

Article

Fuzzy Ontology Reasoning over Sustainable Development Data

Fernando Bobillo ^{1,2}, Giuseppe Filippone ^{3,*}, Gianmarco La Rosa ³, Umberto Straccia ⁴
and Marco Elio Tabacchi ^{3,5}

- ¹ Department of Computer Science and Systems Engineering, University of Zaragoza, C. María de Luna, 1, 50018 Zaragoza, Spain; fbobillo@unizar.es
 - ² Aragon Institute of Engineering Research (I3A), C. Mariano Esquillor Gómez, s/n, 50018 Zaragoza, Spain
 - ³ Dipartimento di Matematica e Informatica, Università degli Studi di Palermo, Via Archirafi 34, 90123 Palermo, Italy; gianmarco.larosa@unipa.it (G.L.R.); marcoelio.tabacchi@unipa.it (M.E.T.)
 - ⁴ Consiglio Nazionale delle Ricerche-Istituto di Scienza e Tecnologie dell'Informazione "Alessandro Faedo" (CNR-ISTI), Via Giuseppe Moruzzi, 1, 56127 Pisa, Italy; umberto.straccia@isti.cnr.it
 - ⁵ Istituto Nazionale di Ricerche Demopolis, 90123 Palermo, Italy
- * Correspondence: giuseppe.filippone01@unipa.it

Abstract

Sustainable-development indicators are inherently graded: a territory is not simply poor or resilient, but exhibits these characteristics to varying degrees. Consequently, policy questions are naturally expressed through vague concepts such as *high unemployment*, *stressed neighbourhood*, or *green-transition readiness*. This work shows how such questions can be answered through reasoning over a fuzzy Web Ontology Language (OWL 2) ontology, rather than by ad hoc data-processing scripts. Our case study is the ontology of the Sustainability Decision Framework (SDF) Project, a five-module fuzzy ontology integrating Sustainable Development Goal (SDG) statistics for the United Nations Economic Commission for Europe (UNECE) region with geopolitical knowledge and a fuzzy semantic layer comprising 35 linguistic datatypes and 21 composite concepts defined through weighted sums and ordered weighted averaging (OWA). The ontology contains almost 400,000 statistical observations, and its complete export into a fuzzyDL Knowledge Base comprises 2.96 million axioms; for this reason, reasoning is performed on the slices of the ontology that are relevant to each query. Reasoning is performed using fuzzy-dl-owl2, our Python re-engineering of the fuzzyDL reasoner and the Fuzzy OWL 2 framework. We present a fully reproducible reasoning pipeline, including slice extraction, OWL 2-to-fuzzyDL translation, and mixed-integer linear programming (MILP) inference using Gurobi. We demonstrate the approach through five applications covering fuzzy territorial profiling, cross-border stress analysis, reasoning over economic groupings, analyst-defined concepts with linguistic hedges and coherence checks of the vocabulary, and temporal trend analysis scenarios. All experiments are conducted on real-world data, and the source code, configuration files, and complete fuzzyDL export of the ontology are publicly available. The results show that fuzzy ontological reasoning provides an expressive, declarative, and reproducible framework for analysing sustainable-development indicators while preserving the semantics of linguistic concepts.

Keywords: fuzzy ontologies; Fuzzy OWL 2; fuzzy description logics; fuzzyDL; sustainable development goals; geopolitical knowledge graphs; ordered weighted averaging; Python



Academic Editor: Ferdinando Di Martino

Received: 11 September 2026

Revised: 6 October 2026

Accepted: 6 October 2026

Published: 9 October 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Ontologies have become the standard means of representing and sharing structured knowledge on the Web. In particular, the World Wide Web Consortium (W3C) Web Ontology Language OWL 2 [1,2], grounded in Description Logics (DLs) [3], provides a well-established formalism together with sound and complete reasoning services. However, classical DLs are inherently crisp and therefore cannot adequately represent the vagueness that characterises many real-world domains [4]. For example, the statement that Greece had high unemployment in 2021 is neither simply true nor false; rather, its validity depends both on the observed unemployment rate (19.6%) and on how the linguistic concept *high* is interpreted.

Fuzzy set theory [5] and fuzzy Description Logics [6–9] address this limitation by interpreting concepts as fuzzy sets and by extending the classical reasoning services to graded semantics. Building on these foundations, *Fuzzy OWL 2* [10] represents fuzzy constructs through OWL 2 annotation axioms. Consequently, a fuzzy ontology remains a valid OWL 2 document that can be processed by standard Semantic Web tools, while specialised fuzzy reasoners, such as fuzzyDL [11], interpret the annotations and perform fuzzy reasoning by computing the corresponding degrees of truth.

Sustainable development is a domain in which such graded reasoning arises naturally. Although the indicators of the 17 Sustainable Development Goals (SDGs) are numerical, the questions posed about them are inherently linguistic and multi-dimensional: *Which territories are under socio-economic stress? Is a country surrounded by stressed neighbours? Do euro-area economies differ from other countries in terms of green-transition readiness?* Answering such questions requires (i) a semantic model specifying the meaning of linguistic terms for each indicator, (ii) an aggregation policy capturing how the different dimensions should be combined (e.g., compensatory versus non-compensatory, or all versus at least one), and (iii) a geopolitical representation describing countries, borders, and economic groupings over which the reasoning is performed. A fuzzy ontology can naturally represent all three aspects, while a fuzzy DL reasoner can answer such queries with a precise and well-defined semantics.

In a companion work [12], we introduced fuzzy-dl-owl2 [13,14], a complete Python re-engineering of the fuzzyDL reasoner and the Fuzzy OWL 2 framework, together with the accompanying pyowl2 library [15], and compared its performance with that of the original Java implementation. The present work applies this software to fuzzy ontological reasoning over a large sustainable-development ontology by operating on query-relevant ontology slices.

More specifically, we briefly present the ontology developed within the Sustainability Decision Framework (SDF) project, a Fuzzy OWL 2 ontology describing the United Nations Economic Commission for Europe (UNECE) SDG statistics. The ontology is organised into five modules (schema, observations, territorial snapshots, geopolitical subjects, and documentation), and its structural metrics are reported after excluding annotation axioms (see Section 4.1). We then describe a fully reproducible reasoning pipeline comprising slice extraction, translation into a fuzzyDL Knowledge Base, and mixed-integer linear programming (MILP) inference using Gurobi. The complete implementation, configuration files, and execution scripts are publicly available (see Section 4.3).

We further present five real-world applications demonstrating how fuzzy ontological reasoning can answer geopolitical queries over sustainable-development data. These applications illustrate the semantics of weighted sums, ordered weighted averaging (OWA), existential restrictions over inverse and transitive roles, linguistic modifiers, graded subsumption and satisfiability (see Sections 5.1–5.5). Finally, we assess the sensitivity of the reasoning results to the ontology breakpoints and aggregation weights (Section 5.6), anal-

use the computational cost of the proposed reasoning pipeline, and discuss its current limitations (Section 5.7).

The complete fuzzyDL translation of the ontology, comprising 2,964,012 axioms (approximately 236 MB), was generated using the same tool chain. Owing to its size, the complete export is not included in the manuscript, but is distributed together with the source code as a reusable artefact. It is not used directly for reasoning in this work, since all reported results are computed on query-relevant ontology slices (Section 4.3).

This work is organised as follows. Section 2 reviews the related work. Section 3 recalls the necessary background on fuzzy Description Logics. Section 4 presents the ontology of the SDF project, the reasoning pipeline, and its implementation. Section 5 presents the five applications and discusses the computational cost and current limitations of the proposed approach. Finally, Section 6 summarises the main findings and concludes the paper.

2. Related Works

Fuzzy Description Logics (DLs) extend classical DLs by introducing truth degrees in $[0, 1]$, fuzzy concept constructors, and graded axioms [7]. Depending on the adopted fuzzy semantics, reasoning is typically reduced either to mixed-integer linear programming (MILP), as in fuzzyDL [11], or to classical DL reasoning, as in DeLorean [16]. Alternative approaches based on tableau calculi and language extensions have also been proposed [17–20]. Decidability and complexity results under different fuzzy semantics are surveyed in [8,9], while the $\mathcal{EL}$ fragment has been studied in [21] separately. The fuzzyDL reasoner, which constitutes the reasoning engine adopted in this work, supports fuzzy *SHIF(D)* with concrete features, fuzzy datatypes (triangular, trapezoidal, and shoulder functions), linguistic modifiers, weighted concepts, weighted sums, OWAs [22,23], and fuzzy integrals.

Fuzzy OWL 2 [10] extends OWL 2 by encoding fuzzy semantics through a dedicated annotation property, `fuzzyLabel`, while preserving syntactic compatibility with standard OWL 2 tools. The accompanying Protégé plug-in [24,25] and the `FuzzyOwl2ToFuzzyDL` translator provide practical support for editing and reasoning over fuzzy ontologies. This framework has been employed to construct fuzzy ontologies from consensuated linguistic datatypes [26] and to learn fuzzy concept inclusions from OWL data [27]. The ontology developed within the SDF project follows the Fuzzy OWL 2 specification directly: fuzzy datatypes, linguistic modifiers, and composite concepts and graded class assertions are represented as OWL 2 entities annotated with `fuzzyLabel`.

The software infrastructure adopted in this work is based on our recent Python re-engineering of fuzzyDL and the Fuzzy OWL 2 framework [12–14,28]. The re-engineering was motivated by the need for a portable implementation supporting multiple MILP backends, including Gurobi [29,30] as well as the open-source solvers HiGHS [31], CBC [32], GLPK [33], and interfaces provided through `mip` [34] and `PuLP` [35]. The parser is generated using `re2c` [36] and `flex` [37], while the computationally intensive components are compiled with Cython [38]. The experiments in this work were carried out with release 1.0.30 of the library.

Knowledge-based approaches to SDG monitoring range from ontology-based representations of indicator metadata to decision-support systems [39,40]. To the best of our knowledge, however, the ontology of the SDF project is the first fuzzy OWL 2 ontology combining the complete SDG statistical layer of an entire region with a geopolitical knowledge layer (including borders, economic groupings, and territorial relationships) and a fuzzy semantic layer designed for automated DL reasoning. More broadly, fuzzy ontologies have been successfully applied to information retrieval [41–43] and Semantic Web

applications [44–46]; the applications presented in this work belong to the latter category, with a particular emphasis on decision support for sustainable development.

3. Background

This section recalls the notions of fuzzy ontologies and fuzzy Description Logics used throughout this work. We refer the reader to [3,47] for classical Description Logics and to [7,48] for comprehensive treatments of fuzzy Description Logics.

Definition 1 (Fuzzy set [5]). A fuzzy set A over a universe X is characterised by a membership function $\mu_A: X \rightarrow [0, 1]$, where $\mu_A(x)$ is the degree to which x belongs to A .

Hence, classical (crisp) sets are the special case $\mu_A: X \rightarrow \{0, 1\}$. Before extending ontologies with fuzzy sets, we recall the classical notion of ontology in Description Logic terms [1,47].

Definition 2 (Ontology). Let N_C , N_R , N_F , and N_I be pairwise disjoint sets of concept names, role names, concrete feature names, and individual names. Let Δ_D be a non-empty concrete domain of data values. An ontology is a Knowledge Base $\mathcal{K} = \langle \mathcal{T}, \mathcal{R}, \mathcal{A} \rangle$, where

- The TBox $\mathcal{T}$ is a finite set of concept inclusions $C \sqsubseteq D$ and concept equivalences $A \equiv C$, where $A \in N_C$ and C, D are concepts constructed from concept names, role names, and concrete feature names using the constructors of the underlying Description Logic;
- The RBox $\mathcal{R}$ is a finite set of role axioms, including role inclusions $R \sqsubseteq S$, inverse-role axioms $R \equiv S^{-}$, and declarations $\text{trans}(R)$, $\text{sym}(R)$, and $\text{func}(R)$, with $R, S \in N_R$;
- The ABox $\mathcal{A}$ is a finite set of assertions of the form $a : C$, $(a, b) : R$, and $(a, v) : f$, where $a, b \in N_I$, $R \in N_R$, $f \in N_F$, and $v \in \Delta_D$.

An interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consists of a non-empty abstract domain $\Delta^{\mathcal{I}}$ and an interpretation function $(\cdot)^{\mathcal{I}}$ such that each individual name $a \in N_I$ is interpreted as an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, each concept name $A \in N_C$ is interpreted as a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, each role name $R \in N_R$ is interpreted as a binary relation $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and each concrete feature $f \in N_F$ is interpreted as a partial function $f^{\mathcal{I}}: \Delta^{\mathcal{I}} \rightharpoonup \Delta_D$. The interpretation function is extended inductively to complex concepts according to the semantics of the constructors of the underlying Description Logic. An interpretation $\mathcal{I}$ is a model of $\mathcal{K}$, written $\mathcal{I} \models \mathcal{K}$, if it satisfies all axioms in $\mathcal{T}$, $\mathcal{R}$, and $\mathcal{A}$. In particular, $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for every $C \sqsubseteq D \in \mathcal{T}$, and $A^{\mathcal{I}} = C^{\mathcal{I}}$ for every $A \equiv C \in \mathcal{T}$. The role axioms are interpreted as follows: $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$ for every $R \sqsubseteq S \in \mathcal{R}$, and $R^{\mathcal{I}} = (S^{\mathcal{I}})^{-1}$ for every $R \equiv S^{-} \in \mathcal{R}$, while $\text{trans}(R)$, $\text{sym}(R)$, and $\text{func}(R)$ require $R^{\mathcal{I}}$ to be, respectively, transitive, symmetric, and functional. Finally, the ABox assertions are satisfied whenever $a^{\mathcal{I}} \in C^{\mathcal{I}}$ for every $a : C \in \mathcal{A}$, $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$ for every $(a, b) : R \in \mathcal{A}$, and $f^{\mathcal{I}}(a^{\mathcal{I}}) = v$ for every $(a, v) : f \in \mathcal{A}$.

