline{T} \overline{x}) \{\cdots\}$ , let signature (FD) denote the type T of field f and signature(MD) the type  $\overline{T} \to T$  of method m.

▶ **Example 1** (Greenhouse digital twin in FSRJ). Let us consider the encoding of the Greenhouse example (Figure 3a) in FSRJ. Sequential composition, which is not part of the syntax of FSRJ, can be encoded. For  $e_1; \dots; e_n$  of  $n \ge 1$  expressions, define an auxiliary class Encode with a method  $T_n \operatorname{sqT1_T2_...Tn} (T_1 x_1, ..., T_n x_n)$  { return  $e_n$ ;}, where  $T_i$  is the type of  $e_i$   $(1 \le i \le n)$ . Namely,  $e_1; \dots; e_n$  can be expressed as new Encode().sqT1...Tn $(e_1, ..., e_n)$ .

Then, if we ignore the keyword abstract, consider void as syntactic sugar for int, consider the string literal "Ocinum basilicum" as syntactic sugar for an integer literal, and consider System.wait(1) as syntactic sugar for new System().wait(1) where System is a class providing a method int wait(int x) {...}, we have that the Java-like code fragment in Figure 3a can be turned into an FSRJ code fragment by replacing the code of the class Main (in lines 16-21) by the following FSRJ code:



© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:9–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

```
FSRJ
```

# 3.2 A FOL Representation of Knowledge Bases

To avoid the technicalities of knowledge representation languages in FSRJ, the knowledge base  $\mathcal{K}$  of an FSRJ program is given as a set of closed FOL formulas over a given signature. Formally, signatures and formulas over signatures are defined as follows.

▶ **Definition 2** (Signatures). A signature is a set of predicate symbols  $\Sigma = \Sigma_0 \cup \Sigma_1 \cup \Sigma_2$ , where  $\Sigma_0$  is a set of 0-ary predicate symbols (constants),  $\Sigma_1$  a set of unary predicate symbols, and  $\Sigma_2$  a set of binary predicate symbols (relations). Let c range over  $\Sigma_0$ , P over  $\Sigma_1$  and R over  $\Sigma_2$ .

▶ **Definition 3** (FOL formulas and terms over a signature). Let term variables z range over constants in a signature  $\Sigma$ . The formulas  $\phi$  and terms t over  $\Sigma$  are defined as follows:

 $\phi ::= \neg \phi \mid \phi \lor \phi \mid \phi \land \phi \mid \phi \oplus \phi \mid \phi \Rightarrow \phi \mid \phi \Leftrightarrow \phi \mid \exists z. \ \phi \mid \forall z. \ \phi \mid P(t) \mid R(t,t) \mid t \doteq t$  $t ::= z \mid c$ 

Thus, in a  $\lambda$ -predicate  $\lambda z \overline{z}.\psi$ ,  $z\overline{z}$  is a (non-empty) list of term variables and  $\phi$  a FOL formula. We denote by  $\mathsf{FV}(t)$  the free variables in a term t and by  $\mathsf{FV}(\phi)$  the free variables in a formula  $\phi$ . The substitution  $\phi[z := t]$  replaces all free occurrences of the term variable z by the term t in the formula  $\phi$  (we omit its inductive definition here). Given a knowledge base  $\mathcal{K}$  and a formula  $\psi$ , we let  $\mathcal{K} \Rightarrow \psi$  abbreviate ( $\wedge_{\phi \in \mathcal{K}} \phi$ )  $\Rightarrow \psi$ .

Recall that the semantics of FOL with equality defines the values of terms to be individuals of a domain with an identity-relation, and the values of (well-formed) formulas to be truth values [4]. Thus, the semantics of a formula over a signature  $\Sigma$  is defined relative to an equality-model  $\mathcal{M} = \langle \mathcal{D}, \mathcal{I}, \approx \rangle$  and a variable assignment  $\beta$  that maps term variables to elements in the domain  $\mathcal{D}$ , such that: (1) the interpretation  $\mathcal{I} : \Sigma \to \mathcal{D}^n$  maps symbols in  $\Sigma_2$ to binary relations over  $\mathcal{D} \times \mathcal{D}$ , symbols in  $\Sigma_1$  to unary predicates over  $\mathcal{D}$ , and symbols in  $\Sigma_0$ to distinct elements in  $\mathcal{D}$ ; and (2) the interpretation  $\mathcal{I}(\doteq)$  is the equality-relation  $\approx$  on  $\mathcal{D}$ . In the following definition of satisfiability, we omit the operators  $\forall, \land, \oplus$  (which denotes exclusive or),  $\Rightarrow$  and  $\Leftrightarrow$ , since their semantics can be derived by the semantics of the other operators.

▶ **Definition 4** (Satisfiability). Let  $\Sigma$  be signature,  $\mathcal{M} = \langle \mathcal{D}, \mathcal{I}, \approx \rangle$  an equality-model such that  $\mathcal{I} : \Sigma \to \mathcal{D}$ , and  $\beta$  a variable assignment. The values of terms t with respect to  $\beta$  in  $\mathcal{M}$ , denoted  $\llbracket t \rrbracket_{\mathcal{M},\beta}$ , are defined by  $\llbracket z \rrbracket_{\mathcal{M},\beta} = \beta(z)$  and  $\llbracket c \rrbracket_{\mathcal{M},\beta} = \mathcal{I}(c)$ . The relation  $\mathcal{M}, \beta \models \phi$ , expressing that a formula  $\phi$  is satisfied by  $\mathcal{M}$  and  $\beta$ , is defined inductively as follows:

$$\begin{split} \mathcal{M}, \beta &\models \neg \phi & iff \ \mathcal{M}, \beta \not\models \phi \\ \mathcal{M}, \beta &\models \phi_1 \lor \phi_2 & iff \ \mathcal{M}, \beta \models \phi_1 \ or \ \mathcal{M}, \beta \models \phi_2 \\ \mathcal{M}, \beta &\models \exists x. \ \phi & iff \ \mathcal{M}, \beta [x \mapsto d] \models \phi \ for \ some \ d \in D \\ \mathcal{M}, \beta &\models P(t) & iff \ \llbracket t \rrbracket_{\mathcal{M}, \beta} \in \mathcal{M}(P) \\ \mathcal{M}, \beta &\models R(t_1, t_2) & iff \ (\llbracket t_1 \rrbracket_{\mathcal{M}, \beta}, \llbracket t_2 \rrbracket_{\mathcal{M}, \beta}) \in \mathcal{M}(R) \end{split}$$

A formula  $\phi$  over  $\Sigma$  is (1) satisfiable if  $\mathcal{M}, \beta \models \phi$  holds for some  $\mathcal{M}$  and  $\beta$ , and (2) valid, denoted  $\models \phi$ , when  $\phi$  is satisfiable for every  $\mathcal{M}$  and variable assignment.

A knowledge base  $\mathcal{K}$  is satisfiable/valid if  $\wedge_{\phi \in \mathcal{K}} \phi$  is satisfiable/valid.

Given a knowledge base  $\mathcal{K}$  and a formula  $\psi$ , we write  $\mathcal{K} \models \psi$  as shorthand for  $\models \mathcal{K} \Rightarrow \psi$ .

© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:10–2:30

licensed under Creative Commons License CC-BY 4.0 39th European Conference on Object-Oriented Programming (ECOOP 2025).

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

**Example 5** (Domain knowledge base for the Greenhouse digital twin in FSRJ). The domain knowledge base for the FSRJ Greenhouse digital twin is given in Figure 2a.

# 3.3 Program Typing

In this section we present the type system for FSRJ, detailing especially on the language aspects related to knowledge bases, semantic reflection and dynamic object reclassification.

The function type, that extracts the type from a field or method definition, is defined by: type(T f) = T and type(T m( $\overline{T} \ \overline{x}$ ){ $\cdots$ }) =  $\overline{T} \rightarrow T$ . The lookup of the type of a field or method a in the program is defined b:y aType(C)(a) = type(aDef(C)(a)). With respect to reclassification, we say that a class  $C \in dom(Prg)$  is:

- standard, written Std(C), if (1) links ∈ dom(Prg(C)); (2) classifies ∉ dom(Prg(C)); (3) retrieves ∉ dom(Prg(C)); and (4) if C has a superclass, then this superclass is standard (i.e., if Prg(C)(extends) = B for some B, then Std(B)).
- adaptable, written Adp(C), if (1) links ∉ dom(Prg(C)); (2) classifies ∈ dom(Prg(C)); (3) retrieves ∉ dom(Prg(C)); and (4) if C has a superclass, then this superclass is standard (i.e., if Prg(C)(extends) = B for some B, then Std(B)).
- reclassifiable, written Rcl(c), if (1) links ∈ dom(Prg(C)); (2) classifies ∈ dom(Prg(C));
   (3) retrieves ∈ dom(Prg(C)); and (4) C has a superclass which is adaptable (that is, Prg(C)(extends) = B and Adp(B), for some B).

By inspecting the program, it is possible to check that there are no cycles in the transitive closure of the **extends** relation; that, for each class C in dom(Prg), the names of the fields in C are distinct from the names of inherited fields; that each class is either standard, or adaptable, or reclassifiable; and that each adaptable class has at least a subclass (which is necessarily reclassifiable). We write **sane(Prg)** to express that the program **Prg** satisfies the well-formedness conditions in Section 3.1 and that the following sanity conditions:

- 1.  $C_1 <: C_2$  implies that, for all method names m, if  $aType(C_2)(m)$  is defined, then  $aType(C_1)(m) = aType(C_2)(m)$ .
- **2.**  $C_1 <: C_2$  and  $C_1 \neq C_2$  imply that, for all field names f, if  $f \in dom(Prg(C_1))$  then  $f \notin dom(Prg(C_2))$ .
- **3.** For every  $c \in dom(Prg)$ , exactly one of Std(c), Adp(c) or Rcl(c) holds.
- 4. For every  $C \in dom(Prg)$ , if Adp(C) then there exists  $D \in dom(Prg)$  s.t. Prg(D)(extends) = C.
- 5. Reclassifiable class names are not used as field types, as method parameters or return types.

The typing rules and field lookup rule for FSRJ are given in Figure 6. Rule T-knowledge enforces that the domain knowledge base consists of closed formulas. The typing of expressions is standard, except for rule T-new, which enforces new objects to be typed either by an adaptable or a standard class, and rule T-adapt, which enforces the **adapt**-expression and the argument of **adapt** to be typed by an adaptable class c. This allows the argument of **adapt** to be dynamically reclassified to any subclass of c.