In a classical interpretation, an individual either belongs to a concept or does not. Fuzzy Description Logics relax this assumption by interpreting concepts as fuzzy sets and roles as fuzzy binary relations.

Definition 3 (Fuzzy interpretation). A fuzzy interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consists of a non-empty domain $\Delta^{\mathcal{I}}$ and an interpretation function that maps each individual name $a \in N_I$ to an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, each concept name $A \in N_C$ to a membership function $A^{\mathcal{I}}: \Delta^{\mathcal{I}} \rightarrow [0, 1]$, each role name $R \in N_R$ to a fuzzy binary relation $R^{\mathcal{I}}: \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \rightarrow [0, 1]$, and each concrete feature $f \in N_F$ to a partial function $f^{\mathcal{I}}: \Delta^{\mathcal{I}} \rightharpoonup \Delta_D$. The interpretation function is extended to complex concepts according to the semantics of the constructors given below.

Just as concept membership and role relations can be graded, axioms and assertions may also hold to a degree, giving rise to the notion of a fuzzy ontology.

Definition 4 (Fuzzy ontology). A fuzzy ontology is a Knowledge Base $\mathcal{K} = \langle \mathcal{T}, \mathcal{R}, \mathcal{A} \rangle$ as in 2, except that concept and role assertions are graded, i.e., of the form $\langle a : C \geq n \rangle$ and $\langle (a, b) : R \geq n \rangle$, and the TBox may contain graded concept inclusions $\langle C \sqsubseteq D \geq n \rangle$, with $n \in [0, 1]$. In particular, every axiom written without a degree is assumed to have a degree 1. A fuzzy interpretation $\mathcal{I}$ (3) satisfies these axioms according to

$$\begin{aligned} \mathcal{I} \models \langle a : C \geq n \rangle &\iff C^{\mathcal{I}}(a^{\mathcal{I}}) \geq n, \\ \mathcal{I} \models \langle (a, b) : R \geq n \rangle &\iff R^{\mathcal{I}}(a^{\mathcal{I}}, b^{\mathcal{I}}) \geq n, \\ \mathcal{I} \models \langle C \sqsubseteq D \geq n \rangle &\iff \inf_{x \in \Delta^{\mathcal{I}}} (C^{\mathcal{I}}(x) \Rightarrow D^{\mathcal{I}}(x)) \geq n, \\ \mathcal{I} \models A \equiv C &\iff A^{\mathcal{I}}(x) = C^{\mathcal{I}}(x) \text{ for all } x \in \Delta^{\mathcal{I}}, \end{aligned} \quad (1)$$

where $(\Rightarrow)$ is the implication of the underlying fuzzy logic. Feature assertions are satisfied if $f^{\mathcal{I}}(a^{\mathcal{I}}) = v$. Role axioms are satisfied if $R^{\mathcal{I}}(x, y) \leq S^{\mathcal{I}}(x, y)$ for $R \sqsubseteq S$, $R^{\mathcal{I}}(x, y) = S^{\mathcal{I}}(y, x)$ for $R \equiv S^{-}$, $R^{\mathcal{I}}(x, y) \otimes R^{\mathcal{I}}(y, z) \leq R^{\mathcal{I}}(x, z)$ for $\text{trans}(R)$, and $R^{\mathcal{I}}(x, y) = R^{\mathcal{I}}(y, x)$ for $\text{sym}(R)$, for all $x, y, z \in \Delta^{\mathcal{I}}$, whereas $\text{func}(R)$ requires every x to have at most one y with $R^{\mathcal{I}}(x, y) > 0$ [11]. Finally, $\mathcal{I}$ is a model of $\mathcal{K}$ if it satisfies all the axioms of $\mathcal{K}$, and $\mathcal{K}$ entails an axiom τ , written $\mathcal{K} \models \tau$, if every model of $\mathcal{K}$ satisfies τ [7,10].

Hence, an ontology (2) is the special case in which all degrees are 1 and interpretations take values in $\{0, 1\}$.

For instance, a territorial snapshot may belong to a concept `EconomicStress` with degree 0.35, rather than simply being classified as stressed or not. The remainder of this section introduces the concept constructors and the fuzzy constructs used in this work.

Complex concepts are obtained from concept names by means of the usual Description Logic constructors: the universal and empty concepts $\top$ and $\perp$, negation $\neg C$, conjunction $C \sqcap D$, disjunction $C \sqcup D$, and the existential and universal restrictions $\exists R.C$ and $\forall R.C$. The latter two describe the individuals related through the role R to, respectively, at least one individual belonging to C or exclusively to individuals belonging to C . Roles may also be declared symmetric, transitive, functional, or inverses of other roles. Their semantics are determined by a fuzzy logic, namely a t -norm $\otimes$, its dual t -conorm $\oplus$, and the corresponding negation and implication operators. In particular,

$$(C \sqcap D)^{\mathcal{I}}(x) = C^{\mathcal{I}}(x) \otimes D^{\mathcal{I}}(x), \quad (\exists R.C)^{\mathcal{I}}(x) = \sup_{y \in \Delta^{\mathcal{I}}} (R^{\mathcal{I}}(x, y) \otimes C^{\mathcal{I}}(y)). \quad (2)$$

Interpretations are required to be *witnessed*, that is, the infimum and supremum in Equations (1) and (2) are attained at some point of the interpretation domain [7].

Numerical data, such as the SDG indicator values of the territorial snapshots, are handled through *fuzzy concrete domains*, which associate linguistic terms with ranges of numerical values.

Definition 5 (Concrete feature and fuzzy datatype). A concrete feature $f \in \mathbf{N}_{\mathbf{F}}$ (e.g., `unemploymentRate`) is interpreted as a partial function $f^{\mathcal{I}}: \Delta^{\mathcal{I}} \rightarrow [k_1, k_2]$ assigning at most one numerical value to each element of the domain, and a fuzzy datatype $\mathbf{d}$ over $[k_1, k_2]$ is a fuzzy set with membership function $\mathbf{d}: [k_1, k_2] \rightarrow [0, 1]$.

This work employs the left-shoulder, right-shoulder, and trapezoidal membership functions, defined for parameters $a < b \leq c < d$ as

$$ls(a, b)(x) = \begin{cases} 1 & x \leq a, \\ \frac{b-x}{b-a} & x \in (a, b], \\ 0 & x > b, \end{cases} \quad rs(a, b)(x) = \begin{cases} 0 & x \leq a, \\ \frac{x-a}{b-a} & x \in (a, b], \\ 1 & x > b, \end{cases} \quad (3)$$

and

$$trz(a, b, c, d)(x) = \begin{cases} 0 & x \leq a \text{ or } x > d, \\ \frac{x-a}{b-a} & x \in (a, b], \\ 1 & x \in (b, c], \\ \frac{d-x}{d-c} & x \in (c, d]. \end{cases} \quad (4)$$

A *datatype restriction* $\exists f.\mathbf{d}$ is then a concept whose degree, for a concrete feature f , is given by the membership of the feature value in the datatype,

$$(\exists f.\mathbf{d})^{\mathcal{I}}(x) = \begin{cases} \mathbf{d}(f^{\mathcal{I}}(x)) & \text{if } f^{\mathcal{I}}(x) \text{ is defined,} \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

whereas crisp comparisons such as $\exists f.\geq_m$, $\exists f.\leq_m$, and $\exists f.=_m$ are special cases whose membership functions take only the values 0 and 1.

A *fuzzy modifier* m transforms the degree of a concept according to

$$(m(C))^{\mathcal{I}}(x) = f_m(C^{\mathcal{I}}(x)), \quad (6)$$

and is used to model linguistic hedges such as *very*. This work employs *linear modifiers*, parameterised by a single value $c > 0$. The corresponding function $f_{lm(c)}$ is the piecewise linear function passing through the points $(0, 0)$, $(\frac{c}{c+1}, \frac{1}{c+1})$, and $(1, 1)$, i.e.,

$$f_{lm(c)}(x) = \begin{cases} \frac{x}{c} & x \leq \frac{c}{c+1}, \\ cx - c + 1 & x > \frac{c}{c+1}, \end{cases} \quad (7)$$

which sharpens concepts for $c > 1$ (preserving 0 and 1 while reducing intermediate degrees) and dilates them for $c < 1$. Modifiers can also be applied to fuzzy datatypes: $(\exists f.m(\mathbf{d}))^{\mathcal{I}}(x) = f_m(\mathbf{d}(f^{\mathcal{I}}(x)))$.

Finally, *fuzzyDL* [11] extends the language with aggregation concepts [22]. A *weighted sum*

$$w_1 \cdot C_1 + \dots + w_k \cdot C_k, \quad (8)$$

with $w_i \in [0, 1]$ and $\sum_i w_i \leq 1$, computes a weighted combination $\sum_i w_i C_i^{\mathcal{I}}(x)$ of the component membership degrees, whereas the *ordered weighted averaging* (OWA) operator

$$\text{OWA}_W(C_1, \dots, C_k) \quad (9)$$

applies the weight vector $W = (w_1, \dots, w_k)$, with $w_i \in [0, 1]$ and $\sum_i w_i = 1$, to the component membership degrees sorted in decreasing order. All these constructs are available in OWL 2 ontologies through the Fuzzy OWL 2 annotation scheme [10], in which fuzzy datatypes, modifiers, aggregation concepts, and graded assertions are encoded as standard OWL 2 annotations. Consequently, a fuzzy ontology remains a valid OWL 2 document that

can be processed by conventional Semantic Web tools while preserving its fuzzy semantics. This annotation scheme is adopted throughout this work (see Section 4.1).

4. Methods

In this section, we present the knowledge representation and reasoning infrastructure developed to support the applications considered in this work. We first describe the modular fuzzy ontology underlying the SDF project, including its main knowledge components and the fuzzy constructs used to represent graded sustainability-related concepts. Next, we justify some design choices regarding the type of fuzzy sets and the parameters of membership functions and aggregation operators. We then introduce a reasoning pipeline that extracts query-specific ontology slices, translates them into fuzzy Description Logic knowledge bases, and executes the required reasoning tasks. Finally, we describe the supporting software infrastructure, with particular emphasis on ontology indexing, slice extraction, query generation, and the reproducibility of the experimental workflow.

4.1. The Fuzzy Ontology of the SDF Project

The ontology developed within the SDF project (a co-funded research project on decision support for sustainability; see the funding statement) is a modular OWL 2 ontology whose import closure comprises five modules (see Figure 1 and Table 1). The *schema* module defines the conceptual model, including SDG goals, targets, indicators and series, observations, subjects (countries, territories, and organisations), territorial systems, phenomena, policies, actors, and the fuzzy layer. The schema module also contains 106 *SubClassOf*, 64 *EquivalentClasses*, and 44 *DisjointClasses* axioms. The individual count reported for the import closure corresponds to the number of distinct individuals. For example, the 56 UNECE countries are declared in the *observations* module and subsequently enriched with additional assertions in the *subjects* module, but they are counted only once. The *observations* module contains the complete UNECE SDG database, comprising the statistics for the 17 SDGs, 253 indicator tables, and 396,592 observations, of which 239,758 correspond to missing values represented as individuals without an associated numerical value.

The *territorial snapshots* module bridges the statistical and fuzzy layers. Each country–year snapshot (an instance of *TerritorialSystem*) stores the values of six selected SDG indicators as concrete features, i.e., *povertyRate* (SDG Indicator 1.1.1), *unemploymentRate* (SDG Indicator 8.5.2), *economicGrowthRate* (SDG Indicator 8.1.1), *foodInsecurityRate* (SDG Indicator 2.1.2), *digitalAccessRate* (SDG Indicator 17.8.1), and *cleanEnergyAccessRate* (SDG Indicator 7.1.2) maintains provenance links to the original observations, and materialises the corresponding membership degrees in the fuzzy concepts. For SDG 8.5.2, the observations module provides unemployment rates separately for persons with and without disabilities. The concrete feature *unemploymentRate* uses the value for persons without disabilities, which does not correspond to the headline unemployment rate and may differ from it. Overall, this module contains 57,656 graded class assertions, computed off-line using *fuzzy-dl-owl2* and stored as lower bounds rounded down to four decimal places.

The *subjects* module provides the geopolitical layer. It contains 887 subjects, including 212 countries (56 of which belong to the UNECE region), former countries, territories, regions, and organisations, together with 216 cities, 145 currencies, 15,714 membership assertions relating subjects to 750 cluster nodes distributed over 31 cluster types (e.g., geographical regions, income groups, development status, partnerships, currency areas, languages, and political systems). The module also defines 29 economic-profile nodes (instances of *Economy*) connected to subjects through *hasEconomy*, and 516 *Distance* individuals encoding the borders between neighbouring countries, together with their cor-

responding border lengths. Finally, the *annotations* module documents the ontology by providing metadata such as SDG rationales, related indicators, and source references, while importing the remaining four modules.

Table 1. Structural metrics of the modules comprising the import closure of the SDF ontology (OWL API style, excluding annotation assertions and annotation properties). LA = logical axioms; Decl = declarations; C = classes; OP/DP = object/data properties; I = individuals; DT = datatypes; CA/OPA/DPA = class, object-property, and data-property assertions. The schema also contains 106 SubClassOf, 64 EquivalentClasses and 44 DisjointClasses axioms. The closure count of individuals is the number of distinct individuals (the 56 UNECE countries are declared in the observations module and enriched in the subjects module).

Module	LA	Decl	C	OP	DP	I	DT	CA	OPA	DPA
schema (SDF_ext)	460	318	181	78	25	0	32	–	–	–
observations (SDF_individuals)	2,778,183	397,180	0	2	3	397,174	1	397,174	1,587,386	793,612
snapshots (SDF_territories)	84,873	3237	0	1	1	3235	0	60,891	11,989	11,989
subjects (SDF_subjects)	100,412	18,291	0	0	4	18,284	0	18,257	57,825	24,323
annotations (SDF_fuzzy_annotated)	0	8	0	0	0	0	0	–	–	–
closure	2,963,928	419,034	181	81	33	418,637	33	476,322	1,657,200	829,924




Figure 1. The five modules comprising the import closure of the SDF ontology and their owl:imports relationships (arrows indicate the imported module). The *schema* module is the only module containing terminological and role axioms (TBox and RBox), together with the fuzzy layer. The remaining three data modules contribute individuals and assertion axioms (ABox), whereas the *annotations* module provides documentation for the ontology.

The fuzzy layer is encoded according to the Fuzzy OWL 2 specification. It comprises 32 *fuzzy datatypes*, which define the semantics of the linguistic terms associated with the numerical features by means of left-shoulder, right-shoulder, and trapezoidal membership functions over xsd:decimal. Listing 1 reports the corresponding fuzzyDL translation for some representative datatypes.

On top of these datatypes, the ontology defines 37 *bridge classes*, each relating a linguistic term to a concrete feature through an existential datatype restriction. For example, *TerritoryWithHighUnemployment* represents the class of *TerritorialSystem* individuals whose *unemploymentRate* belongs to the fuzzy datatype *HighUnemployment*.

Crisp concepts can then be defined by combining these bridge classes; for instance, *SevereBasicNeedsDeprivation* represents territories simultaneously characterised by high food insecurity and high poverty. Both kinds of definitions are illustrated in Listing 1, while Figure 2 depicts the membership functions associated with two representative features.

The ontology further contains 21 *composite concepts*, which aggregate bridge classes according to three different policies (see Listing 2). The first is a compensatory weighted sum (*w-sum*); the second is a non-compensatory weighted sum (*w-sum-zero*), which evaluates to zero whenever at least one component is zero and is therefore used to model zero-tolerance concepts; the third is an OWA operator with weights (0.5, 0.25, 0.125, 0.125), which favours the best-satisfied dimensions and is used to define the broad concepts.

Listing 1. Fuzzy datatypes, a bridge class, and a crisp severe concept translated into fuzzyDL by fuzzy-dl-owl2 (excerpt). For each fuzzy datatype, the first two parameters specify the domain of the corresponding feature, whereas the remaining parameters define the associated membership function.

```

1 (define-fuzzy-logic lukasiewicz)
2 (define-fuzzy-concept LowPoverty      left-shoulder(0.0, 100.0, 2.0, 10.0))
3 (define-fuzzy-concept HighPoverty     right-shoulder(0.0, 100.0, 20.0, 30.0))
4 (define-fuzzy-concept ModeratePoverty trapezoidal(0.0, 100.0, 2.0, 10.0, 20.0, 30.0))
5 (define-fuzzy-concept HighUnemployment right-shoulder(0.0, 100.0, 12.0, 18.0))
6 (define-fuzzy-concept EconomicDecline left-shoulder(-40.0, 40.0, -5.0, 0.0))
7 (define-fuzzy-concept PositiveEconomicGrowth right-shoulder(-40.0, 40.0, 2.0, 5.0))
8 (define-fuzzy-concept StableEconomicGrowth trapezoidal(-40.0, 40.0, -5.0, 0.0, 2.0, 5.0))
9 (define-concept TerritoryWithHighUnemployment
10   (and TerritorialSystem (some unemploymentRate HighUnemployment)))
11 (define-concept SevereBasicNeedsDeprivation
12   (and TerritoryWithHighFoodInsecurity TerritoryWithHighPoverty))

```




Figure 2. Membership functions associated with the linguistic terms of two of the six features ((left): povertyRate, representing the proportion of the population living below the international poverty line; (right): unemploymentRate). For each feature, the linguistic terms are modelled by left-shoulder, trapezoidal, and right-shoulder membership functions with shared breakpoints, ensuring that the membership degrees of adjacent terms sum to one over their region of overlap.

4.2. Design Choices

In this section, we present and motivate the main modelling choices underlying the fuzzy ontology, including the adoption of type-1 fuzzy sets, the definition of membership functions, and the specification of aggregation weights. We also clarify the interpretation and adaptability of these parameters within the proposed framework.

Listing 2. Composite concepts translated into fuzzyDL by fuzzy-dl-owl2 (excerpt), illustrating compensatory, non-compensatory, and OWA aggregation.

```

1 (define-concept BasicNeedsStress
2   (w-sum (0.4 TerritoryWithHighFoodInsecurity) (0.35 TerritoryWithHighPoverty)
3     (0.25 TerritoryWithHighUnemployment)))
4 (define-concept EconomicStress
5   (w-sum (0.4 TerritoryInEconomicDecline) (0.35 TerritoryWithHighUnemployment)
6     (0.25 TerritoryWithHighPoverty)))
7 (define-concept SDGSynergyPotential
8   (w-sum-zero (0.25 TerritoryWithHighDigitalAccess) (0.25
9     TerritoryWithHighCleanEnergyAccess)
10     (0.25 TerritoryWithStrongInstitutionalCapacity
11   ) (0.25 TerritoryWithStrongResilience)))
12 (define-concept BroadSocioEconomicStress
13   (owa (0.5 0.25 0.125 0.125)
14     (TerritoryWithHighPoverty TerritoryWithHighInequality
15       TerritoryWithHighUnemployment TerritoryWithHighFoodInsecurity)))

```

The ontology relies on classical (type-1) fuzzy sets [5], rather than on extensions such as intuitionistic [49], interval-valued [50], interval-valued intuitionistic [51], or type-2 [52] fuzzy sets (see [53] for an overview of their relationships). This choice is motivated by three considerations. First, the SDG indicators are point-valued statistics with no associated bounds or standard errors, and therefore provide no independent evidence for a distinct non-membership degree. Setting the non-membership degree to $1 - \mu$ yields an intuitionistic fuzzy set with zero hesitation, which reduces to the type-1 case. Second, type-1 membership degrees represent the gradualness of linguistic terms, whereas interval-valued memberships represent uncertainty about the membership degree itself, that is, ill-known membership [54]. When a feature value is unknown or only partially specified, min-instance? and max-instance? already provide lower and upper bounds on the corresponding membership degree [55]. Third, Fuzzy OWL 2 and fuzzyDL represent each degree as a single value in $[0, 1]$ and, under Łukasiewicz logic, the constructors and membership functions used in this work are piecewise linear, which enables the reduction in reasoning to mixed-integer linear programming (MILP) [11]. Fuzzy Description Logics over lattices of truth degrees have been studied [56] but, to the best of our knowledge, no comparable OWL 2-compatible reasoning system is currently available for intuitionistic, interval-valued, or type-2 fuzzy Description Logics.

The membership functions and aggregation weights were defined by the domain experts of the SDF project, informed by the typical values of the corresponding SDG indicators rather than by a formal elicitation procedure. The breakpoints are anchored to the absolute scale of each indicator, allowing the same linguistic vocabulary to be applied across regions and years and thereby preserving the temporal comparability of the trajectories analysed in Section 5.5. For example, the breakpoints of HighUnemployment (12% and 18%) correspond approximately to the 78th and 90th percentiles of the UNECE values observed between 2015 and 2023. By contrast, terms calibrated on a global scale, such as HighPoverty, are never activated by the UNECE values observed over the same period. For poverty, unemployment, food insecurity, digital access, and clean-energy access, adjacent linguistic terms form strong fuzzy partitions [57], so that their membership degrees sum to one over their regions of overlap.

The weights of each composite concept sum to one, while dimensions for which data are not yet available retain their assigned weights, ensuring that the resulting degrees remain comparable when additional data sources become available. The OWA weight vector (0.5, 0.25, 0.125, 0.125) has an orness (see Remark 1) of 0.708, corresponding to an aggregation behaviour that emphasises the presence of at least one significant factor.

These parameters should therefore be regarded as modelling choices rather than empirical estimates. An analyst may replace the membership functions, breakpoints, or aggregation weights with alternative definitions at query time, provided that they are introduced under new concept names; otherwise, the membership degrees already materialised in the ontology act as lower bounds. Such alternatives can be introduced without modifying the underlying ontology, as illustrated in Section 5.4, while Section 5.6 quantifies the sensitivity of the reasoning results to these modelling choices.

Remark 1. *The attitude of an OWA operator is quantified by its orness [58], defined as*

$$\text{orness}(W) = \frac{1}{k-1} \sum_{i=1}^k (k-i) w_i \in [0, 1]. \quad (10)$$

The orness equals 1 for $W = (1, 0, \dots, 0)$, corresponding to the maximum of the component degrees, and equals 0 for $W = (0, \dots, 0, 1)$, corresponding to their minimum. For the arithmetic mean, with $W = (1/k, \dots, 1/k)$, the orness is 0.5. Values greater than 0.5 therefore indicate an increasingly disjunctive aggregation behaviour, in which the best-satisfied components exert greater influence on the aggregated degree.

4.3. Reasoning Pipeline

Reasoning over the entire SDF ontology, comprising 2.96 million logical axioms, is neither necessary nor practical for the types of queries considered in this work, since each query is solved as a single MILP over the entire ABox of the corresponding Knowledge Base (Section 5.7). In practice, each geopolitical query involves only a small subset of the ontology, typically consisting of a limited number of territorial snapshots together with their associated countries, observations, economic profiles, and neighbouring entities. For this reason, the proposed reasoning pipeline operates on *slices* of the ontology. Proposition 1 formalises the guarantees provided by answers computed on a slice with respect to those obtained from the complete ontology.

The pipeline consists of five steps. First, the individuals relevant to the query are selected (e.g., all territorial snapshots for the year 2021 containing the six SDG features, neighbouring country pairs, or other application-specific entities). Second, a temporary OWL 2 ontology is assembled by combining the complete TBox and RBox of the schema module with the ABox assertions of the selected individuals, the individuals they reference (such as countries, observations, and economic profiles), and the corresponding owl:Axiom annotations encoding graded assertions. Third, fuzzy-dl-owl2 translates the resulting slice into a fuzzyDL Knowledge Base. Fourth, any additional axioms defined by the analyst, written directly in fuzzyDL syntax, are appended to the Knowledge Base. Finally, the reasoning queries are added and solved using the configured MILP backend.

Figure 3 summarises the overall workflow. The second step relies on a lightweight SQLite index of the ontology import closure, automatically generated once by the pipeline, whereas the remaining steps are implemented directly through the fuzzy-dl-owl2 API. The configuration file used by the library is reported in Listing 3, while Listing 4 presents the two functions implementing the pipeline and reused by all application scripts. The complete implementation, together with the five application scripts, is distributed with this work.

Table 2 summarises the fragment of the ontology used throughout the applications. For each class, role, and concrete feature, it reports the corresponding Description Logic definition together with the textual description provided by the annotations module.




Figure 3. Overview of the reasoning pipeline. Grey boxes represent data artefacts and blue boxes represent processing steps.