The typing of links clauses ensures that the linkage formula only contains one free variable z, which ensures that the formula can be closed by instantiating z. Similarly, the typing of the adapts clause ensures that the free variables of the classifies- and retrieves-formulas are known. The typing rules for method and class definitions and for programs are standard recursively checking the different syntactic elements; the latter rule additionally checks that the program is sane. In the type system, the auxiliary function fields(C) retrieves all field definitions of a given class C in the program Prg, including the inherited ones. To type the null value, we use the special type  $\bot$ , which is not a class name, not part of the program and a subtype of any class name. In the sequel, we use the following convention:



© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:11–2:30

<sup>39</sup>th European Conference on Object-Oriented Programming (ECOOP 2025).

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



**Figure 6** FSRJ: typing rules and field lookup rule.

the metavariable T denotes either a non-reclassifiable class name or int; and
 the metavariable S denotes either a non-reclassifiable class name, int or ⊥.
 We say that a program Prg is well-typed to mean that ⊢ Prg OK holds.

▶ Lemma 6 (Syntactic guarantees for well-typed programs). Let  $\vdash Prg \ OK$ . Then sane(Prg) and if new C(...) appears in Prg then C is not adaptable.

# 3.4 Code, Linkage and Program Knowledge Bases

The purpose of this section is to explain how to formalise program-specific knowledge bases in FOL. Given a well-typed program  $\Pr g = \mathcal{K} \overline{CD} e$  with domain knowledge base  $\mathcal{K}$ , we first define a knowledge base  $\mathcal{K}^{code}$  that reflects part of the code of  $\Pr g$  in FOL. The knowledge base  $\mathcal{K}^{code}$  includes general axioms and symbols common to all FSRJ programs, as well as knowledge specifically derived from  $\Pr g$ , such as knowledge about declared classes and fields.

▶ Definition 7 (Code knowledge bases). Given a well-typed FSRJ program Prg with domain knowledge base  $\mathcal{K}$  over the signature  $\Sigma$ , let Cls = dom(Prg) be the set of class names in Prg, Fls(C) the set of qualified field names C.f declared (not inherited) in class C, type(C.f) the type of the field C.f, and  $\mathbb{Z}_k = \{z \mid -k \leq z \leq k\}$  (where k is natural number) a set of integers. We denote by  $\Sigma^{code}$  the following signature (where subscripts denote the arity of non-constant symbols):

© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:12–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

A1.  $\forall x. (x \doteq \mathsf{int}) \oplus (x \doteq \mathsf{null}) \oplus \mathsf{isInt}(x) \oplus \mathsf{isObj}(x) \oplus \mathsf{isCls}(x),$ A2. $\forall x. isCls(x) \Leftrightarrow (isStd(x) \oplus isAdp(x) \oplus isRcl(x))$ A3. $\mathsf{isInt}(n)$  for all  $n \in \mathbb{Z}_k$ isCls(c) for all  $c \in Cls$ A4.A5.isStd(c) for all  $c \in Cls$  such that Std(c)A6.isAdp(c) for all  $c \in Cls$  such that Adp(c)isRcl(c) for all  $c \in Cls$  such that Rcl(c)A7. $\forall x. \ \left( \wedge_{n \in \mathbb{Z}_k} \neg (x \doteq n) \right) \Rightarrow \neg \mathsf{isInt}(x)$ A8. $\forall x. \ \left( \land_{\mathsf{C} \in Cls} \neg (x \doteq \mathsf{C}) \right) \Rightarrow \neg \mathsf{isCls}(x)$ A9. $\forall x, y.$  subclass $(x, y) \Rightarrow (isCls(x) \land isCls(y))$ A10.A11. subclass(c, c') for all  $c, c' \in Cls$  such that c <: c'A12. $\neg$ subclass(C, C') for all C, C'  $\in$  Cls such that C  $\measuredangle$ : C' A13.  $\forall x, y. \text{ compatible}(x, y) \Leftrightarrow \exists z. \text{ subclass}(x, z) \land \text{subclass}(y, z) \land \text{isAdp}(z)$ A14.  $\forall x, y. (instOf(x, y) \lor in(x, y)) \Rightarrow (isObj(x) \land isCls(y))$  $\forall x. isObj(x) \Rightarrow \exists y. isCls(y) \land instOf(x, y) \land \neg isAdp(y)$ A15.A16.  $\forall x, y, z. \ (instOf(x, y) \land instOf(x, z)) \Rightarrow (y \doteq z)$  $\forall x, y. in(x, y) \Leftrightarrow \exists z.instOf(x, z) \land subclass(z, y)$ A17. A18.  $\forall x, y$ . hasType $(x, y) \Leftrightarrow (in(x, y) \lor (isInt(x) \land (y \doteq int)) \lor ((x \doteq null) \land isCls(y)))$ A19.  $\forall x. in(x, c) \Leftrightarrow \exists y. c_f(x, y) \text{ for all } c \in Cls \text{ and } c_f \in Fls(c)$ A20.  $\forall x, y, z$ .  $(C_f(x, y) \land C_f(x, z)) \Rightarrow (y \doteq z)$  for all  $c \in Cls$  and  $c.f \in Fls(c)$ A21.  $\forall x, y. \ C_f(x, y) \Rightarrow in(x, c) \land hasType(y, type(C_f)) \text{ for all } c \in Cls \text{ and } c.f \in Fls(c)$ 

**Figure 7** The code knowledge base.

 $\begin{aligned} \mathbb{Z}_k \cup \{\mathsf{int}, \mathsf{null}\} \cup \mathit{Cls} \cup \{\mathtt{C_{f_2}} \mid \mathtt{C.f} \in \mathit{Fls}(\mathtt{C})\} \\ \cup \{\mathsf{isInt}_1, \mathsf{isObj}_1, \mathsf{isCls}_1, \mathsf{isStd}_1, \mathsf{isAdp}_1, \mathsf{isRcl}_1, \mathsf{subclass}_2, \mathsf{in}_2, \mathsf{compatible}_2, \mathsf{instOf}_2, \mathsf{hasType}_2 \} \end{aligned}$ 

and (without loss of generality) assume that  $\Sigma \cap \Sigma^{\mathsf{code}} = \mathbb{Z}_k \cup \{\doteq\}$ . The code knowledge base  $\mathcal{K}^{\mathsf{code}}$  of  $\mathsf{Prg}$  is the set consisting of the formulas over  $\Sigma^{\mathsf{code}}$  in Figure 7.

We briefly discuss the formulas of  $\mathcal{K}^{\mathsf{code}}$ . Axiom A1 expresses that each element in the domain must be either the type name int, the value null, an integer value, an object or a class; and A2 that each class is either standard, adaptable or reclassifiable. Axioms A3 declare that each constant  $n \in \mathbb{Z}_k$  is an integer and A4 declare that each constant  $c \in Cls$  is a class. Axioms A5-A7 declare the kind (standard, adaptable, reclassifiable) of each constant  $c \in Cls$ . Axioms A8 express that each integer in the domain is mapped to a constant  $n \in \mathbb{Z}_k$  and A9 that each class in the domain is mapped to a constant  $c \in Cls$ . Axioms A10-A12 express that the subclass relation relates classes to classes and capture the <: relation of Prg. Axiom A13 expresses that two (reclassifiable) classes are compatible if and only if they are both subclasses of the same adaptable class. Axiom A14 expresses that both the instOf and the in relations relate objects with classes. Axioms A15 and A16 declare that each object is an instOf of exactly one non-adaptable class. Axiom A17 declares that an object is in a class if and only if it is an instOf of a subclass of that class. Axiom A18 declares that only null, integers and objects have types, and which are their types. Axioms A19-A21 express that the predicate C\_f models the field f of objects of any subclass of C, each object has exactly the fields (declared or inherited) in its class, and each field contains exactly one value of the declared type.

**Example 8** (Code knowledge base for the Greenhouse digital twin in FSRJ). A subset of the

<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0 39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:13–2:30 Leibniz International Proceedings in Informatics Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

code knowledge base for the FSRJ program in Example 1 (namely, the subset containing axioms A4 and A11) is given in Figure 2c.

We define the linkage knowledge base to express that each object satisfies the predicate specified by the **links** clause of its class.

▶ Definition 9 (Linkage knowledge bases). Let Prg be a well-typed FSRJ program with domain knowledge base  $\mathcal{K}$  over  $\Sigma$ . Let Cls = dom(Prg). The linkage knowledge base  $\mathcal{K}^{link}$  of Prg is the set of formulas over the signature  $\Sigma \cup \Sigma^{code}$  defined by  $\forall x$ . instOf $(x, C) \Rightarrow \phi[z := x]$  for all  $C \in Cls$  and  $Prg(C)(links) = \lambda z.\phi$ .

**Example 10** (Linkage knowledge base for the Greenhouse digital twin in FSRJ). The linkage knowledge base for the FSRJ program of Example 1 is given in Figure 3b.

Observe that by combining  $\Sigma$  and  $\Sigma^{code}$ , predicates in  $\mathcal{K}^{link}$  may express connections between an FSRJ program's runtime values (e.g., instances of a class c) and the domain knowledge base. We define a knowledge base that comprises the three knowledge bases defined above:

▶ Definition 11 (Program knowledge bases). Let Prg be a well-typed FSRJ program with domain knowledge base  $\mathcal{K}$  over the signature  $\Sigma$ . The program knowledge base of Prg is the knowledge base  $\mathcal{K}^{Prg} = \mathcal{K} \cup \mathcal{K}^{code} \cup \mathcal{K}^{link}$  over the signature  $\Sigma \cup \Sigma^{code}$ .

▶ Lemma 12 (Satisfiability of the program knowledge base). For every well-typed program Prg with domain knowledge base  $\mathcal{K}$  over the signature  $\Sigma$ . If  $\mathcal{K}$  is satisfiable, then the knowledge base  $\mathcal{K}^{\mathsf{Prg}} = \mathcal{K} \cup \mathcal{K}^{\mathsf{code}} \cup \mathcal{K}^{\mathsf{link}}$  over the signature  $\Sigma \cup \Sigma^{\mathsf{code}}$  is satisfiable.

#### 3.5 **Program Coherence**

This section considers requirements on FSRJ programs Prg to ensure that reasoning with the program knowledge base  $\mathcal{K}^{Prg}$  is in fact meaningful for declarative object reclassification. These requirements are needed to ensure type soundness; even with a satisfiable domain knowledge base, linking axioms can make the knowledge base incoherent (e.g., define a linkage predicate to introduce false into the knowledge base). Below, let  $\lambda z.\phi^{C} = Prg(C)(classifies)$ and Children(C) = {D  $\in dom(Prg) | Prg(D)(extends) = C$ } for classes C. We say that a well-typed program Prg with domain knowledge base  $\mathcal{K}$  over a signature  $\Sigma$  is *coherent* if: *Coh1:*  $\mathcal{K}$  is satisfiable. This requirement ensures that models exist for the domain knowledge base.

**Coh2:**  $\mathcal{K} \cup \mathcal{K}^{\text{code}} \models \mathcal{K}^{\text{links}}$ . This requirement ensures that the user-defined linkage axioms exclude no interpretation that satisfies the union of the domain and code knowledge bases

 $\mathcal{K} \cup \mathcal{K}^{\mathsf{code}}$ ; thus ensuring that each object satisfies the linkage predicate of its class.

**Coh3:** In every model satisfying  $\mathcal{K}^{Prg}$ , for every object x instance of a reclassifiable subclass of an adaptable class c, the classification predicate  $\phi^{c}$  of c must hold for x:

 $\mathcal{K}^{\mathsf{Prg}} \models \forall x. \text{ in}(x, \mathtt{C}) \Rightarrow \phi^{\mathtt{C}}[z := x], \text{ for all } \mathtt{C} \in Cls \text{ such that } \mathsf{isAdp}(\mathtt{C}).$ 

This requirement ensures that the classification-predicates in the knowledge base respect the subclass hierarchy in the program.

**Coh4:** In every model satisfying  $\mathcal{K}^{\mathsf{Prg}}$ , for every object x in an adaptable class  $\mathsf{c}$ , classification predicate  $\phi^{\mathsf{D}}$  of one of its subclasses  $\mathsf{D}$  must hold for x:

$$\mathcal{K}^{\mathsf{Prg}} \models \forall x. \text{ in}(x, \mathsf{C}) \Rightarrow \bigvee_{\mathsf{D} \in \mathrm{Children}(\mathsf{C})} \phi^{\mathsf{D}}[z := x], \text{ for all } \mathsf{C} \in \mathit{Cls} \, \mathrm{s.t.} \, \mathsf{isAdp}(\mathsf{C}).$$

This requirement ensures that for every instance of a subclass of c, the classification query of at least one of the subclasses of c always holds.



39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:14–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

**Coh5:** In every model satisfying  $\mathcal{K}^{Prg}$ , for every object x in an adaptable class C and subclass D of C, the **retrieves** predicate of D must hold for x when its **classifies** predicate holds:

 $\mathcal{K}^{\mathsf{Prg}} \models \forall x. (\mathsf{in}(x, \mathsf{C}) \land \phi^{\mathsf{D}}[z := x]) \Rightarrow ((\exists \overline{y}. \psi^{\mathsf{D}}[z\overline{z} := x\overline{y}]) \land (\forall \overline{y}. \psi^{\mathsf{D}}[z\overline{z} := x\overline{y}] \Rightarrow \mathsf{hasType}(\overline{y}, \overline{\mathsf{T}}))),$ for all  $\mathsf{C}, \mathsf{D} \in Cls$  such that  $\mathsf{isAdp}(\mathsf{C})$  and  $\mathsf{subclass}(\mathsf{D}, \mathsf{C})$ 

where  $\lambda z, \overline{z}.\psi^{\mathsf{D}} = \mathsf{Prg}(\mathsf{D})(\mathsf{retrieves})$ ,  $\mathsf{fields}(\mathsf{D}) = \overline{\mathsf{T}} \ \overline{\mathsf{f}}$  (for some  $\overline{\mathsf{f}}$ ), and  $\mathsf{hasType}(\overline{y},\overline{\mathsf{T}})$  is shorthand for  $\mathsf{hasType}(y_1,\mathsf{T}_1) \land \cdots \land \mathsf{hasType}(y_n,\mathsf{T}_n)$ . This requirement ensures that a reclassified object can be instantiated correctly.

Conditions **Coh1** and **Coh2** are basic requirements on the domain knowledge base  $\mathcal{K}$  and the linkage axioms. They ensure that reasoning over the combined knowledge base is possible. In contrast, **Coh3–Coh5** are required for type soundness (cf. Theorem 22): first **Coh3** ensures that subclassing in the program and the knowledge base are aligned; then **Coh4** ensures that reclassification always succeeds in finding a target class D; and finally **Coh5** ensures that values for the fields in the target class D can always be retrieved in the knowledge base and respect the declared types of the fields of class D in the program.

#### 3.6 Heap Lifting, Synchronisation and Operational Semantics

To model the operational semantics of FSRJ, we introduce the concepts of *address* (of an object in a heap), *stack*, *values*, *object*, *heap*, *lifted-heap knowledge base*, *environment*, *synchronisation knowledge base*, *overall knowledge base*, *runtime expressions* and *runtime configurations*. *Addresses*, ranged over by the metavariable  $\iota$ , are elements of the denumerable set *I*. *Stacks* are sequences  $\bar{\iota}$ , recording the addresses of objects on which methods under execution have been invoked (the address of the object of the most recent invocation is rightmost). Values, ranged over by the metavariable v, are either addresses, integers, or the special value null. *Objects* are tuples  $\langle C, \bar{f} = \bar{v} \rangle$ , where C is the class of the object,  $\bar{f}$  the names and  $\bar{v}$  the values of the fields. A heap  $\mathcal{H}$  maps addresses to objects; the empty heap is denoted by  $\emptyset$ . A value v occurs in  $\mathcal{H}$  if  $v \in \{\iota, \bar{v}\}$  for some  $\mathcal{H}(\iota) = \langle C, \bar{f} = \bar{v} \rangle$ , where  $\iota \in dom(\mathcal{H})$ .

▶ Definition 13 (Heaps for well-typed programs). A heap  $\mathcal{H}$  for a well-typed FSRJ program Prg is a mapping from addresses to objects such that

- $= \{\iota \mid \iota \text{ occurs in } \mathcal{H}\} = dom(\mathcal{H}), \text{ and}$
- $\blacksquare$   $\mathcal{H}$  only contains instances of non-adaptable classes defined in Prg.

The following definition formalises the concept of *lifting a heap* into a knowledge base.

▶ Definition 14 (Lifted-heap knowledge base). Given a heap  $\mathcal{H}$  for a well-typed FSRJ program Prg, the lifted-heap knowledge base  $\mathcal{K}^{\text{heap}}(\mathcal{H})$ , which lifts  $\mathcal{H}$  to FOL, is the set consisting of the following formulas over the signature  $\Sigma^{\text{heap}}(\mathcal{H}) = \Sigma^{\text{code}} \cup \text{dom}(\mathcal{H})$ :

L1. instOf( $\iota$ , c) for all  $\iota \in dom(\mathcal{H})$  such that  $\mathcal{H}(\iota) = \langle c, f_1 = v_1, \dots, f_n = v_n \rangle$ 

L2.  $C_{j}f_{i}(\iota, v_{j})$  for all  $\iota \in dom(\mathcal{H})$  and for all  $C_{j}f_{i}$  such that

$$\mathcal{H}(\iota) = \langle \mathsf{C}, \mathsf{f}_1 = v_1, \dots, \mathsf{f}_n = v_n \rangle, \ \mathsf{C} <: \mathsf{C}_j, \ \mathsf{C}_j.\mathsf{f}_j \in Fls(\mathsf{C}_j) \text{ and } 1 \leq j \leq n$$
  
L3.  $\forall x. \left( \bigwedge_{\iota \in \mathit{dom}(\mathcal{H})} (x \neq \iota) \right) \Rightarrow \neg \mathsf{isObj}(x)$ 

The following definition formalises the concept of stating the values of the current properties of the physical assets into a knowledge base.

▶ Definition 15 (Environment and synchronisation knowledge base). An environment  $\mathcal{E}$  for a well-typed FSRJ program Prg is a reference to a mapping that provides suitable



licensed under Creative Commons License CC-BY 4.0

<sup>39</sup>th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:15–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

current values for the properties of the assets described in the domain knowledge base of Prg. The synchronisation knowledge base  $\mathcal{K}^{\text{sync}}(\mathcal{E})$  of Prg provides a witness for each formula  $\exists x.\mathsf{ctx\_prop}(\mathsf{ctx\_asset}, x) \text{ in } \mathcal{K} \text{ which states that the property } \mathsf{ctx\_prop} \text{ of the asset } \mathsf{ctx\_asset}$ has some value x.

We now define a knowledge base that combines the two knowledge bases defined above, which model the dynamic part of the system (the program heap and the current values of the properties physical assets, respectively) with the domain and code knowledge base, which model the static part of the system (domain knowledge and the structure of the code, respectively). The combined knowledge base is operationalised in the operational semantics of FSRJ (see rules R-adapt-y and R-adapt-n in Figure 8).

**Definition 16** (Overall knowledge base). Given a heap  $\mathcal{H}$  and an environment  $\mathcal{E}$  for a welltyped FSRJ program Prg, the overall knowledge base  $\mathcal{K}^{\text{overall}}(\mathcal{H}, \mathcal{E}) = \mathcal{K}^{\text{Prg}} \cup \mathcal{K}^{\text{heap}}(\mathcal{H}) \cup \mathcal{K}^{\text{sync}}(\mathcal{E})$ of Prg, over the signature  $\Sigma^{\mathsf{Prg}} \cup \Sigma^{\mathsf{heap}}(\mathcal{H})$ , represents domain knowledge, code, heap and value of the asset properties for the program Prg with heap  $\mathcal{H}$  and environment  $\mathcal{E}$ .

We briefly discuss the formulas of  $\mathcal{K}^{\mathsf{heap}}(\mathcal{H})$ . The axiom L1 introduces individuals  $\iota$  as instances of classes c (Axiom A14 of Definition 7 then allows us to infer isObj( $\iota$ )). Axiom L2 populates the relation  $c_f$  for all fields f of the classes c. Axiom L3 states that only addresses from the heap can be considered as objects in the knowledge base  $\mathcal{K}^{\text{overall}}(\mathcal{H}, \mathcal{E})$ (thus, L3 is similar to the closure axioms A8 and A9 for integers and classes in Figure 7).