Listing 3. CONFIG.ini: configuration file containing the settings read by fuzzy-dl-owl2 from the working directory.

```

1 [DEFAULT]
2 debugPrint = False
3 epsilon = 0.001
4 maxIndividuals = -1
5 owlAnnotationLabel = fuzzyLabel
6 milpProvider = gurobi

```

Listing 4. sdf_fuzzy.py (excerpt): translation of an ontology slice into a fuzzyDL Knowledge Base and execution of the corresponding reasoning queries.

```

1 import os
2 import shutil
3
4 def build_kb(individuals, workdir, axioms=""):
5     """Slice OWL file -> fuzzyDL KB (+ analyst axioms in fuzzyDL syntax)."""
6     shutil.copy(CODE / "CONFIG.ini", workdir / "CONFIG.ini") # read from the cwd
7     owl = workdir / "slice.owl"
8     stats = build_slice(individuals, owl)
9     # schema modules + individuals + degrees (sdf_index.py)
10    os.chdir(workdir)
11    from fuzzy_dl_owl2.fuzzyowl2.fuzzyowl2_to_fuzzydl import \
12        FuzzyOwl2ToFuzzyDL
13    FuzzyOwl2ToFuzzyDL(str(owl), "kb.fdl", base_iri=BASE).translate_owl2ontology()
14    kb = workdir / "results" / "kb.fdl"
15    kb.write_text(kb.read_text() + "\n# ---- analyst axioms ----\n" + axioms + "\n")
16    return kb, stats
17
18 def solve(kb, queries, workdir):
19     """Append the queries to the KB and answer them."""
20    kbq = workdir / "kb_query.fdl"
21    kbq.write_text(kb.read_text() + "\n# ---- queries ----\n" + "\n".join(queries) + "\n")
22
23    from fuzzy_dl_owl2.fuzzydl.parser import DLParseFast
24    kb_obj, qs = DLParseFast.get_kb(str(kbq))
25    kb_obj.solve_kb()
26    out = []
27    for q in qs:
28        sol = q.solve(kb_obj) # one MILP per query
29        out.append({"query": str(q),
30                    "value": sol.get_solution() if sol.is_consistent_kb() else None})
31    return out

```

Table 2. Excerpt of the ontology: the classes, roles, and concrete features used in the applications, together with their Description Logic definitions or characteristics and the corresponding textual descriptions provided by the annotations module. Weighted sums denote w -sum concepts, whereas $OWA_{\vec{w}}$ denotes the ordered weighted averaging operator.

Entity	DL Definition/Characteristics	Description (From the Annotations)
<i>Primitive classes</i>		
TerritorialSystem	(primitive)	Territorial units or spatial systems used as the primary objects of SDG-related reasoning; country–year snapshots are instances of this class.
CountrySubject	(primitive)	Country-level geographic subjects.
Distance	(primitive)	Distance entities modelling land-border relationships between geographic subjects, with one individual for each pair of neighbouring countries.
Economy	(primitive)	Represents economic entities and classifications associated with geographic subjects (e.g., the euro area and income groups).
<i>Distance classes</i>		
BorderingDistance	$\equiv \exists \text{hasDistanceValue}.\equiv_0$	Bordering or adjacent geographic entities.
ShortDistance	$\equiv \exists \text{hasDistanceValue}.\leq_{200}$	Short distances between geographic entities.
MidDistance	$\equiv \exists \text{hasDistanceValue}.\leq_{1000} \sqcap \exists \text{hasDistanceValue}.\geq_{200}$	Intermediate distances between geographic entities.
LongDistance	$\equiv \exists \text{hasDistanceValue}.\geq_{1000}$	Long distances between geographic entities.
<i>Bridge classes (37 in total, one per feature and linguistic term)</i>		
TerritoryWith⟨T⟩	$\equiv \text{TerritorialSystem} \sqcap \exists f.T$, e.g. $\text{TerritoryWithHighUnemployment} \equiv \text{TerritorialSystem} \sqcap \exists \text{unemploymentRate.HighUnemployment}$	Represents territorial systems characterised by the linguistic term T associated with the concrete feature f , as defined by the corresponding fuzzy datatype.
TerritoryInEconomicDecline	$\equiv \text{TerritorialSystem} \sqcap \exists \text{economicGrowthRate.EconomicDecline}$	Represents the class of territorial systems whose economic growth rate has a positive membership degree in the EconomicDecline fuzzy datatype.

Table 2. Cont.

Entity	DL Definition/Characteristics	Description (From the Annotations)
<i>Composite fuzzy concepts (21 in total; those used in the applications)</i>		
BasicNeedsStress	$\equiv 0.40 \text{ TerritoryWithHighFoodInsecurity} + 0.35 \text{ TerritoryWithHighPoverty} + 0.25 \text{ TerritoryWithHighUnemployment}$	Represents basic-needs stress arising from food insecurity, poverty, and unemployment.
EconomicStress	$\equiv 0.40 \text{ TerritoryInEconomicDecline} + 0.35 \text{ TerritoryWithHighUnemployment} + 0.25 \text{ TerritoryWithHighPoverty}$	Represents economic stress arising from economic decline, unemployment, and poverty.
StructuralDevelopmentStress	$\equiv 0.25 \text{ TerritoryWithHighPoverty} + 0.20 \text{ TerritoryWithHighUnemployment} + 0.20 \text{ TerritoryWithHighFoodInsecurity} + 0.20 \text{ TerritoryWithLowDigitalAccess} + 0.15 \text{ TerritoryWithLowCleanEnergyAccess}$	Accumulated structural development constraints.
BroadSocioEconomicStress	$\equiv \text{OWA}_{(0.5, 0.25, 0.125, 0.125)} (\text{TerritoryWithHighPoverty}, \text{TerritoryWithHighInequality}, \text{TerritoryWithHighUnemployment}, \text{TerritoryWithHighFoodInsecurity})$	Represents broad socio-economic stress arising from multiple stress dimensions, with OWA weights favouring the most strongly satisfied dimensions.
GreenTransitionReadiness	$\equiv 0.35 \text{ TerritoryWithHighCleanEnergyAccess} + 0.25 \text{ TerritoryWithStrongInstitutionalCapacity} + 0.20 \text{ TerritoryWithHighDigitalAccess} + 0.20 \text{ TerritoryWithStrongResilience}$	Represents readiness for the green transition through clean-energy access, institutional capacity, digital access, and resilience.
SustainableEconomicPerformance	$\equiv 0.30 \text{ TerritoryWithHighEconomicGrowth} + 0.25 \text{ TerritoryWithLowPoverty} + 0.25 \text{ TerritoryWithLowInequality} + 0.20 \text{ TerritoryWithHighCleanEnergyAccess}$	Represents sustainable economic performance characterised by low poverty, low inequality, and clean-energy access.
<i>Crisp combinations of bridge classes</i>		
SevereBasicNeedsDeprivation	$\equiv \text{TerritoryWithHighFoodInsecurity} \sqcap \text{TerritoryWithHighPoverty}$	Severe deprivation through poverty and food insecurity.
StructuralLockInRisk	$\equiv \text{TerritoryInEconomicDecline} \sqcap \text{TerritoryWithHighPoverty} \sqcap \text{TerritoryWithHighUnemployment} \sqcap \text{TerritoryWithLowDigitalAccess}$	Structural lock-in risk under economic decline, poverty, unemployment and low digital access.

Table 2. Cont.

Entity	DL Definition/Characteristics	Description (From the Annotations)
<i>Object properties</i>		
hasLocation	inverse isLocationOf; range CountrySubject	Relates a subject (e.g., a territorial snapshot) to its country-level geographic location.
hasDistance	symmetric, transitive; range Subject	Represents the distance relation between country-level geographic subjects, whose transitive closure captures reachability by land.
distanceFrom, distanceTo	range Subject	Relate a Distance individual to the two subjects between which the distance is measured.
hasEconomy	inverse isEconomyOf; range Economy	Relates a subject to its economic-profile nodes.
<i>Concrete features (data properties)</i>		
povertyRate, unemploymentRate, economicGrowthRate, foodInsecurityRate, digitalAccessRate, cleanEnergyAccessRate	functional; xsd:decimal	The six SDG features of a territorial snapshot (percentages; growth rates in $[-40, 40]$), copied from the underlying UNECE observations.
hasDistanceValue	functional; xsd:decimal	Represents the distance value, in kilometres, associated with a Distance individual, with value 0 for pairs of bordering countries.
<i>Fuzzy datatypes (32 in total)</i>		
Low/Moderate/High(feature), EconomicDecline, PositiveEconomicGrowth, ...	left-shoulder, trapezoidal and right-shoulder membership functions	Represents the linguistic terms defined over the concrete features, whose membership functions are reported in Listing 1 and Figure 2.

The reference slice used throughout the applications comprises the 31 country–year snapshots for 2021 for which all six SDG features are available. After including the neighbouring entities required for reasoning, the slice contains 248 individuals and 12,048 RDF triples, which are translated into a fuzzyDL Knowledge Base comprising approximately 1900 axioms. The countries included in the slice are Armenia, Austria, Belgium, Croatia, Czechia, Denmark, Estonia, Finland, France, Germany, Greece, Ireland, Israel, Italy, Latvia, Lithuania, Luxembourg, the Netherlands, Norway, Poland, Portugal, the Republic of Moldova, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, the United Kingdom, and the United States. Listing 5 illustrates the knowledge available to the reasoner for one of these snapshots.

Listing 5. Assertions associated with Italy_2021 in the slice Knowledge Base (excerpt), including feature values, crisp bridge-class memberships, and materialised membership degrees.

```

1 (instance Italy_2021 TerritorialSystem 1.0)
2 (related Italy_2021 Italy hasLocation 1.0)
3 (instance Italy_2021 (= povertyRate 0.9) 1.0)
4 (instance Italy_2021 (= unemploymentRate 11.9) 1.0)
5 (instance Italy_2021 (= economicGrowthRate 9.3) 1.0)
6 (instance Italy_2021 (= foodInsecurityRate 2.2) 1.0)
7 (instance Italy_2021 (= digitalAccessRate 74.9) 1.0)
8 (instance Italy_2021 (= cleanEnergyAccessRate 100.0) 1.0)
9 (instance Italy_2021 TerritoryWithLowPoverty 1.0)
10 (instance Italy_2021 TerritoryWithModerateUnemployment 1.0)
11 (instance Italy_2021 TerritoryWithHighEconomicGrowth 1.0)
12 (instance Italy_2021 SustainableEconomicPerformance 0.75)
13 (instance Italy_2021 GreenTransitionReadiness 0.55)
14 (instance Italy_2021 BroadSustainabilityOpportunity 0.75)

```

All queries considered in the following sections are expressed using the fuzzyDL query language [59]. In particular, $(\text{min-instance? } a \ C)$ returns the greatest lower bound n such that the Knowledge Base entails $C(a) \geq n$, whereas $(\text{max-instance? } a \ C)$ returns the least upper bound m such that the Knowledge Base entails $C(a) \leq m$; $(\text{min-subs? } C \ D)$ returns the greatest degree n such that the Knowledge Base entails $\langle D \sqsubseteq C \geq n \rangle$; $(\text{max-sat? } C)$ returns the supremum of $C(x)$ over all models; and $(\text{min-related? } a \ b \ R)$ returns the lower bound of $R(a, b)$. Reasoning is based on Łukasiewicz semantics, with

$$a \otimes b = \max(0, a + b - 1), \quad \neg a = 1 - a, \quad a \Rightarrow b = \min(1, 1 - a + b), \quad (11)$$

and

$$\exists R.C(x) = \sup_y (R(x, y) \otimes C(y)). \quad (12)$$

Since each slice contains the complete TBox and RBox but only a subset of the ABox, the relationship between the answers obtained from a slice and those obtained from the complete ontology can be characterised formally.

Proposition 1. Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{R}, \mathcal{A} \rangle$ be the consistent Knowledge Base of the complete ontology, and let $\mathcal{K}_S = \langle \mathcal{T}, \mathcal{R}, \mathcal{A}_S \rangle$, with $\mathcal{A}_S \subseteq \mathcal{A}$, be the Knowledge Base of a slice, both extended with the same analyst-defined axioms. Then, for every individual a and concept C , $(\text{min-instance? } a \ C)$ returns on $\mathcal{K}_S$ a degree no greater than that returned on $\mathcal{K}$, whereas $(\text{max-instance? } a \ C)$ returns a degree no smaller than that returned on $\mathcal{K}$.

Proof. Since $\mathcal{K}_S \subseteq \mathcal{K}$, every model of $\mathcal{K}$ is also a model of $\mathcal{K}_S$. Consequently, the infimum and supremum of $C^{\mathcal{I}}(a^{\mathcal{I}})$ are taken over a larger set of models for $\mathcal{K}_S$ than for $\mathcal{K}$. Enlarging the set of models can only decrease the infimum and increase the supremum, which proves the result. $\square$

Hence, the degrees returned by `min-instance?` on a slice are sound lower bounds on the corresponding degrees for the complete ontology, while those returned by `max-instance?` are sound upper bounds. These bounds coincide with those obtained from the complete ontology whenever the degree of C at a depends only on assertions contained in the slice. This condition holds for all concepts that do not traverse roles, including the composite concepts considered in Applications 1, 4, and 5 and the fuzzy concepts used in Application 3. Their degrees depend only on the feature values and graded assertions associated with a , all of which are preserved in the slice. Moreover, no schema axiom propagates degrees to a territorial snapshot from other individuals. Indeed, the schema contains neither universal restrictions nor nominals, no role assertion in the ontology has a territorial snapshot as its object, and the only role asserted for a snapshot in the translated Knowledge Base, `hasLocation`, does not occur in any schema concept definition, apart from its domain and range axioms. Accordingly, slices containing a single snapshot reproduce the corresponding membership degrees reported in these applications, as confirmed by the 276 single-snapshot slices evaluated in Section 5.7.