▶ Example 17 (Lifted-heap and synchronisation knowledge bases for the Greenhouse digital twin in FSRJ). The lifted-heap and synchronisation knowledge bases for the FSRJ program in Figure 3a and Example 1 after the first execution of the pu.plant=this.pl expression in the body of the loop method, are given in Figures 2b and 2d, respectively.

Recall that (by Lemma 12) the program knowledge base  $\mathcal{K}^{\mathsf{Prg}}$  of a well-typed program  $\mathsf{Prg}$ is satisfiable. The following theorem shows that, if  $\mathcal{K}^{\mathsf{Prg}}$  satisfies conditions Coh1 and Coh2for program coherence (see Section 3.5), then its extension  $\mathcal{K}^{\mathsf{overall}}(\mathcal{H}, \mathcal{E})$  with a lifted-heap knowledge base  $\mathcal{K}^{\mathsf{heap}}(\mathcal{H})$  and synchronization knowledge base  $\mathcal{K}^{\mathsf{sync}}(\mathcal{E})$  for Prg remains satisfiable (thus, reasoning over  $\mathcal{K}^{\mathsf{overall}}(\mathcal{H}, \mathcal{E})$  is possible).

**Theorem 18** (Satisfiability of the overall knowledge base). Let  $\mathcal{H}$  be a heap and  $\mathcal{E}$  and environment for the well-typed program Prg with domain knowledge base K is over the signature  $\Sigma$ , and let  $\mathcal{K}^{\mathsf{Prg}} = \mathcal{K} \cup \mathcal{K}^{\mathsf{code}} \cup \mathcal{K}^{\mathsf{link}}$  be the program knowledge base of  $\mathsf{Prg}$ . If  $\mathcal{K}$  is satisfiable and  $\mathcal{K} \cup \mathcal{K}^{\mathsf{code}} \models \mathcal{K}^{\mathsf{links}}$ , then the overall knowledge base  $\mathcal{K}^{\mathsf{overall}}(\mathcal{H}, \mathcal{E})$  is satisfiable.

*Runtime expressions* are ranged over by e. Their syntax is obtained from the syntax of expressions (cf. Figure 5) by removing variables x (recall that this is a variable) and adding addresses  $\iota$  and ret-expressions (of the form ret(e), that represent the execution of a method). A runtime configuration (configuration, for short) is a triple consisting of a heap, a stack and a runtime expression, i.e.,  $\mathcal{H} | \bar{\iota} | e$ . The reduction relation has the form  $\mathcal{H} | \bar{\iota} | e \to \mathcal{H}' | \bar{\iota}' | e'$ , where the configuration  $\mathcal{H}[\bar{\iota}]e$  gets reduced to  $\mathcal{H}'[\bar{\iota}']e'$  in one step. A configuration  $\mathcal{H}[\bar{\iota}]e$ that cannot be reduced is called a *normal form*. The initial configuration of a program is  $\emptyset \mid \bullet \mid \bullet$ , where  $\bullet$  is the main expression of the program. Remark that configurations only include stacks to ensure that the operational semantics can model that execution gets stuck when the reclassification of an object on which a method under execution has been invoked is attempted (i.e., that a configuration that is going to perform such a reclassification is a normal form).

The reduction system for FSRJ is given in Figure 8. The reduction rules use the auxiliary function mBody (m, C) (for method body lookup) and the congruence rule (which uses evaluation



<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:16-2:30

licensed under Creative Commons License CC-BY 4.0 39th European Conference on Object-Oriented Programming (ECOOP 2025).

Leibniz International Proceedings in Informatics

Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

#### Computation rules



**Figure 8** FSRJ operational semantics: computation rules, method body lookup, evaluation context, and congruence rule.

contexts) are also given in Figure 8. The rules R-field, R-new and R-assign for field access, object creation and assignment are standard. Rule R-invoke pushes the address of the called object to the stack and a ret-expression that stems from method lookup for the called method in the class of the called object, as the new runtime expression. Here, actual parameters and the self-reference are inlined by substitution on the runtime expression. Rule R-ret then pops the object address from the stack, returning the computed value as the new runtime expression.

In FSRJ an object is hot (see Section 2) if and only if its address occurs on the stack. Both rules for dynamic object reclassification, R-adapt-y and R-adapt-n, check whether an object needs to be reclassified and cause the execution to get stuck if an attempt to reclassify a hot object is made. Rule R-adapt-y applies if the object needs to be reclassified (i.e., it cannot be classified as a member of its current class); rule R-adapt-n applies otherwise. We now explain the interaction between execution and the knowledge base in these rules. Both rules check whether the object  $\iota$  of class D is cold and, if so, lift the heap  $\mathcal{H}$  to construct the knowledge base  $\mathcal{K}^{\text{overall}}(\mathcal{H}, \mathcal{E})$ . If the classification predicate  $\phi$  of D holds when instantiated with  $\iota$ , then rule R-adapt-n is applied and the object is not reclassified. Otherwise rule R-adapt-y is applied. Recall from Axiom A13 of Definition 7 that two classes are compatible if they have the same superclass. Thus, the premises of rule R-adapt-y identify a class D' that is compatible with D, for which the classification predicate  $\phi'$  holds when instantiated with  $\iota$ . The last premise ensures that the **retrieves**-predicate  $\psi'$  of D' holds for values  $\iota, \overline{v}'$ , and the resulting heap on the right-hand side binds  $\iota$  to an object of class D' in which the fields  $\overline{\mathbf{f}}$  are bound to  $\overline{v}'$ .

#### 3.7 Type Soundness for Coherent Programs

In order to prove type soundness for coherent FSRJ programs by a subject reduction theorem and a progress theorem for the small-step operational semantics, we need to formulate a type system for runtime expressions. Expressions containing either a *stupid selection*, i.e., a field selection null.f or a method invocation null.m(...), or a *stupid reclassification*, i.e.,



© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:17–2:30 Leibniz International Proceedings in Informatics



#### **Runtime expression typing** $\Theta \Vdash e : S$

**Figure 9** FSRJ run-time typing: runtime expressions, well-formed heaps and configurations.

an object reclassification adapt (null), are not well-typed according to the FSRJ (source level) type system. However, a runtime expression without stupid selections and stupid reclassifications may reduce to a runtime expression containing stupid selections or stupid reclassifications. Therefore, the type system for runtime expressions contains a rule for assigning to the value null either the type  $\perp$  (like in the source level typing) or any class name.

The type rules for runtime expressions are shown in Figure 9 (left); these rules are of the form  $\Theta \Vdash e : S$ , where the environment  $\Theta$  is a finite (possibly empty) mapping from addresses to class names. Figure 9 (bottom) also presents the notions of well-formed-heaps and well-formed configurations. The notion of well-formed heap ensures that the environment  $\Theta$  maps all the addresses in the heap into the class of the corresponding object and that, for every object stored in the heap, the fields of the object contain appropriate values. The notion of well-formed configuration ensures that the heap is well-formed, that the addresses in the stack are defined in the heap, that the runtime expression is well-typed and has a structure that is compatible with the stack — this last check is performed by exploiting evaluation contexts for FSRJ runtime expressions (see Figure 8).

Then, type soundness can be proved by using the standard technique of subject reduction and progress theorems [50]; however, we need to account for declarative object reclassification.

#### ▶ Lemma 19 (From expression typing to runtime expression typing). If $\bullet \vdash e : T$ , then $\bullet \Vdash e : T$ .

We now consider subject reduction for the execution of semantically reflected programs. Here, program coherence (Section 3.5) and Theorem 18 (which assumes **Coh1** and **Coh2** and is independent from *Coh3–Coh5*) are crucial for the case of R-adapt-y:



licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:18-2:30

Leibniz International Proceedings in Informatics

- an inconsistent (i.e., not satisfiable) knowledge base, or
- a consistent knowledge base that satisfies Coh1-Coh4 and does not satisfy Coh5 because the retrieved values for the fields have a wrong type

would allow us to, e.g., infer erroneous values for retrieves  $\lambda$ -predicates (see the premise  $\mathcal{K}^{\text{overall}}(\mathcal{H}, \mathcal{E}) \models \psi'[z\overline{z} := \iota \overline{\nu}']$  of rule R-adapt-y), potentially leading to typing errors for runtime configurations after reclassification. Lifting a heap that is not well-typed, could make the knowledge base inconsistent (see Condition L2 in Definition 14 and Axiom A21 in Figure 7). Thus, well-typed runtime configurations ensure consistent overall knowledge bases, which in turn (together with **Coh5**) ensure that (whenever rule R-adapt-y does not get stuck) dynamic object reclassification produces well-typed runtime configurations.

▶ Theorem 20 (Subject reduction). If  $\Theta \Vdash \mathcal{H} | \overline{\iota} | e : S$  and  $\mathcal{H} | \overline{\iota} | e \to \mathcal{H}' | \overline{\iota}' | e'$  then there exists  $\Theta' \supseteq \Theta$  such that  $\Theta' \Vdash \mathcal{H}' | \overline{\iota}' | e' : S'$  for some S' such that  $S' \leq S$ .

Progress relies on program coherence, but not on consistency of the knowledge base. Concretely, Coh3-Coh5 guarantee that there exists a class D' such that the following three premises of rule R-adapt-y hold:

$$\mathcal{K}^{\mathsf{code}} \models \mathsf{compatible}(\mathtt{D}, \mathtt{D}') \qquad \mathcal{K}^{\mathsf{overall}}(\mathcal{H}, \mathcal{E}) \models \phi'[x := \iota] \qquad \mathcal{K}^{\mathsf{overall}}\mathcal{H}, \mathcal{E}) \models \psi'[z\overline{z} := \iota \overline{v}']$$

The last case in the statement of the theorem states that the adaptation of hot objects gets stuck (cf. DC3 in Section 4).

▶ Theorem 21 (Progress). Let  $\mathcal{H} | \bar{\iota} | e$  be a well-typed normal form. Then

- **1.** either e is a value and  $\overline{\iota} = \bullet$ ; or
- **2.** for some evaluation context E we can express e as
  - **a.** either  $E[\mathsf{null.} f]$  for some f, or
  - **b.**  $E[\textit{null.m}(\overline{v})]$  for some *m* and  $\overline{v}$ , or
  - c. E[null.f = v] for some f and v, or
  - **d.** E[adapt(null)], or
  - e.  $E[adapt(\iota)]$  for some  $\iota \in \overline{\iota}$ .