For concepts that traverse roles, equality with the complete ontology is guaranteed only when the slice contains all individuals relevant to the corresponding role chains. For this reason, Application 3 explicitly selects the countries together with their snapshots, ensuring that the associated economic-profile nodes are included in the slice. Similarly, Application 2 restricts the neighbouring snapshots to the year of interest and includes all of them in the slice (Section 5.2). Finally, since the ontology contains no nominals, `fuzzy-dl-owl2` evaluates the TBox queries `min-subs?` and `max-sat?` on a copy of the Knowledge Base with the ABox removed. Consequently, their answers are independent of the selected ontology slice.

4.4. Implementation Details

The reasoning pipeline is implemented in two small Python modules distributed together with this work. In addition `fuzzy-dl-owl2`, its only external dependency is `rdflib`. The code developed for this work is available in the GitHub repository indicated in the Data Availability Statement. In order to reproduce the experiments, download release v1.0.1, which includes the compressed SQLite index used in place of the complete OWL ontology to reduce data-loading and slice-extraction times.

The first module builds a lightweight SQLite index of the ontology import closure, which is created once and subsequently reused by all applications. The five ontology modules are parsed with `rdflib`, resulting in approximately 4.5 million RDF triples, a process that requires about five minutes on the reference platform. The extracted information is organised into the four tables shown in Listing 6: the ontology modules together with their role (schema, ABox, or annotations), the IRIs of all named individuals, the RDF statements having a named individual as subject, and the `owl:Axiom` reifications encoding the materialised membership degrees. Overall, the index contains 418,637 named individuals, approximately 4.2 million RDF statements, and 57,656 graded assertions, occupying about 2.6 GB on disk. Since the index is built only once, its construction cost is amortised across all subsequent executions.

Slice extraction is implemented by the function summarised in Listing 7. The complete schema modules are always included in the slice so that every extracted ontology contains the full TBox and RBox. Before translation, three sanitisation steps are applied to ensure compatibility with the `fuzzyDL` translator: `owl:imports` declarations are removed, `rdf:PlainLiteral` ranges are mapped to `xsd:string`, and class IRIs appearing inside Fuzzy OWL 2 annotations are replaced by their local names. The ABox is then assembled by performing a breadth-first traversal of the SQLite index starting from the selected in-

dividuals. All assertions associated with the selected individuals are copied into the slice together with the corresponding owl:Axiom annotations, while referenced individuals are recursively included up to one hop, thereby incorporating the required countries, observations, and economic-profile individuals without importing the entire ontology. Finally, entity names that are not valid fuzzyDL identifiers are renamed through a reversible mapping. For the reference slice (248 individuals and 12,048 RDF triples), the extraction step requires less than one second (see Section 5.7).

Listing 6. Schema of the SQLite index of the ontology import closure.

```

1 CREATE TABLE modules(file TEXT PRIMARY KEY, role TEXT);
2 -- 'schema' / 'abox' / 'annotations'
3 CREATE TABLE individuals(iri TEXT PRIMARY KEY);
4 -- named individuals of the closure
5 CREATE TABLE stmt(s TEXT, p TEXT, o TEXT,
6 -- statements whose subject is an
7     lit TEXT, dt TEXT, lang TEXT);
8 -- individual (object IRI or literal)
9 CREATE TABLE axiom_ann(s TEXT, p TEXT, o TEXT,
10 -- owl:Axiom reifications: the
11
12 prop TEXT, value TEXT);
13 -- materialised membership degrees
14 CREATE INDEX stmt_s ON stmt(s);
15 CREATE INDEX stmt_p ON stmt(p);
16 CREATE INDEX axiom_ann_spo ON axiom_ann(s, p, o);

```

Listing 7. Slice extraction (condensed): construction of an ontology slice by including the complete TBox and RBox, sanitising the ontology, expanding the ABox by one hop with graded assertions, and renaming entities to valid fuzzyDL identifiers.

```

1 def build_slice(individuals, out):
2     """Slice of the closure for the selected individuals, written to the OWL file 'out'."""
3     c = sdf_index.connect() # the SQLite index, built on first use
4     g = rdflib.Graph()
5     schema = [f for (f,) in c.execute("SELECT file FROM modules WHERE role = 'schema'")]
6     for f in schema: # full TBox/RBox of the schema modules
7         g.parse(ROOT / f, format="xml")
8     sanitise(g) # drop owl:imports; rdf:PlainLiteral -> xsd:string;
9     # "Class(IRI)" in fuzzy labels -> local names
10    queue, seen = list(individuals), set()
11    while queue: # selected individuals + 1-hop neighbourhood
12        iri = queue.pop()
13        if iri in seen:
14            continue
15        seen.add(iri)
16        for p, o, lit, dt, lang in c.execute(
17            "SELECT p,o,lit,dt,lang FROM stmt WHERE s=?", (iri,)
18        ):
19            add_statement(g, iri, p, o, lit, dt, lang)
20            if o and iri in individuals and is_individual(c, o):
21                queue.append(o) # only the selected pull their neighbours
22            copy_degree_annotations(
23                g, c, iri, p, o
24            ) # owl:Axiom with the materialised degree
25    rename_unsafe_local_names(g) # FuzzyDL identifiers only (fdl_safe)
26    g.serialize(destination=str(out), format="xml")

```

All five applications follow the same execution pattern. First, the required territorial snapshots are selected through SQL queries on the SQLite index. Second, the analyst-defined axioms are specified directly in fuzzyDL syntax. Third, the reasoning queries are generated automatically, producing one min-instance?, min-subs?, or max-sat? query for each combination of individual and concept. Finally, a common execution driver invokes the slice extraction and reasoning pipeline, stores the result of every query as a JSON record

together with the execution statistics discussed in Section 5.7, and generates the tables presented throughout this paper directly from the recorded results.

The execution driver also supports two optional features. The `patches` option applies documented textual substitutions to the translated fuzzyDL knowledge base, whereas `prune` removes cluster-membership assertions from the extracted slice (Applications 2 and 3; see Section 5.7). Consequently, each application can be reproduced by executing a single script, without requiring any manual intervention.

5. Applications

In this section, we evaluate the proposed framework through a set of complementary applications designed to examine its reasoning capabilities from different analytical perspectives. The evaluation covers fuzzy territorial profiling, spatial reasoning across neighbouring territories, comparisons among economic groupings, the dynamic definition of analyst-defined concepts together with coherence checks of the ontology vocabulary, and the analysis of socio-economic trajectories over time. Taken together, these applications demonstrate how the same ontology and reasoning infrastructure can support heterogeneous forms of fuzzy, relational, and temporal analysis, while accommodating different aggregation strategies and concepts introduced at query time. The section concludes with a sensitivity analysis, an assessment of the computational cost of the pipeline, and a discussion of the main methodological and implementation limitations affecting its current applicability and scalability.

5.1. Application 1: Fuzzy Territorial Profiling

The first application investigates the degree to which the UNECE territories were affected by different forms of socio-economic stress in 2021, and analyses how the answer changes under different aggregation policies.

For each of the 31 territorial snapshots, the reasoner evaluates the query (`min-instance? s C`) for the four composite concepts `BasicNeedsStress` (BNS), `EconomicStress` (ES), `StructuralDevelopmentStress` (SDS), and `BroadSocioEconomicStress` (BSES). The first three concepts are defined through compensatory weighted sums, whereas the fourth is defined by the OWA operator described in Listing 2. Overall, the application requires the evaluation of 124 queries on a single Knowledge Base.

The results are reported in Table 3, which lists the territorial snapshots having a positive membership degree in at least one of the four concepts. The remaining 26 territories receive degree 0 in all cases. Inspection of the underlying indicators (see Table 4) shows that, in 2021, none of the considered territories exhibited high poverty (defined as at least 20% of the population living below the international poverty line), high food insecurity, or negative economic growth. Consequently, unemployment is the only indicator contributing to the computed stress degrees. Greece (19.6%) and Serbia (29.8%) belong completely to the concept `HighUnemployment`, since the right-shoulder membership function reaches degree 1 at 18%. Spain (17.8%), Croatia (16.0%), and France (12.1%) belong to the same concept with degrees 0.97, 0.67, and 0.02, respectively.

For the three weighted-sum concepts, the resulting membership degree is simply the product of the unemployment membership degree and the corresponding criterion weight, namely 0.25 for BNS, 0.35 for ES, and 0.20 for SDS. This behaviour reflects the compensatory semantics of the `w-sum` operator: since only one of the three dimensions is active, it contributes only its assigned weight to the final degree. The OWA-based concept behaves differently. Using the weight vector (0.5, 0.25, 0.125, 0.125) applied to the membership degrees sorted in decreasing order, a single fully satisfied dimension

already yields a degree of 0.5. Consequently, BSES captures the interpretation that at least one significant form of socio-economic stress is present, assigning Greece and Serbia a degree of 0.5, compared with 0.35 under ES. The two aggregation policies therefore answer different questions: weighted sums quantify the overall amount of socio-economic stress, whereas the OWA operator highlights the presence of at least one prominent stress factor. The appropriate interpretation depends on the policy objective and is determined by the concept selected by the analyst.

Table 3. Lower bounds returned by the (min-instance?) queries for the membership degrees of the 2021 territorial snapshots in the four socio-economic stress concepts. Snapshots with degree 0 in all four concepts are omitted.

Snapshot	BNS	ES	SDS	BSES (OWA)
Greece 2021	0.250	0.350	0.200	0.500
Serbia 2021	0.250	0.350	0.200	0.500
Spain 2021	0.242	0.338	0.193	0.483
Croatia 2021	0.167	0.233	0.133	0.333
France 2021	0.004	0.006	0.003	0.008

Table 4. Feature values (%) of the five 2021 territorial snapshots with a positive membership degree in HighUnemployment, together with their corresponding membership degrees. Pov. = poverty; Unemp. = unemployment rate of persons without disability (SDG 8.5.2); Growth = economic growth; Food = food insecurity; Digital = digital access; Energy = clean-energy access; μ_{HU} = membership degree in HighUnemployment.

Snapshot	Pov.	Unemp.	Growth	Food	Digital	Energy	μ_{HU}
Greece 2021	0.7	19.6	9.9	6.3	78.5	100.0	1.00
Serbia 2021	1.7	29.8	9.1	14.8	81.2	81.1	1.00
Spain 2021	0.7	17.8	6.6	8.0	93.9	100.0	0.97
Croatia 2021	0.3	16.0	13.5	9.7	81.3	100.0	0.67
France 2021	0.1	12.1	6.6	6.6	86.1	100.0	0.02

5.2. Application 2: Stress Across Land Borders

The second application investigates the extent to which a territory is surrounded by economically stressed neighbouring territories and identifies the most asymmetric land borders in the 2021 UNECE dataset.

The slice contains the 31 complete territorial snapshots for 2021, together with 20 incomplete 2021 snapshots corresponding to countries that border them. It also includes the associated 51 countries, the relevant geopolitical entities, and 96 Distance individuals representing the bordering country pairs within this set. Two concepts are defined at query time (see Listing 8). The first, StressedNeighbourhood, relates a territorial snapshot to the EconomicStress of its neighbouring territories through the chain

`hasLocation ◦ borders ◦ isLocationOf,`

where `isLocationOf` is the inverse of `hasLocation`. Since each country is also the location of snapshots from other years and of its observations, the innermost filler is restricted to `Year2021 ⊓ EconomicStress`.

Listing 8. Analyst-defined axioms and queries (excerpt). The script generates one `borders` assertion for each bordering pair identified by the data property `sharesBorder`. The crisp analyst-defined concept `Year2021` is defined by the datatype restriction (`= referenceYear 2021`), while `LandReachableStress` is defined using the transitive role `hasDistance` provided by the schema.

```

1 (symmetric borders)
2 (related Italy France borders 1.0)
3 (define-concept Year2021 (= referenceYear 2021))
4 (define-concept StressedNeighbourhood
5   (some hasLocation (some borders (some isLocationOf (and Year2021 EconomicStress)))))
6 (define-concept LandReachableStress
7   (some hasLocation (some hasDistance (some isLocationOf (and Year2021 EconomicStress)))))
8 (define-concept ShortDistanceCheck (some hasDistanceValue ShortDistance))
9 (min-instance? Italy_2021 EconomicStress)
10 (min-instance? Italy_2021 StressedNeighbourhood)
11 (min-instance? Armenia_2021 LandReachableStress)
12 (min-related? Italy France hasDistance)
13 (min-instance? Distance_France_Italy BorderingDistance)
14 (min-instance? Distance_France_Italy ShortDistanceCheck)
15 (min-instance? Distance_France_Italy MidDistanceCheck)
16 (min-instance? Distance_France_Italy LongDistanceCheck)

```

The crisp analyst-defined concept `Year2021` is defined by the datatype restriction (`= referenceYear 2021`). Since the data property `referenceYear` is declared in the schema, this restriction is translated directly and no explicit `Year2021` assertions are required. Since the roles and `Year2021` are crisp, the existential restriction reduces to the maximum `EconomicStress` degree among the neighbouring 2021 snapshots.