Type soundness follows from Lemma 19 and Theorems 20 and 21. After Theorem 20, this theorem is a second place where consistency (needed only for Theorem 20) and program coherence (needed for Theorem 21) interact: We need the overall knowledge base  $\mathcal{K}^{\text{overall}}(\mathcal{H}, \mathcal{E})$ to be consistent so that Coh3-Coh5 indeed select the correct class D' to reclassify.

▶ Theorem 22 (Type soundness). Let  $Prq = \mathcal{K} \overline{CD} e_0$  be well-typed and coherent FSRJ program such that  $\bullet \vdash e_0 : T_0$ , and  $\emptyset \mid \bullet \mid e_0 \to^* \mathcal{H} \mid \overline{\iota} \mid e \text{ with } \mathcal{H} \mid \overline{\iota} \mid e \text{ being a normal form. Then } e \text{ is}$ 

- 1. either an integer n and  $T_0$  is int and  $\bar{\iota} = \bullet$ ; or
- **2.** null and  $T_0$  is either  $\perp$  or a class name and  $\bar{\iota} = \bullet$ ; or
- **3.** an address  $\iota$  such that  $\mathcal{H}(\iota) = \langle C, \cdots \rangle$  with  $C \leq T_0$  and  $\overline{\iota} = \bullet$ ; or
- 4. an expression containing either null f or null. $m(\overline{v})$  or or null f = v or adapt(null) or  $adapt(\iota)$  for some  $f, m, \overline{v}, v$  and  $\iota$  such that  $\iota \in \overline{\iota}$ .

#### 4 Discussion: Design Choices, Implementation and Performance

This section discusses design choices for declarative dynamic object reclassification. We first consider overall design choices for the declarative reclassification mechanism (as reflected in FSRJ), then design choices related to the implementation of declarative dynamic object reclassification (as reflected in our prototype implementation), and finally the implementation itself.



© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:19-2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

### 4.1 Overall Design Choices

Our main goal with declarative dynamic object reclassification is to realise a separation of concerns and decoupling between the *application logic*, which describes program behaviour within a given context, and the *adaptation logic*, which describes how the behaviour of the program changes according to context changes. To focus on the basic interaction between adaptation and application logic, we aimed for a simple programming construct for dynamic object reclassification and considered the following design choices:

**DC1**: The adaptation logic is expressed in a declarative way, leveraging domain knowledge. **DC2**: The application logic is expressed by standard class-based object-oriented code.

DC3: Adaptation works on (cold) objects in isolation and hot object adaptation gets stuck.

Concerning DC1, to focus on the general mechanisms of semantic reflection and declarative dynamic object reclassification, FSRJ formalises the knowledge base as a set of FOL formulas. Observe that FSRJ does not impose any restrictions on queries to the knowledge base. For example, **classifies**-predicates may depend on the class of other objects; thus, the reclassification of one object may lead to another object becoming inconsistent. We believe that this flexibility is convenient in a digital twin setting; e.g., a change of requirement monitor for a plant may trigger a change in the controller for the water pump. An alternative design choice would be to restrict queries to, e.g., only depend on the external context (i.e., the domain and synchronisation knowledge). In this case, we avoid the reclassification of one object triggering the reclassification of another, but we could potentially reclassify one object too early, violating the program logic (e.g., changing the watering policy before changing the humidity monitor in the greenhouse). Furthermore, self-adaptation in FSRJ is not deterministic: an object may be reclassifiable to more than one class. If determinism is desired, one can add the additional constraint that the classification queries of its reclassifiable subclasses of all adaptable classes are pairwise disjoint.

Concerning DC2, it can be worth observing that (1) the links, classifies and retrieves clauses could be expressed as, e.g., program annotations or even in a separate file; (2) typing just enforces a programming pattern that treats an adaptable class as an abstract class and its reclassifiable subclasses as final classes that cannot be used as types; and (3) the adapt expression could be implemented as a library function.

Concerning DC3, there are several ways to relax this restriction. First, the reclassification of a hot object could be delayed until the object becomes cold (this raises the issue of whether the ordering between delayed adaptations of different objects should be preserved). Second, the **adapt** expression could be replaced or complemented by a class-level adaptation expression to trigger the adaptation of all objects of a class. Third, the **adapt** expression could be externalised; e.g., replaced or complemented by an adaptation triggered by the runtime system, which performs adaptation on all objects of adaptable classes, and delays adaptation for hot objects until they become cold. Relaxing DC3 suggests several interesting directions for further work, including

- a type and effect system [3,57] to ensure that hot object adaptation cannot occur;
- coordinated adaptation of object aggregates (i.e., groups of interconnected objects);
- reclassifiable classes extended by other reclassifiable classes (cf. [15, 20]); and
- adaption of hot objects, e.g., the expression adapt(this) (cf. [15,20]) or a method defined in a class c with an expression adapt(e) where e has type c (cf. [16]).

For the last point, the adaptation logic could depend on the values of fields in the target object; when the object is cold, reclassification might rely on these values satisfying some



39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:20–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

class invariant (following a program verification methodology, e.g., [1,48]). So, hot object adaptation should require to devise a suitable invariant that the object must satisfy when, e.g., a specific occurrence of the expression adapt(this) is executed. For this reason, we decided that even rule R-adapt-n gets stuck on hot objects.

### 4.2 Implementation and Performance

**Implementation.** To interact with context in an implementation, we need to be able to check entailment and satisfiability of first-order formulas, and to retrieve witnesses when evaluating the satisfiability of  $\lambda$ -predicates (for **retrieves**-queries). Whereas such functionality is available in, e.g., satisfiability modulo theory solvers, the querying mechanism for declarative dynamic object reclassification would additionally benefit from decidability and availability of formalised domain knowledge. Hence, we consider two possibilities: ontology-based knowledge graphs and logic programming for **DC2**.

- Knowledge Graphs are widely used triple-based data formalisms with numerous variants, of which the RDF/OWL stack is well-suited for reclassification. First, these knowledge graphs are used in numerous industrial ontologies and is an established, core technology for autonomous systems, e.g., in digital twins or robotics [7,42]. Thus, they enable reuse and provide general knowledge across industries, especially in engineering. Second, these knowledge graphs have a foundation in Description Logics [5], which are decidable fragments of FOL with efficient reasoners and advanced tool-supported pragmatics for non-experts [28].
- Logic Programming is widely used and perhaps more rooted in programming and artificial intelligence than knowledge graphs, and also used for autonomous systems [7,32]; numerous implementations are available for knowledge representation and reasoning, e.g., Golog [45]. Logic programming has several variants, such as Prolog-style languages or Answer Set Programming [10, 46], which easily produce the witnesses we need for our queries. Logic programming is based on rules that encode domain knowledge, and powerful reasoning engines to achieve high efficiency. However, they are less prevalent for ontologies (i.e., general knowledge with few or no individuals), and seem better suited for concrete planning tasks.

We aim to leverage domain knowledge for self-adaptation in autonomous systems, e.g., digital twins. Observe that FSRJ's knowledge base and queries fit into Description Logic (specifically, the logic SROIQ underlying OWL 2 [30]). In particular, program coherence (Section 3.5) here maps to *decidable reasoning operations*: *Coh1* expresses satisfiability of a SROIQ knowledge base, *Coh2* the entailment of two sets of SROIQ axioms, and *Coh3–Coh5* the entailment of a subconcept relation from a set of axioms. In Description Logic syntax, we have  $\exists in.\{C\} \sqsubseteq cl_C$  (*Coh3*) and  $\exists in.\{C\} \sqcap cl_C \sqsubseteq \bigsqcup_D cl_D$  (*Coh4*), where  $cl_C$  is the classification concept of class C. *Coh5* is a special case of type checking knowledge graph access under semantic reflection [39]. Hence, our implementation uses knowledge graphs that connect to existing ontologies.

Further, an implementation of declarative dynamic object reclassification needs to realise semantic lifting to integrate program and knowledge base, and dynamic reclassification via adapt-expressions. For these reasons, our implementation is based on SMOL [38], an interpreted research language used to explore connections between programming languages and knowledge graphs, that natively supports heap lifting into a knowledge base. The language is restricted to Description Logics and OWL-based knowledge graphs. The backend uses HermiT [24] for reasoning and Apache Jena (https://jena.apache.org/) for graph data management. SMOL was designed to support semantic reflection for the implementation of digital twins [40], but has not yet offered direct programmatic support for reclassification as proposed in this paper.



OWL

ctx:NormalOps EquivalentTo: ctx:Pump and (ctx:temp some xsd:double[<= 50.0])
ctx:Overheating EquivalentTo: ctx:Pump and (ctx:temp some xsd:double[> 50.0])
FunctionalDataProperty(ctx:id)

(a) General knowledge as OWL ontology \_\_SMOL\_\_\_\_\_

```
1 abstract class Pump (domain String id, Int GpioPin, Plant plant)
2 links "a ctx:Pump";
3 Unit pump() /* ... */ end end
4 class Normal extends Pump () classifies "<ctx:NormalOps>";
5 /* methods */ end
6 class Overheating extends Pump (Int maximal) classifies "<ctx:Overheating>";
7 retrieves "?this ctx:hasMaintananceProfile [ctx:hasMaximalPower ?maximal]";
8 /* methods */ end
```

(b) SMOL code

RDF

```
1 ctx:pump1 a ctx:Pump; ctx:id "2"; ctx:temp 52.5. run:ob1 a prog:Seedling.
2 run:ob2 a prog:NormalPump; prog:id "2"; prog:GpioPin "7", prog:plant run:ob1.
3 run:ob2 smol:links [a ctx:Pump; ctx:id "1"].
```

(c) Lifting in RDF

**Figure 10** Instantiation of the pump in the Greenhouse digital twin in the SMOL implementation.

Our implementation, extending SMOL, goes slightly beyond the design decisions of FSRJ concerning minimality. The main differences are: (1) an explicit constructor that is called after reclassification to establish the classification predicate of the new class; (2) links, classifies and retrieves clauses are inherited; and (3) a more elaborate lifting mechanism, which can be encoded in the one of FSRJ. The constructor ensures a form of reclassification stability, while inheritance of classification predicates ensures a form of behavioural subtyping corresponding to *Coh3* (see Section 3.5). Our implementation supports both the adapt-expression of FSRJ and an option for reclassification to be triggered periodically by the runtime system. For static program analysis, we exploit the decidability of description logics underlying knowledge graphs to ensure program coherence. Figure 10 gives the SMOL version of our running example, as well as the knowledge graph used as the knowledge base.