Since the roles are crisp, the existential restriction

$$\exists R.C(x) = \sup_y (R(x, y) \otimes C(y)) \quad (13)$$

reduces to the maximum membership degree of `EconomicStress` among the 2021 snapshots of the neighbouring countries, regardless of the semantics adopted. The second concept, `LandReachableStress`, replaces `borders` with the transitive role `hasDistance`. Consequently, it measures the maximum stress degree among all territories reachable through land borders rather than only direct neighbours. Finally, three crisp queries verify the geopolitical layer by checking whether Italy and France share a border and by classifying the corresponding `Distance` individual according to the fuzzy datatype definitions.

The results are summarised in Table 5. Under the direct-border interpretation, two stress corridors clearly emerge. In the Balkans, Serbia ($ES = 0.35$) is the most stressed neighbour of Croatia and Romania, which therefore inherit neighbourhood degrees of 0.35 for `StressedNeighbourhood` and 0.25 for `BasicNeedsNeighbourhood`. Croatia ($ES = 0.23$) is in turn the most stressed neighbour of both Slovenia and Serbia. In the Iberian Peninsula, Spain ($ES = 0.34$) is the most stressed neighbour of Portugal and France, whereas Spain itself receives a neighbourhood degree of only 0.006, inherited from France. Among the borders connecting snapshots with determined membership degrees, the largest asymmetries occur between Serbia and Romania (0.35 versus 0) and between Spain and Portugal (0.34 versus 0), followed closely by Spain and France (0.34 versus 0.006). Countries whose most stressed neighbour is France (Belgium, Germany, Italy, Luxembourg, and Spain) all inherit the same degree of 0.006. Greece, whose four neighbours (Albania, Bulgaria, North Macedonia, and Türkiye) have only incomplete 2021 snapshots, inherits a degree of 0.047 for `StressedNeighbourhood` and 0.033 for `BasicNeedsNeighbourhood` from Bulgaria, whose stress degrees are determined despite one missing feature; the undetermined stress degrees of the other three neighbours do not increase the corresponding lower bounds. Similarly, Switzerland and Austria inherit a degree of 0.004 from Liechtenstein, while

Armenia inherits 0.040 for BasicNeedsNeighbourhood from Georgia. These lower bounds coincide with those obtained using all 2021 snapshots in the complete ontology.

The transitive interpretation yields a substantially different result. Since hasDistance is declared transitive, both Italy and Armenia obtain LandReachableStress = 0.35, as each is connected by land to Serbia and Greece through the transitive closure of the border relation.

Unlike the neighbourhood degrees, this value is, in general, only a lower bound (Proposition 1), since the slice does not contain the entire connected component of hasDistance. In this case, however, it already equals the largest lower bound of EconomicStress among all 2021 snapshots.

Finally, the crisp queries, summarised in Table 6, verify the expected semantics of the geopolitical layer. In particular, Italy and France are related through hasDistance with degree 1, and Distance_France_Italy is classified as both BorderingDistance and ShortDistance, but not as LongDistance.

Table 5. Lower bounds of the membership degrees of the 31 complete 2021 snapshots in EconomicStress, StressedNeighbourhood, and BasicNeedsNeighbourhood, together with the neighbouring territory attaining the corresponding supremum. The neighbourhood also includes the incomplete 2021 snapshots of bordering countries. For Armenia, the neighbour attaining BasicNeedsNeighbourhood is reported. Snapshots with degree 0 in all three concepts are omitted.

Snapshot	ES (Own)	ES (Neighbourhood)	BNS (Neighbourhood)	Most Stressed Neighbour
Romania 2021	0.000	0.350	0.250	Serbia
Croatia 2021	0.233	0.350	0.250	Serbia
Portugal 2021	0.000	0.338	0.242	Spain
France 2021	0.006	0.338	0.242	Spain
Slovenia 2021	0.000	0.233	0.167	Croatia
Serbia 2021	0.350	0.233	0.167	Croatia
Greece 2021	0.350	0.047	0.033	Bulgaria
Switzerland 2021	0.000	0.008	0.004	Liechtenstein
Austria 2021	0.000	0.008	0.000	Liechtenstein
Spain 2021	0.338	0.006	0.004	France
Belgium 2021	0.000	0.006	0.004	France
Germany 2021	0.000	0.006	0.004	France
Italy 2021	0.000	0.006	0.004	France
Luxembourg 2021	0.000	0.006	0.004	France
Armenia 2021	0.000	0.000	0.040	Georgia (BNS)

Table 6. Verification queries for the geopolitical layer and the transitive closure of the border relation, together with the corresponding lower bounds returned by the reasoner.

Query	Result
(min-related? Italy France hasDistance)	1.00
(min-instance? Distance_France_Italy BorderingDistance)	1.00
(min-instance? Distance_France_Italy ShortDistance)	1.00
(min-instance? Distance_France_Italy LongDistance)	0.00
(min-instance? Italy_2021 LandReachableStress)	0.35
(min-instance? Armenia_2021 LandReachableStress)	0.35

5.3. Application 3: Economic Groupings

The third application investigates whether euro-area territories differ from the other high-income and middle-income UNECE territories in terms of green-transition readiness and sustainable economic performance.

The geopolitical layer of the ontology associates each country with one or more economic profiles through the role hasEconomy, whose fillers are instances of the class Economy

(e.g., EURO_AREA, HIGH_INCOME, and UPPER_MIDDLE_INCOME), derived from the World Bank and Organisation for Economic Co-operation and Development (OECD) classifications. Since fuzzyDL does not support nominals, the analyst introduces three crisp concepts representing the economic groups by asserting them on the corresponding economy nodes and combines them with the fuzzy concepts already defined in the ontology (see Listing 9). For example,

$$\text{EuroAreaGreenReadiness} \equiv \text{EuroAreaTerritory} \sqcap \text{GreenTransitionReadiness}.$$

Since EuroAreaTerritory is crisp, this concept inherits exactly the membership degree of GreenTransitionReadiness for euro-area territories, while evaluating to 0 for all other territories, regardless of the semantics adopted. Overall, eight concepts are queried over the 31 territorial snapshots, and the resulting membership degrees are subsequently aggregated by economic group.

Listing 9. Analyst-defined group concepts on the economy nodes and their composition with the fuzzy concepts of the ontology.

```

1 (instance CL10344_EURO_AREA EuroArea 1.0)
2 (instance CL10351_1_HIGH_INCOME HighIncome 1.0)
3 (instance CL10352_2_UPPER_MIDDLE_INCOME MiddleIncome 1.0)
4 (instance CL10353_3_LOWER_MIDDLE_INCOME MiddleIncome 1.0)
5 (define-concept EuroAreaTerritory (some hasLocation (some hasEconomy EuroArea)))
6 (define-concept HighIncomeTerritory (some hasLocation (some hasEconomy HighIncome)))
7 (define-concept MiddleIncomeTerritory (some hasLocation (some hasEconomy MiddleIncome)))
8 (define-concept EuroAreaGreenReadiness (and EuroAreaTerritory GreenTransitionReadiness))
9 (define-concept MiddleIncomeStructuralStress
10 (and MiddleIncomeTerritory StructuralDevelopmentStress))

```

The results show that the crisp component of the query is evaluated exactly as expected. Seventeen snapshots belong to the euro area, nine correspond to high-income territories outside the euro area (Croatia, Czechia, Denmark, Israel, Norway, Poland, Sweden, Switzerland, and the United States), and four belong to the middle-income group (Armenia, the Republic of Moldova, Romania, and Serbia). All remaining snapshots receive degree 0 in the corresponding grouping concepts because countries are connected only to their own economic-profile nodes.

The fuzzy component, summarised in Table 7, provides a more interesting interpretation. The concept GreenTransitionReadiness is defined as a weighted sum of four dimensions, although only two are currently supported by the available data sources, namely clean-energy access (weight 0.35) and digital access (weight 0.20). Both indicators are almost fully satisfied throughout the UNECE region in 2021. Consequently, the concept attains the same degree (0.55) for all euro-area and other high-income territories, whereas only the middle-income group has a lower mean degree (0.47). This difference is entirely attributable to Serbia, whose clean-fuel access (81.1%) lies within the transition region of HighCleanEnergyAccess, yielding a membership degree of 0.07 and, consequently, a readiness degree of 0.23. Armenia, the Republic of Moldova, and Romania instead attain a readiness degree of 0.55, equal to that of the high-income territories.

A similar behaviour is observed for SustainableEconomicPerformance, whose weighted-sum definition reaches the maximum attainable value of 0.75 for most high-income territories. Lower values are observed for nine territories. Finland (0.500), Norway (0.590), Germany (0.610), Czechia (0.670), Austria (0.690), Sweden (0.710), and Switzerland (0.730) have growth rates (2.5%, 3.4%, 3.6%, 4.2%, 4.4%, 4.6%, and 4.8%, respectively) that do not fully satisfy PositiveEconomicGrowth, whose transition region ends at 5%. Serbia attains 0.565 because its clean-fuel access only marginally satisfies HighCleanEnergyAccess,

while Armenia attains 0.747 because its poverty rate (2.1%) lies just above the lower break-point of *LowPoverty*. The lower mean degree of the middle-income group (0.70) is therefore driven primarily by Serbia. The composed concepts reported in Table 8 behave exactly as predicted by the Łukasiewicz conjunction. In particular, *EuroAreaGreenReadiness* coincides with the corresponding readiness degree for the 17 euro-area territories and evaluates to 0 elsewhere, whereas *MiddleIncomeStructuralStress* is positive only for Serbia, reflecting the unemployment component of *StructuralDevelopmentStress*.

Table 7. Lower bounds of the membership degrees of the 2021 territorial snapshots in *GreenTransitionReadiness* (GTR), *SustainableEconomicPerformance* (SEP), and *StructuralDevelopmentStress* (SDS), aggregated by economic grouping. For each group, the mean and maximum lower bounds are reported.

Group (2021 Snapshots)	n	GTR		SEP		SDS
		Mean	Max	Mean	Max	Mean
Euro area	17	0.550	0.550	0.724	0.750	0.023
High income, outside the euro area	10	0.550	0.550	0.720	0.750	0.013
Middle income	4	0.469	0.550	0.703	0.750	0.050

Table 8. Membership degrees in the composed concepts *EuroAreaGreenReadiness* and *MiddleIncomeStructuralStress*, together with *SustainableEconomicPerformance* (SEP), for the 2021 territorial snapshots. Only snapshots with a positive degree in at least one of the two composed concepts are reported.

Snapshot	<i>EuroAreaGreenReadiness</i>	<i>MiddleIncomeStructural-Stress</i>	SEP
Austria 2021	0.550	0.000	0.690
Belgium 2021	0.550	0.000	0.750
Estonia 2021	0.550	0.000	0.750
Finland 2021	0.550	0.000	0.500
France 2021	0.550	0.000	0.750
Germany 2021	0.550	0.000	0.610
Greece 2021	0.550	0.000	0.750
Ireland 2021	0.550	0.000	0.750
Italy 2021	0.550	0.000	0.750
Latvia 2021	0.550	0.000	0.750
Lithuania 2021	0.550	0.000	0.750
Luxembourg 2021	0.550	0.000	0.750
Netherlands 2021	0.550	0.000	0.750
Portugal 2021	0.550	0.000	0.750
Slovakia 2021	0.550	0.000	0.750
Slovenia 2021	0.550	0.000	0.750
Spain 2021	0.550	0.000	0.750
Serbia 2021	0.000	0.200	0.565

Overall, the application shows that, according to the current ontology and the 31 complete territorial snapshots for 2021, euro-area and other high-income territories are indistinguishable with respect to the readiness concepts, whereas the four middle-income snapshots differ primarily in exhibiting higher levels of socio-economic stress. Since these four snapshots represent only 4 of the 19 middle-income countries in the region, however, this difference cannot be generalised to the group as a whole (Section 5.7).

At the same time, the results reveal a limitation of the current ontology: because the institutional-capacity and resilience dimensions are not yet represented, the readiness concepts provide only limited discrimination among advanced economies. The complete set of 248 reasoning queries is evaluated in a single execution (see Section 5.7).

5.4. Application 4: Analyst-Defined Concepts, Hedges and Semantic Checks

The fourth application illustrates how an analyst can extend the ontology with new concepts at query time and use the reasoner not only to classify territorial snapshots, but also to verify the semantic properties of the ontology itself.

Four concepts are introduced dynamically (see Listing 10). The first, *JoblessGrowth*, combines positive economic growth with high unemployment through a Łukasiewicz conjunction of two datatype restrictions. The second, *VeryHighUnemployment*, applies the linguistic hedge *very* to the fuzzy datatype *HighUnemployment*. The third, *SocialAlarm*, aggregates the three stress concepts by means of an OWA operator assigning weights (0.5, 0.3, 0.2) to the ordered membership degrees. Finally, *InclusiveRebound* defines a new compensatory weighted sum over bridge classes. Besides the instance queries, a number of TBox queries investigate the semantic relations among the concepts already defined in the ontology, including graded subsumption and satisfiability. All instance and TBox queries are evaluated in a single reasoning session.