**Performance.** The most expensive operations of FSRJ are on the lifted heap. Experimental results [35] that compare querying a description logic knowledge base with lifted heaps to querying a representation of the same knowledge expressed in program code, show that reasoning in the description logic knowledge base scales significantly better than reasoning with the knowledge represented in the program, both for more complex knowledge bases (i.e., more possible classification targets) and for more elements (i.e., more constants or individuals in the knowledge graph). The direct representation in the program only performs better for simple (<10 classification targets) and small (<2000 individuals) knowledge base every time the lifted heap is accessed. Instead it uses *virtualisation* [59]: the implementation determines which parts of the knowledge base are relevant for a query, and only generates this part.



@ Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:22-2:30

Leibniz International Proceedings in Informatics

#### 5 **Evaluation of the Prototype Implementation**

To evaluate the prototype implementation of declarative dynamic object reclassification, we aim to demonstrate that the reclassification works in practice as expected from the calculus. Concretely, we consider the following research questions:

**RQ1**: Does derivation in the knowledge base identify the correct class for reclassification? *RQ2:* Does derivation in the knowledge base instantiate state correctly during reclassification?

These RQs address essential aspects of the declarative reclassification process, namely that the classifies-and retrieves-predicates are resolved correctly in the knowledge base at runtime.

**Experimental Design and Setup.** This evaluation is based on GreenhouseDT [41], an opensource digital twin exemplar of a mini greenhouse developed in SMOL as a sensor-driven system that controls the temperature and humidity on different shelves of a greenhouse, as well as the soil moisture level of the plants on these shelves. The digital twin uses these sensor readings to determine the control policy of pumps (the actuators of the digital twin) for watering purposes. In our experiments, we focus on the behaviour of the actuators, i.e., the pumps used to water the plants, and use declarative object reclassification to dynamically determine which controller implementation to use for the pumps and to adopt the controller objects accordingly.

We extend GreenhouseDT to capture the following condition-based maintenance scenario: to ensure that pumps work properly:

- 1. the pumps need to operate within a given temperature range, and
- 2. the pumps should not exceed their expected life time in hours of operation.

If the temperature in an old pump exceeds the temperature thresholds, the pumps could get damaged. For that reason, we implement an extension of the motivating example discussed in Section 2, in which pumps can have one of the following distinct modes of behaviour: normal operating mode for pumps that work normally, overheating mode for pumps whose temperature is too high, underheating mode for pumps whose temperature is too low, and maintenance mode for pumps whose temperature is outside the operating range and that have exceeded the expected life span (and therefore need to be stopped). We model the different behavioural modes as reclassifiable subclasses of an adaptable class Pump. The exact values for the temperature and age thresholds depend on the type of pump considered; thus, the exact reclassification parameters differ, depending on the specifications of the different pumps. These extensions to the example are realistic as different actuators typically come with different specifications. The extensions also add complexity to the adaptation logic; it is easy to see that as this complexity increases, then a clear separation of concerns between the adaptation logic and the program logic also becomes increasingly attractive, as promoted in the approach of this paper.

Our evaluation is performed in terms of two sets of experiments: in the first set, we provide different temperature measurements to evaluate how the system adapts pump controller objects in response to changes in the temperature. We considered scenarios with three pumps with different specifications, and three temperature measurements for each pump: one below the specified temperature range, one above the range and one within the range. In the second set, we additionally let the age of the pumps change to capture a more complex reclassification scenario. In this set, we expanded the previous scenarios with additional age update events that enable the reclassification. In the experiments, we opted to keep the knowledge base invariant; i.e., the actual components that constitute the greenhouse remain the same over time in the domain knowledge base, and only the behaviour of the pump controllers in the virtual layer of the digital twin will change and will require dynamic



<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:23-2:30

Leibniz International Proceedings in Informatics

Pump ID	Model	Time (H)	Temperature	Expected Class	Actual Class			
After Instantiation								
Pump $1$	R385	0	5	Operational	Operational			
Pump $2$	WPS27	30	80	Overheating	Overheating			
Pump 3	R365	3000	94	Maintenance	Maintenance			
After First Reclassification								
Pump 1	R385	0	-5	Underheating	Underheating			
Pump $2$	WPS27	30	41	Overheating	Overheating			
Pump 3	R365	2000	50	Operational	Operational			
After Second Reclassification								
Pump 1	R385	2500	-5	Maintenance	Maintenance			
Pump $2$	WPS27	30000	41	Maintenance	Maintenance			
Pump 3	R365	300	1	Underheating	Underheating			

**Table 1** Pump states based on input temperature and reclassification cycles.

object reclassification in the digital twin. For RQ1, we then inspected the class targeted by the reclassification in each experiment and, for RQ2, the resulting object state.

**Results.** Table 1 illustrates the results from the experiments, comparing the expected class of each pump with the actual class after running **adapt**-expressions on objects with different pump specifications in the domain knowledge base in different scenarios. The actual class of each pump in the experiments is retrieved via the API, whereas the expected class of the pumps in the system is determined from the specifications of the different pumps, combined from constraints for the adaptation logic as expressed in the domain knowledge base, and the actual state of the objects in the system after the reclassification process.

To answer **RQ1**, our experiments used different inputs to adapt the Pump objects in the model. When calling **adapt** on an object, the system changes the class of the object depending on the input event, the object state and the domain knowledge base. We checked that the reclassification process worked as formalised in FSRJ by comparing the returned class to the expected ones after running the **adapt** on the Pump objects. We ran the two sets of experiments described above on our prototype implementation, in all cases the dynamic object reclassification determined by the **classifies**-predicate selected the correct reclassifiable subclass of the Pump class.

To answer RQ2, we ran the same sets of experiments as for RQ1, and inspected the state of the Pump objects after the reclassification process, comparing the values of the fields to the expected ones. In all the conducted experiments, the actual state of the Pump objects was consistent with the expected state, demonstrating that the system was able to derive correct values for the **retrieves**-predicates from the domain knowledge base at runtime. The second set of experiments showed that when considering more complex reclassification constraints, our prototype implementation was able to find a correct state.

**Threats to Validity.** The evaluation of the prototype implementation of declarative dynamic object reclassification has been done on a specific system, **GreenhouseDT**. Although the adaptation logic involves multiple classes in the knowledge base, the solutions to knowledge base queries have not been very complex. In particular, queries have deterministic solutions in our evaluation setup. We do not foresee particular challenges with non-deterministic queries,



© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:24–2:30

Leibniz International Proceedings in Informatics

as any solution would then be correct (although not necessarily optimal). Furthermore, although the queries in our evaluation involve multiple constraints, they do not involve complex constraint solving (i.e., constraints with many variables). Although well-suited for knowledge representation, knowledge graphs are not particularly strong at complex constraint solving. As our prototype implementation relies completely on the decidable fragments of FOL available in knowledge graphs to answer the queries, the prototype implementation could potentially fail to address RQ2 in cases where solutions might exist. Remark that we do not see this limitation of the prototype as a limitation of the proposed approach, as constraint solvers such as SMT could be integrated in a more refined query-answering mechanism.

### 6 Related Work

We focus this discussion of related research on two complementary aspects of our work: dynamic object reclassification techniques and programming with external domain knowledge. To the best of our knowledge, the combination of these aspects is novel to our paper.

There is a body of work on dynamic reclassification and related programming mechanisms (e.g., [2,8,13–16,20,26]), that addresses "self-extension" [13] and extensions of Java and use of the JVM [8,14,16,26], to the design of novel programming languages [2,20]. This line of work shares the following limitation: the requirements for the reclassification process are handled internally, as part of the program. In contrast, our paper proposes a novel approach to dynamic reclassification, in which we shift the reclassification process from ensuring internal consistency, with an operational description, to external consistency with a declarative description. This way, our work enables a separation of concerns between the adaptation logic and the application logic of the program, that has not been considered in previous work. In contrast to the works mentioned above, we have considered a very basic programming construct for dynamic object reclassification that does not require an advanced type (and effect) system.

Dynamic Software Updating (DSU) [27, 49, 51, 55] addresses runtime changes to the program code, such as software patching and task updates in real-time. DSU leverages techniques such as aspect-oriented programming [52], or changes Java program behaviour [49, 51] at the JVM or Bytecode level by modifying the code base or the instantiated objects, without halting execution. As code elements change, ensuring consistency becomes a major concern: new objects must allow the system to remain consistent with the specification. This can be solved by modifying old objects or by enforcing type renaming [27]. For DSU on asynchronous systems, statically collected type constraints can be enforced at runtime, delaying upgrades until they are type-safe [33]. Compared to reclassification mechanisms, including our work, DSU addresses externally triggered runtime changes to the program code, while dynamic reclassification programmatically changes the class of objects based on changes in the heap or context, but does not modify the class table. Although software updates are triggered externally, DSU does not consider external consistency as done in our work.

Programming with external domain knowledge is common for self-adaptive systems [11, 44, 58] in, e.g., robotics [7,9] and digital twins [17], but these systemic approaches generally lack the guarantees that can be provided by language abstractions. Semantic reflection, originally introduced in SMOL [38], has been used for many purposes in digital twins [40], including to repair external consistency by means of self-adaptive techniques [35–37, 40, 41], but so far the adaptation logic had to be programmed by hand. Semantic reflection has also been applied in digital twin architectures [23] to detect internal inconsistencies that trigger reconfigurations, and to drive domain-aware simulations [53]. Whereas this line of work exploits the lifting of runtime states, the self-adaptation has been concerned with the



licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:25–2:30

Leibniz International Proceedings in Informatics

composition of components and not with changing their internal behaviour. In all cases, the adaptation logic was implemented manually without language or type support. Our work overcomes this limitation by means of linguistic support for declarative reclassification.

The adaptation and retrieval queries of our paper can be seen as a form of LINQ [47], used for reflection with knowledge bases for the underlying serialisation. We are not aware of previous applications of such techniques for reclassification. Golog [45] uses FOL to examine and pick elements from its own state and allows decidable model checking when restricted to description logics [6], but has no type or class system. The systems of Fagin et al. [18] and Calvanese et al. [12] include knowledge bases that can be manipulated and updated by the program through epistemic operators, but have no type or class system either. Here, restrictions to description logics gives tractable knowledge base revision [21,22,43]. We are not aware of other work on core calculi for semantically reflected programs, as studied in our paper.

#### 7 **Conclusion and Future Work**

This paper presents declarative dynamic object reclassification and its formalisation in Featherweight Semantically Reflected Java, a minimal core calculus that formalises declarative reclassification in terms of interaction between a program's runtime configurations, domain knowledge and environment. We further show how to implement declarative dynamic object reclassification, and how to use these features in practice for adaptable digital twins. A key advantage of this approach is that developers can specify the reclassification process in a declarative way, disentangled from the business code, using queries to determine when and how an object should be reclassified. This makes it easier to reason about the reclassification process, and to ensure that the system remains consistent with an external context. Unlike previous approaches, declarative reclassification is based on external consistency; i.e., reclassification is driven by external domain knowledge, rather than by the internal state of the system. This makes our approach better suited for applications that must adapt to changes in the environment; it further provides a separation of concerns between the program's business code and its adaptation logic.

Some limitations and possible lines for future work are mentioned in Section 4. This paper does not address the issue of ensuring that the reclassification process is triggered at the right time; instead, reclassification is explicitly invoked. An implicit trigger might extend the range of applications for declarative reclassification; e.g., reclassification could be automatically triggered when certain conditions are met, at fixed time intervals, or during the lifting process. Such systems would be more dynamic and responsive to changes in the environment. When doing so, it would be crucial to ensure that the system is "free" to act; e.g., the target objects are cold or wait until reclassification is safe. Such constraints point towards DSU and asynchronous update mechanisms; in fact, integrating such mechanisms with a declarative adaptation logic would be an interesting extension of our work. Another interesting line of work is to consider more flexible class hierarchies, e.g., by allowing adaptable subclasses or reclassifiable subclasses of reclassifiable classes, as well as static restrictions to ensure progress, such as type and effect systems to eliminate reclassification attempts on hot objects. Finally, this paper focused on adaptation of cold objects in isolation. In future work we would like to generalise declarative reclassification to reclassification of hot objects and to synchronised reclassification of multiple objects (e.g., all instances of a class or object aggregates) tagged for reclassification, while maintaining the separation of concerns and type soundness of the language.



© Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025). Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:26-2:30

Leibniz International Proceedings in Informatics

#### — References

- 1 Wolfgang Ahrendt, Bernhard Beckert, Richard Bubel, Reiner Hähnle, Peter H. Schmitt, and Mattias Ulbrich, editors. *Deductive Software Verification - The KeY Book - From Theory to Practice*, volume 10001 of *Lecture Notes in Computer Science*. Springer, 2016. doi:10.1007/978-3-319-49812-6.
- 2 Jonathan Aldrich, Joshua Sunshine, Darpan Saini, and Zachary Sparks. Typestate-oriented programming. In Proc. 24th Conference Companion on Object Oriented Programming Systems Languages and Applications (OOPSLA 2009), page 1015–1022. ACM, 2009. doi:10.1145/ 1639950.1640073.
- 3 Torben Amtoft, Hanne Riis Nielson, and Flemming Nielson. Type and effect systems behaviours for concurrency. Imperial College Press, 1999.
- 4 Peter B. Andrews. An introduction to mathematical logic and type theory: to truth through proof, volume 27 of Applied Logic Series. Kluwer Academic Publishers, 2 edition, 2002. doi:10.1007/978-94-015-9934-4.
- 5 Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, 2003. doi:10.1017/CB09780511711787.
- 6 Franz Baader and Benjamin Zarrieß. Verification of Golog programs over description logic actions. In Pascal Fontaine, Christophe Ringeissen, and Renate A. Schmidt, editors, Proc. 9th International Symposium on Frontiers of Combining Systems (FroCoS 2013), volume 8152 of Lecture Notes in Computer Science, pages 181–196. Springer, 2013. doi:10.1007/978-3-642-40885-4\\_12.
- 7 Michael Beetz, Daniel Beßler, Andrei Haidu, Mihai Pomarlan, Asil Kaan Bozcuoglu, and Georg Bartels. Know Rob 2.0 - A 2nd generation knowledge processing framework for cognitionenabled robotic agents. In Proc. International Conference on Robotics and Automation (ICRA 2018), pages 512–519. IEEE, 2018. doi:10.1109/ICRA.2018.8460964.
- 8 Lorenzo Bettini, Sara Capecchi, and Betti Venneri. Extending Java to dynamic object behaviors. In Viviana Bono and Michele Bugliesi, editors, Proc. Workshop on Object Oriented Developments (WOOD 2003), volume 82 (8) of Electronic Notes in Theoretical Computer Science, pages 33–52. Elsevier, 2003. doi:10.1016/S1571-0661(04)80801-6.
- 9 Darko Bozhinoski, Mario Garzon Oviedo, Nadia Hammoudeh Garcia, Harshavardhan Deshpande, Gijs van der Hoorn, Jon Tjerngren, Andrzej Wąsowski, and Carlos Hernández Corbato. MROS: runtime adaptation for robot control architectures. Advanced Robotics, 36(11):502–518, 2022. doi:10.1080/01691864.2022.2039761.
- 10 Gerhard Brewka, Thomas Eiter, and Miroslaw Truszczynski. Answer set programming at a glance. Commun. ACM, 54(12):92–103, 2011. doi:10.1145/2043174.2043195.
- 11 Yuriy Brun, Giovanna Di Marzo Serugendo, Cristina Gacek, Holger Giese, Holger M. Kienle, Marin Litoiu, Hausi A. Müller, Mauro Pezzè, and Mary Shaw. Engineering self-adaptive systems through feedback loops. In Betty H. C. Cheng, Rogério de Lemos, Holger Giese, Paola Inverardi, and Jeff Magee, editors, Software Engineering for Self-Adaptive Systems, volume 5525 of Lecture Notes in Computer Science, pages 48–70. Springer, 2009. doi: 10.1007/978-3-642-02161-9\\_3.
- 12 Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati. Actions and programs over description logic knowledge bases: A functional approach. In Gerhard Lakemeyer and Sheila A. McIlraith, editors, *Knowing, Reasoning, and Acting: Essays in Honour of Hector J. Levesque.* College Press, 2011.
- 13 Alberto Ciaffaglione, Pietro Di Gianantonio, Furio Honsell, and Luigi Liquori. A protoypebased approach to object evolution. J. Object Technol., 20(2):4:1–24, 2021. doi:10.5381/JOT. 2021.20.2.A4.
- 14 Tal Cohen and Joseph Gil. Three approaches to object evolution. In Ben Stephenson and Christian W. Probst, editors, Proc. 7th International Conference on Principles and Practice of Programming in Java (PPPJ 2009), pages 57–66. ACM, 2009. doi:10.1145/1596655.1596665.

39th European Conference on Object-Oriented Programming (ECOOP 2025).

<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

licensed under Creative Commons License CC-BY 4.0

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:27–2:30 Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

- 15 Sophia Drossopoulou, Ferruccio Damiani, Mariangiola Dezani-Ciancaglini, and Paola Giannini. Fickle : Dynamic object re-classification. In Proc. 15th European Conference, on Object-Oriented Programming (ECOOP 2001), volume 2072 of Lecture Notes in Computer Science, pages 130–149. Springer, 2001. doi:10.1007/3-540-45337-7\\_8.
- 16 Sophia Drossopoulou, Ferruccio Damiani, Mariangiola Dezani-Ciancaglini, and Paola Giannini. More dynamic object reclassification: Fickle<sub>||</sub>. ACM Trans. Program. Lang. Syst., 24(2):153–191, 2002. doi:10.1145/514952.514955.
- 17 Romina Eramo, Francis Bordeleau, Benoit Combemale, Mark van Den Brand, Manuel Wimmer, and Andreas Wortmann. Conceptualizing digital twins. *IEEE Software*, 39(2):39–46, 2021. doi:10.1109/MS.2021.3130755.
- 18 Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. Knowledge-based programs. Distributed Comput., 10(4):199–225, 1997. doi:10.1007/S004460050038.
- 19 Kevin Feichtinger, Karl Kegel, Romain Pascual, Uwe Aßmann, Bernhard Beckert, and Ralf H. Reussner. Towards formalizing and relating different notions of consistency in cyber-physical systems engineering. In Proc. 27th International Conference on Model Driven Engineering Languages and Systems, MODELS Companion 2024. ACM, 2024. doi:10.1145/3652620. 3688565.
- 20 Ronald Garcia, Éric Tanter, Roger Wolff, and Jonathan Aldrich. Foundations of typestateoriented programming. ACM Trans. Program. Lang. Syst., 36(4):12:1–12:44, 2014. doi: 10.1145/2629609.
- 21 Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On the update of description logic ontologies at the instance level. In *Proc. 21st Conference* on Artificial Intelligence (AAAI 2006), pages 1271–1276. AAAI Press, 2006. URL: http: //www.aaai.org/Library/AAAI/2006/aaai06-199.php.
- 22 Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On the approximation of instance level update and erasure in description logics. In *Proc. 22nd Conference on Artificial Intelligence (AAAI 2007)*, pages 403–408. AAAI Press, 2007. URL: http://www.aaai.org/Library/AAAI/2007/aaai07-063.php.
- 23 Santiago Gil, Eduard Kamburjan, Prasad Talasila, and Peter Gorm Larsen. An architecture for coupled digital twins with semantic lifting. Software and Systems Modeling, Nov 2024. doi:10.1007/s10270-024-01221-w.
- 24 Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. Hermit: An OWL 2 reasoner. J. Autom. Reason., 53(3):245–269, 2014. doi:10.1007/S10817-014-9305-1.
- Barbara Hayes-Roth. A blackboard architecture for control. Artificial intelligence, 26(3):251–321, 1985. doi:10.1016/0004-3702(85)90063-3.
- 26 Chengwan He, Zhijie Nie, Bifeng Li, Lianlian Cao, and Keqing He. Rava: Designing a Java extension with dynamic object roles. In Proc. International Symposium and Workshop on Engineering of Computer Based Systems (ECBS 2006), pages 453–459. IEEE Computer Society, 2006. doi:10.1109/ECBS.2006.57.
- 27 Michael W. Hicks, Jonathan T. Moore, and Scott Nettles. Dynamic software updating. In Michael Burke and Mary Lou Soffa, editors, Proc. Conference on Programming Language Design and Implementation (PLDI 2001), pages 13–23. ACM, 2001. doi:10.1145/378795.378798.
- 28 Pascal Hitzler. A review of the semantic web field. Commun. ACM, 64(2):76-83, 2021. doi:10.1145/3397512.
- 29 Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d'Amato, Gerard de Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan F. Sequeda, Steffen Staab, and Antoine Zimmermann. Knowledge graphs. ACM Comput. Surv., 54(4):71:1–71:37, 2022. doi:10.1145/3447772.
- 30 Ian Horrocks, Oliver Kutz, and Ulrike Sattler. The even more irresistible SROIQ. In Patrick Doherty, John Mylopoulos, and Christopher A. Welty, editors, Proc. 10th International

licensed under Creative Commons License CC-BY 4.0

39th European Conference on Object-Oriented Programming (ECOOP 2025)

<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen;

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:28–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Conference on Principles of Knowledge Representation and Reasoning (KR 2006). AAAI Press, 2006. URL: http://www.aaai.org/Library/KR/2006/kr06-009.php.