The instance-query results are reported in Table 9, whereas the TBox queries are summarised in Table 10. The concept *JoblessGrowth* identifies territories combining strong post-pandemic economic growth with high unemployment. Greece (growth 9.9%, unemployment 19.6%) and Serbia (growth 9.1%, unemployment 29.8%) satisfy the concept completely, while Spain and Croatia belong to it with degrees 0.97 and 0.67, respectively. Under the Łukasiewicz conjunction, the resulting degree coincides with the membership degree in *HighUnemployment*, because *PositiveEconomicGrowth* is fully satisfied for growth rates above 5%.

The linguistic hedge *very*, implemented as a linear modifier with parameter 2 (i.e., $lm(2)$; see Section 3), sharpens the corresponding fuzzy concept by reducing intermediate membership degrees while preserving the extreme values. Consequently, Spain decreases from 0.97 to 0.93, Croatia from 0.67 to 0.33, and France from 0.02 to 0.01, whereas Greece and Serbia remain at degree 1. The concept *SocialAlarm* aggregates the three stress indicators through the specified OWA operator. For example, Greece has the ordered stress vector (0.35, 0.25, 0.20), yielding the aggregated degree

$$0.5 \cdot 0.35 + 0.3 \cdot 0.25 + 0.2 \cdot 0.20 = 0.29.$$

Finally, *InclusiveRebound* evaluates to high degrees for most territories in 2021, reflecting the widespread post-pandemic economic recovery combined with generally low poverty levels. The Republic of Moldova and Romania attain degree 1, whereas most euro-area territories reach 0.75 because their unemployment levels are not classified as *low*. Greece, Serbia, and Spain also receive degree 0.75, illustrating that a compensatory weighted sum may mask high unemployment when the remaining dimensions are fully satisfied.

The TBox queries assess the internal coherence of the ontology vocabulary independently of the ABox data. The crisp concept *SevereBasicNeedsDeprivation* implies *BasicNeedsStress* to degree 0.75, corresponding to the sum of the poverty and food-insecurity weights in the weighted sum, while the complementary satisfiability query confirms that a territory can satisfy *SevereBasicNeedsDeprivation* without fully satisfying *BasicNeedsStress*, with degree 0.25. Likewise, *StructuralLockInRisk* implies *EconomicStress* to degree 1, since the minimum of its four defining conjuncts never exceeds the weighted sum of the three they share, whose weights total one, whereas *SevereBasicNeedsDeprivation* implies *StructuralDevelopmentStress* only to degree 0.45. The satisfiability queries also show that opposite linguistic terms associated with the same feature cannot hold simultaneously under the Łukasiewicz conjunction. In particular, *LowPoverty* and *HighPoverty* are completely disjoint, while *ModeratePoverty* and

HighPoverty overlap only where their membership degrees sum to one, and, thus, the conjunction degree is in both cases 0. Finally, the concept JoblessGrowth is satisfiable to degree 1, confirming that the ontology consistently admits territories simultaneously characterised by positive economic growth and high unemployment, a situation that is indeed observed in the underlying data.

Listing 10. Analyst-defined concepts, axioms, and TBox queries (excerpt).

```

1 (define-modifier very linear-modifier(2.0))
2 (define-concept JoblessGrowth
3   (and (some economicGrowthRate PositiveEconomicGrowth) (some unemploymentRate
4     HighUnemployment)))
5 (define-concept VeryHighUnemployment (some unemploymentRate (very HighUnemployment)))
6 (define-concept SocialAlarm
7   (owa (0.5 0.3 0.2) (BasicNeedsStress EconomicStress StructuralDevelopmentStress)))
8 (define-concept InclusiveRebound
9   (w-sum (0.5 TerritoryWithHighEconomicGrowth) (0.25 TerritoryWithLowUnemployment)
10    (0.25 TerritoryWithLowPoverty)))
11 (min-subs? BasicNeedsStress SevereBasicNeedsDeprivation)
12 (min-subs? EconomicStress StructuralLockInRisk)
13 (max-sat? (and (some povertyRate LowPoverty) (some povertyRate HighPoverty)))
14 (max-sat? (and (some povertyRate ModeratePoverty) (some povertyRate HighPoverty)))
15 (max-sat? (and SevereBasicNeedsDeprivation (not BasicNeedsStress)))

```

Table 9. Membership degrees of the 2021 territorial snapshots in the analyst-defined concepts. Only snapshots with a positive degree in at least one stress-related concept or a high degree of InclusiveRebound are reported.

Snapshot	JoblessGrowth	VeryHigh-Unemployment	SocialAlarm	InclusiveRebound
Greece 2021	1.000	1.000	0.290	0.750
Serbia 2021	1.000	1.000	0.290	0.750
Spain 2021	0.967	0.933	0.280	0.750
Croatia 2021	0.667	0.333	0.193	0.750
France 2021	0.017	0.008	0.005	0.750
Israel 2021	0.000	0.000	0.000	0.915
Luxembourg 2021	0.000	0.000	0.000	0.955
Netherlands 2021	0.000	0.000	0.000	0.930
Republic of Moldova 2021	0.000	0.000	0.000	1.000
Romania 2021	0.000	0.000	0.000	1.000
United Kingdom 2021	0.000	0.000	0.000	0.935

Table 10. TBox queries and the corresponding graded subsumption and satisfiability results.

Query	Answer
(min-subs? BasicNeedsStress SevereBasicNeedsDeprivation)	0.75
(min-subs? EconomicStress StructuralLockInRisk)	1.0
(min-subs? StructuralDevelopmentStress SevereBasicNeedsDeprivation)	0.45
(max-sat? (and (some povertyRate LowPoverty) (some povertyRate HighPoverty)))	0.0
(max-sat? (and (some povertyRate ModeratePoverty) (some povertyRate HighPoverty)))	0.0
(max-sat? (and (some unemploymentRate ModerateUnemployment) (some unemploymentRate HighUnemployment)))	0.0
(max-sat? (and SevereBasicNeedsDeprivation (not BasicNeedsStress)))	0.25
(max-sat? (and StructuralLockInRisk (not EconomicStress)))	0.0
(max-sat? JoblessGrowth)	1.0

5.5. Application 5: Temporal Trajectories

The fifth application analyses the temporal evolution of socio-economic conditions in southern Europe, investigating whether economic stress has decreased since 2015 and which territories have shown the greatest improvements in sustainable economic performance.

The ontology slice contains every complete territorial snapshot of Greece, Spain, Italy, Croatia, and Portugal between 2015 and 2023, whenever all six SDG features are available. Four concepts are evaluated for each country–year snapshot: the ontology concepts `EconomicStress` and `SustainableEconomicPerformance`, together with the analyst-defined concepts `JoblessGrowth` (introduced in Section 5.4) and

$$\text{Recovery} \equiv \text{SustainableEconomicPerformance} \sqcap \neg \text{EconomicStress},$$

which identifies territories simultaneously exhibiting high sustainable economic performance and low economic stress (Listing 11). The resulting membership degrees are interpreted as temporal trajectories and are reported in Tables 11 and 12.

The trajectories reveal two dominant patterns. The first corresponds to the economic impact of the COVID-19 pandemic. In 2020, the degree of `EconomicStress` increases sharply, reaching 0.75 for Spain, Greece, and Croatia, and 0.48 for Portugal. This increase is driven primarily by the activation of the `TerritoryInEconomicDecline` component (weight 0.40), whose left-shoulder membership function, with its transition ending at −5%, maps the economic growth rates of Spain (−11.4%), Greece (−9.0%), Croatia (−7.6%), and Portugal (−8.4%) to the maximum membership degree of 1.00. These values combine with the unemployment component, which is fully active in Spain, Greece, and Croatia, and partially active in Portugal. At the same time, `SustainableEconomicPerformance` drops uniformly to 0.45, corresponding to the contribution of the two dimensions unaffected by economic growth, namely low poverty (weight 0.25) and clean-energy access (weight 0.20).

The second pattern is the gradual reduction in unemployment-driven stress following the pandemic. Greece remains at the maximum contribution of the unemployment component (0.35) from 2015 to 2021, decreasing only in 2022 (0.24) and 2023 (0.16). Spain follows a similar trajectory, with degrees of 0.34 in 2021 and 0.21 in 2022, before rising again to 0.24 in 2023 as its unemployment rate increases from 15.6% to 16.2%. By contrast, Portugal leaves the stressed region much earlier, with its degree decreasing to 0.10, 0.02, and 0.00 between 2017 and 2019.

The analyst-defined concepts provide additional insight into these temporal patterns. `JoblessGrowth` identifies periods in which strong economic growth coexists with high unemployment. This situation occurs for Spain in 2015 and for Croatia between 2016 and 2018, but disappears during the pandemic because the positive-growth condition is no longer satisfied. The concept `Recovery` is more restrictive, requiring both high sustainable economic performance and the absence of economic stress. Consequently, it reaches its maximum degree of 0.75 only for Italy and Portugal in 2021–2022 and approximately 0.70 for Croatia in 2022–2023, precisely when sustainable economic performance attains its maximum value while economic stress has vanished under the Łukasiewicz conjunction. Since all yearly evaluations are computed using the same ontology, the same rule base, and the same reasoning procedure, the resulting trajectories provide a homogeneous comparison across time.

Listing 11. Analyst-defined axioms.

```

1 (define-concept JoblessGrowth
2   (and (some economicGrowthRate PositiveEconomicGrowth) (some unemploymentRate
3     HighUnemployment)))
4 (define-concept Recovery (and SustainableEconomicPerformance (not EconomicStress)))
5 (min-instance? Greece_2015 EconomicStress)
6 (min-instance? Greece_2015 Recovery)

```

Table 11. Lower bounds of the membership degrees of the five countries in EconomicStress (ES) and SustainableEconomicPerformance (SEP) by year. A dash (–) denotes an incomplete territorial snapshot.

Year	Greece		Spain		Italy		Croatia		Portugal	
	ES	SEP	ES	SEP	ES	SEP	ES	SEP	ES	SEP
2015	0.35	0.45	0.35	0.66	–	–	0.35	0.57	0.35	0.45
2016	0.35	0.45	0.35	0.53	–	–	0.35	0.71	0.25	0.47
2017	0.35	0.45	0.35	0.51	–	–	0.35	0.72	0.10	0.60
2018	0.35	0.48	0.35	0.45	–	–	0.35	0.68	0.02	0.55
2019	0.35	0.49	0.35	0.45	–	–	0.23	0.66	0.00	0.51
2020	0.75	0.45	0.75	0.45	–	–	0.75	0.45	0.48	0.45
2021	0.35	0.75	0.34	0.75	0.00	0.75	0.23	0.75	0.00	0.75
2022	0.24	0.75	0.21	0.75	0.00	0.75	0.05	0.75	0.00	0.75
2023	0.16	0.65	0.24	0.50	0.00	0.45	0.00	0.61	0.00	0.55

Table 12. Lower bounds of the membership degrees in the analyst-defined concepts JoblessGrowth (JG) and Recovery (Rec.) by year.

Year	Greece		Spain		Italy		Croatia		Portugal	
	JG	Rec.	JG	Rec.	JG	Rec.	JG	Rec.	JG	Rec.
2015	0.00	0.10	0.70	0.31	–	–	0.40	0.22	0.00	0.10
2016	0.00	0.10	0.27	0.18	–	–	0.87	0.36	0.00	0.22
2017	0.00	0.10	0.20	0.16	–	–	0.90	0.37	0.00	0.49
2018	0.10	0.13	0.00	0.10	–	–	0.77	0.33	0.00	0.53
2019	0.13	0.14	0.00	0.10	–	–	0.35	0.43	0.00	0.51
2020	0.00	0.00	0.00	0.00	–	–	0.00	0.00	0.00	0.00
2021	1.00	0.40	0.97	0.41	0.00	0.75	0.67	0.52	0.00	0.75
2022	0.70	0.51	0.60	0.54	0.00	0.75	0.13	0.70	0.00	0.75
2023	0.12	0.49	0.00	0.26	0.00	0.45	0.00	0.61	0.00	0.55

5.6. Sensitivity Analysis

We assessed how the conclusions of Sections 5.1 and 5.3–5.5 depend on the aggregation weights of the eight weighted-sum and OWA operators involved, ranging from moderate perturbations of the ontology weights to arbitrary weight vectors, following the uncertainty-analysis approach used for composite indicators [60]. For each weight vector, we generated 500 vectors $\mathbf{u}$ uniformly over the simplex of non-negative weights summing to 1 and replaced the ontology weights $\mathbf{w}_0$ with

$$(1 - \lambda)\mathbf{w}_0 + \lambda\mathbf{u},$$

so that λ controls the share of the total weight redistributed at random, with $\lambda = 1$ yielding an arbitrary weight vector. In each scenario, the perturbed weights replace those in the corresponding concept definitions, the membership degrees materialised in the ontology for those concepts are removed, and the reasoner recomputes the degrees on single-snapshot slices (Section 4.3). With the original ontology weights, this procedure reproduces all the degrees reported in the applications.