- 31 Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Trans. Program. Lang. Syst.*, 23(3):396–450, 2001. doi:10.1145/503502.503505.
- 32 Félix Ingrand and Malik Ghallab. Deliberation for autonomous robots: A survey. Artif. Intell., 247:10-44, 2017. doi:10.1016/j.artint.2014.11.003.
- 33 Einar Broch Johnsen, Marcel Kyas, and Ingrid Chieh Yu. Dynamic classes: Modular asynchronous evolution of distributed concurrent objects. In Ana Cavalcanti and Dennis Dams, editors, Proc. Second World Congress on Formal Methods (FM 2009), volume 5850 of Lecture Notes in Computer Science, pages 596–611. Springer, 2009. doi:10.1007/978-3-642-05089-3\\_38.
- Einar Broch Johnsen and Ingrid Chieh Yu. Dynamic software updates and context adaptation for distributed active objects. In *Principled Software Development*, pages 147–164. Springer, 2018. doi:10.1007/978-3-319-98047-8\\_10.
- 35 Eduard Kamburjan, Nelly Bencomo, Silvia Lizeth Tapia Tarifa, and Einar Broch Johnsen. Declarative lifecycle management in digital twins. In Proc. 1st International Conference on Engineering Digital Twins (EDTconf 2024), MODELS Companion'24, pages 353—363. ACM, 2024. doi:10.1145/3652620.3688248.
- 36 Eduard Kamburjan and Einar Broch Johnsen. Knowledge structures over simulation units. In Annual Modeling and Simulation Conference (ANNSIM 2022), pages 78–89. IEEE, 2022. doi:10.23919/ANNSIM55834.2022.9859490.
- 37 Eduard Kamburjan, Vidar Norstein Klungre, Rudolf Schlatte, David Cameron, S. Lizeth Tapia Tarifa, and Einar Broch Johnsen. Digital twin reconfiguration using asset models. In Proc. 11th International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (ISoLA 2022), volume 13704 of Lecture Notes in Computer Science, pages 71–88. Springer, 2022. doi:10.1007/978-3-031-19762-8\_6.
- 38 Eduard Kamburjan, Vidar Norstein Klungre, Rudolf Schlatte, Einar Broch Johnsen, and Martin Giese. Programming and debugging with semantically lifted states. In Proc. Extended Semantic Web Conference (ESWC 2021), volume 12731 of Lecture Notes in Computer Science, pages 126–142. Springer, 2021. doi:10.1007/978-3-030-77385-4\\_8.
- 39 Eduard Kamburjan and Egor V. Kostylev. Type checking semantically lifted programs via query containment under entailment regimes. In *Description Logics*, volume 2954 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2021. URL: https://ceur-ws.org/Vol-2954/paper-19.pdf.
- 40 Eduard Kamburjan, Andrea Pferscher, Rudolf Schlatte, Riccardo Sieve, Silvia Lizeth Tapia Tarifa, and Einar Broch Johnsen. Semantic reflection and digital twins: A comprehensive overview. In Mike Hinchey and Bernhard Steffen, editors, *The Combined Power of Research, Education, and Dissemination: Essays Dedicated to Tiziana Margaria on the Occasion of Her 60th Birthday*, volume 15240 of *Lecture Notes in Computer Science*, pages 129–145. Springer, 2025. doi:10.1007/978-3-031-73887-6\_11.
- 41 Eduard Kamburjan, Riccardo Sieve, Chinmayi Prabhu Baramashetru, Marco Amato, Gianluca Barmina, Eduard Occhipinti, and Einar Broch Johnsen. GreenhouseDT: An exemplar for digital twins. In Proc. 19th Intl. Symp. on Software Engineering for Adaptive and Self-Managing Systems (SEAMS'24), page 175–181. ACM, 2024. doi:10.1145/3643915.3644108.
- 42 Erkan Karabulut, Salvatore F. Pileggi, Paul Groth, and Victoria Degeler. Ontologies in digital twins: A systematic literature review. *Future Gener. Comput. Syst.*, 153:442–456, 2024. doi:10.1016/j.future.2023.12.013.
- 43 Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In James F. Allen, Richard Fikes, and Erik Sandewall, editors, Proc. 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR 1991), pages 387–394. Morgan Kaufmann, 1991.
- 44 Jeffrey O. Kephart and David M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003. doi:10.1109/MC.2003.1160055.

<sup>©</sup> Riccardo Sieve and Eduard Kamburjan and Ferruccio Damiani and Einar Broch Johnsen; licensed under Creative Commons License CC-BY 4.0

<sup>39</sup>th European Conference on Object-Oriented Programming (ECOOP 2025).

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:29–2:30

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

- 45 Hector J. Levesque, Raymond Reiter, Yves Lespérance, Fangzhen Lin, and Richard B. Scherl. GOLOG: A logic programming language for dynamic domains. J. Log. Program., 31(1-3):59-83, 1997. doi:10.1016/S0743-1066(96)00121-5.
- 46 Vladimir Lifschitz. What is answer set programming? In Dieter Fox and Carla P. Gomes, editors, Proc. 23rd Conference on Artificial Intelligence (AAAI 2008), pages 1594–1597. AAAI Press, 2008. URL: http://www.aaai.org/Library/AAAI/2008/aaai08-270.php.
- Erik Meijer, Brian Beckman, and Gavin M. Bierman. LINQ: reconciling object, relations 47 and XML in the .NET framework. In Surajit Chaudhuri, Vagelis Hristidis, and Neoklis Polyzotis, editors, Proc. International Conference on Management of Data, page 706. ACM, 2006. doi:10.1145/1142473.1142552.
- 48 Bertrand Meyer. Object-Oriented Software Construction. Prentice-Hall, 2 edition, 1997. URL: http://www.eiffel.com/doc/oosc/page.html.
- 49 Alessandro Orso, Anup Rao, and Mary Jean Harrold. A technique for dynamic updating of Java software. In Proc. 18th International Conference on Software Maintenance (ICSM 2002), pages 649-658. IEEE Computer Society, 2002. doi:10.1109/ICSM.2002.1167829.
- 50 Benjamin C. Pierce. Types and programming languages. MIT Press, 2002.
- 51 Luís Pina, Luís Veiga, and Michael W. Hicks. Rubah: DSU for Java on a stock JVM. In Andrew P. Black and Todd D. Millstein, editors, Proc. International Conference on Object Oriented Programming Systems Languages & Applications (OOPSLA 2014), pages 103–119. ACM, 2014. doi:10.1145/2660193.2660220.
- Susanne Cech Previtali and Thomas R. Gross. Dynamic updating of software systems based 52 on aspects. In Proc. 22nd International Conference on Software Maintenance (ICSM 2006), pages 83-92. IEEE Computer Society, 2006. doi:10.1109/ICSM.2006.23.
- 53 Yuanwei Qu, Eduard Kamburjan, Anita Torabi, and Martin Giese. Semantically triggered qualitative simulation of a geological process. Applied Computing and Geosciences, 21:100152, 2024. doi:10.1016/j.acags.2023.100152.
- Tim Rentsch. Object oriented programming. ACM SIGPLAN Notices, 17(9):51-57, 1982. 54 doi:10.1145/947955.947961.
- Habib Seifzadeh, Hassan Abolhassani, and Mohsen Sadighi Moshkenani. A survey of dynamic 55 software updating. J. Softw. Evol. Process., 25(5):535-568, 2013. doi:10.1002/smr.1556.
- 56 Bjarne Stroustrup. What is object-oriented programming? IEEE Softw., 5(3):10-20, 1988. doi:10.1109/52.2020.
- 57 Jean-Pierre Talpin and Pierre Jouvelot. The type and effect discipline. In Proc. 7th Symposium on Logic in Computer Science (LICS '92), pages 162–173. IEEE Computer Society, 1992. doi:10.1109/LICS.1992.185530.
- 58 Danny Weyns. An Introduction to Self-Adaptive Systems: A Contemporary Software Engineering Perspective. John Wiley & Sons, UK, 2020. doi:10.1002/9781119574910.
- 59 Guohui Xiao, Diego Calvanese, Roman Kontchakov, Domenico Lembo, Antonella Poggi, Riccardo Rosati, and Michael Zakharyaschev. Ontology-based data access: A survey. In Proc. Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI 2018), pages 5511-5519. ijcai.org, 2018. doi:10.24963/IJCAI.2018/777.
- 60 Jinzhi Lu Xiaochen Zheng and Dimitris Kiritsis. The emergence of cognitive digital twin: vision, challenges and opportunities. International Journal of Production Research, 60(24):7610-7632, 2022. doi:10.1080/00207543.2021.2014591.

© Riccardo Sieve and Eduard Kamburian and Ferruccio Damiani and Einar Broch Johnsen: licensed under Creative Commons License CC-BY 4.0 39th European Conference on Object-Oriented Programming (ECOOP 2025)

Editors: Jonathan Aldrich and Alexandra Silva; Article No. 2; pp. 2:30-2:30

Leibniz International Proceedings in Informatics