In Table 13, we report the proportion of scenarios in which each weight-dependent qualitative conclusion remains valid. Three conclusions hold for any strictly positive weights, up to the four-decimal precision of the degrees returned by the reasoner. The remaining three fail when either an individual weight or a ratio between two weights crosses the threshold reported in the final column, which also gives the corresponding value under the ontology weights. By contrast, replacing the weighted sums with their non-compensatory variants based on Łukasiewicz or Gödel conjunction causes the stress concepts to evaluate to 0 for all snapshots. The scripts, random seeds, and complete results are available in the public repository.

In Table 14, we compare the degrees obtained under the random weight scenarios with the baseline values. No pair of snapshots is ever reversed for the 2021 stress concepts, GTR, or SEP in Application 5. By contrast, up to 13% of snapshot pairs are reversed for InclusiveRebound and Recovery, while the maximum reversal rate is 4% for SEP in 2021 and for ES in Application 5.

Table 13. Share (%) of the 500 random scenarios in which each qualitative conclusion of the applications holds, for each proportion λ of the total weight redistributed at random ($\lambda = 1$: arbitrary weights). *Fails when*: condition on the weights under which the conclusion fails, with the corresponding value under the ontology weights reported in brackets. w_{HU} and w_{ED} denote the weights of HighUnemployment and EconomicDecline in EconomicStress; w_{HCEA} and w_{PG} denote the weights of HighCleanEnergyAccess and PositiveEconomicGrowth in SustainableEconomicPerformance. The symbol $\lesssim$ denotes an approximate condition that is violated in fewer than 1% of the scenarios, while >99 denotes a share of at least 99.5% but below 100%. Conclusions that do not depend on aggregation weights, namely those involving JoblessGrowth and the unemployment-ceiling results, are unaffected.

App.	Conclusion Stated in the Paper	Weights	
		$\lambda = 0.25/0.5/1$	Fails When
1	Greece and Serbia have the highest ES, strictly above every other territory BSES $\geq$ ES for every territory	100/100/ >99	only if $w_{HU} < 0.003$, degrees tied at 4 decimals (0.35)
		87/57/36	first OWA weight of BSES $< w_{HU}$ (0.5, 0.35)
3	GTR: euro area = other high income; middle income lowest SEP: middle income lowest	100/100/100	never
		97/87/74	$w_{HCEA}/w_{PG} \lesssim 0.43$ (0.67)
5	ES peaks in 2020 in all four countries Recovery maximal only for Italy and Portugal, 2021–2022	86/66/51	$w_{ED}/w_{HU} \leq 0.78$, Portugal (1.14)
		100/100/100	never

Table 14. Comparison between the degrees obtained under the random weight scenarios and the baseline, for each value of λ . *Band*: widest 5th-to-95th percentile interval of the membership degree of any snapshot (lower is better). *Rev.*: largest proportion of snapshot pairs whose ordering is reversed relative to the baseline (lower is better). Pairs that become tied are not counted as reversals, since the reasoner returns degrees rounded to four decimal places.

Concept	Snapshots	$\lambda = 0.25$		$\lambda = 0.5$		$\lambda = 1$	
		Band	Rev. (%)	Band	Rev. (%)	Band	Rev. (%)
BNS	2021 (31)	0.191	0.0	0.382	0.0	0.765	0.0
ES	2021 (31)	0.189	0.0	0.378	0.0	0.756	0.0
SDS	2021 (31)	0.134	0.0	0.267	0.0	0.535	0.0
BSES (OWA)	2021 (31)	0.159	0.0	0.317	0.0	0.634	0.0
SocialAlarm (OWA)	2021 (31)	0.115	0.0	0.231	0.0	0.516	0.0
GTR	2021 (31)	0.178	0.0	0.357	0.0	0.714	0.0
SEP	2021 (31)	0.169	0.9	0.337	1.7	0.675	3.8
InclusiveRebound	2021 (31)	0.185	9.5	0.369	10.6	0.738	12.6
ES	App. 5 (39)	0.189	2.9	0.378	3.9	0.756	4.2
SEP	App. 5 (39)	0.176	0.0	0.352	0.0	0.703	0.0
Recovery	App. 5 (39)	0.254	5.2	0.486	9.1	0.858	12.5

5.7. Cost and Limitations

Table 15 reports the execution time of each step of the pipeline for the five applications on a laptop (Apple Silicon, Python 3.11, fuzzy-dl-owl2 1.0.30, and Gurobi 13.0.2). Slice extraction from the SQLite index requires only a few seconds. The OWL 2-to-fuzzyDL translation is the dominant step. For slices containing between 1 and 31 snapshots from 2021, its cost grows approximately quadratically with slice size because the translator retrieves the annotations of each graded axiom through a separate query over the entire

slice graph. Parsing the resulting fuzzyDL Knowledge Base and constructing the initial MILP require negligible time. By contrast, each reasoning query corresponds to solving a single MILP instance over the complete ABox of the slice, with an observed cost that grows approximately linearly with slice size, from about 0.1 s for a single-snapshot slice to a few seconds for the 31-snapshot reference slice. Overall, the proposed pipeline is suitable for interactive reasoning over ontology slices containing a few hundred individuals. Concepts that do not traverse roles can instead be evaluated on single-snapshot slices, which yield the same membership degrees (Section 4.3) and can be processed in parallel. In a re-measurement on the same laptop, a single-snapshot slice required approximately 11 s for translation and 0.15 s per reasoning query. Using four worker processes, the 276 complete snapshots from 2015–2023, with six concepts evaluated for each snapshot, were processed in approximately 30 minutes. By contrast, the complete fuzzyDL translation of the ontology, comprising 2.96 million axioms, is distributed as an export but is not used for reasoning in this work. All reported results are computed on ontology slices, so the computational cost of the pipeline is governed by the size of the selected slice rather than by that of the complete ontology.

Several limitations should nevertheless be noted. First, the degrees reported in the applications are the lower bounds returned by *min-instance?*. A missing value is not interpreted as zero: it contributes nothing to the lower bound of a weighted sum and forces the lower bound of a *w-sum-zero* concept to 0. For this reason, the reference analysis is restricted to the 31 of the 56 UNECE snapshots for 2021 for which all six features are available. Of the remaining 25 snapshots, 22 lack the unemployment rate. This restriction is not neutral: the reference slice contains 17 of the 19 euro-area countries but only 4 of the 19 middle-income countries, so the conclusions of Applications 1–4 primarily concern European Union and high-income territories. Under open-world semantics, combining *min-instance?* with *max-instance?* provides lower and upper bounds on the degrees of the excluded snapshots. In particular, 11 of them have an upper bound of 0.60 for *EconomicStress*, exceeding the degree 0.35 reported for Greece and Serbia. The reported degrees should therefore be interpreted as certified lower bounds rather than as a ranking of the entire region. Application 2 additionally incorporates the incomplete snapshots of neighbouring countries (Section 5.2). Moreover, five dimensions represented in the ontology, namely inequality, climate vulnerability, institutional capacity, adaptive capacity, and resilience, currently have no associated data source. Consequently, the lower bounds of GTR and SEP cannot exceed 0.55 and 0.75, respectively.

Second, the degrees depend on the breakpoints and weights defined in the ontology, which were established through expert judgement by the SDF project team without a formal elicitation procedure (Section 4.2), as well as on the choice of aggregation operators. The sensitivity of the results to these modelling choices is quantified in Section 5.6. Since additional concepts can be introduced at query time, analysts may define alternative breakpoints, weighting schemes, and aggregation operators without modifying the underlying ontology, as illustrated in Application 4.

Finally, the slice-construction algorithm follows only the object-property assertions of explicitly selected individuals. Consequently, Applications 2 and 3 explicitly include both the territorial snapshots and the corresponding countries. Each selected country is associated with approximately one hundred cluster-membership assertions describing its participation in the various classification schemes represented in the ontology. Although these assertions are irrelevant to the applications, they substantially increase the translation time of the Fuzzy OWL 2 converter. In our experiments, translating the unpruned slice containing all cluster memberships did not complete after 80 minutes. The pipeline therefore removes these assertions before translation through the *prune_slice* procedure while

preserving the economic-profile and border relationships required for reasoning. The sizes of the resulting slices are reported in Table 15.

Table 15. Cost of the five applications: number of individuals and RDF triples in each ontology slice (for Applications 2 and 3, before pruning the cluster-membership assertions, 23,431 triples are removed), together with the execution times for slice extraction, OWL 2-to-fuzzyDL translation, and query evaluation (total and average per query).

Application	Indiv.	Triples	Slice (s)	Transl. (s)	Queries	Query Time (s): Total/Mean
1—profiling	248	12,048	0.8	125.8	124	309.5/2.50
2—borders	2653	34,212	2.2	224.2	99	265.8/2.68
3—groupings	2616	33,953	2.2	199.6	248	670.1/2.70
4—analyst concepts	248	12,048	0.8	124.7	133	301.3/2.27
5—trajectories	278	12,730	0.6	175.9	156	468.5/3.00

6. Conclusions

We have shown that a Fuzzy OWL 2 ontology and a fuzzy DL reasoner can answer geopolitical questions about sustainable development with a semantics that is explicit in every step: the linguistic terms are fuzzy datatypes, the aggregation policies are weighted sums or OWA operators, the geopolitical structure is a set of roles along which existential restrictions propagate degrees, and the analyst can extend the vocabulary at query time without touching the ontology. On the 2021 UNECE data, the reasoner located the territories under socio-economic stress and explained their degrees, propagated stress across land borders, contrasted economic groupings and verified the coherence of the linguistic vocabulary, and traced trajectories over time, all with fuzzy-d1-owl2 and a MILP solver on slices of the ontology of the SDF project. Future work includes populating the five dimensions for which no data source is currently available, formally eliciting the membership functions and aggregation weights with domain experts, the use of the crisp reasoners on the same ontology for the non-fuzzy part of the questions, and a translation of the complete closure that avoids the quadratic cost of the current OWL 2-to-fuzzyDL step.

Author Contributions: F.B.: Supervision, formal analysis, writing—review and editing. G.F.: Conceptualisation, methodology, formal analysis, visualization, writing—original draft, writing—review and editing. G.L.R.: Methodology, formal analysis, visualization, writing—original draft, writing—review and editing. U.S.: Supervision, formal analysis, writing—review and editing. M.E.T.: Supervision, project administration, funding acquisition, writing—review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: Giuseppe Filippone, Gianmarco La Rosa, and Marco Elio Tabacchi acknowledge financial support from the **Sustainability Decision Framework (SDF)** Research Project—CUP **B79J23000540005**—Grant Assignment Decree No. **5486** adopted on **4 August 2023** by the Italian Ministry of University and Research (MUR). Gianmarco La Rosa also acknowledges support from the **Gruppo Nazionale per le Strutture Algebriche, Geometriche e le loro Applicazioni (GNSAGA)** of the **Istituto Nazionale di Alta Matematica (INdAM)** Francesco Severi. Marco Elio Tabacchi also acknowledges financial support under the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, Call for tender No. 1409 published on 14 September 2022 by the Italian Ministry of University and Research (MUR), funded by the European Union—NextGenerationEU—Project Title **Quantum Models for Logic, Computation and Natural Processes (QM4NP)**—CUP **B53D23030160001**—Grant Assignment Decree No. **1371** adopted on **1 September 2023** by the Italian Ministry of University and Research (MUR). Fernando Bobillo and Umberto Straccia acknowledge support from the project **PID2024-159530OB-I00** funded by MICIU/AEI/10.13039/501100011033/FEDER, UE. Fernando Bobillo also acknowledges support from project **UZ2024-IyA-02** (funded by the University of Zaragoza).

Informed Consent Statement: Not applicable.

Data Availability Statement: The Python code, together with the scripts, random seeds, and results of the analyses added in this revision (sensitivity analysis, bounds for incomplete snapshots, slicing checks, revised Application 2, and cost measurements), is freely accessible from the GitHub repository at https://github.com/SDF-Unipa/sdf_reasoning_under_vagueness accessed on 14 September 2026.

Acknowledgments: We would like to express our gratitude to Vincenzo Reina, a student on the bachelor's degree course in Computer Science at the Department of Mathematics and Computer Science at the Università degli Studi di Palermo, for his contribution to the testing phase on the GitHub repository at the following address: https://github.com/renvins/fuzzy_dl_owl2_testing accessed on 14 September 2026. In particular, the student tested the process of translating from FDL code to the respective fuzzy ontology in Fuzzy OWL 2, and vice versa.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DL	Description Logic
MILP	Mixed-Integer Linear Programming
OWA	Ordered Weighted Averaging
SDG	Sustainable Development Goal
UNECE	United Nations Economic Commission for Europe
W3C	World Wide Web Consortium
OWL	Web Ontology Language
RDF	Resource Description Framework
IRI	Internationalised Resource Identifier
TBox, RBox, ABox	Terminological, role and assertional part of a Knowledge Base
OECD	Organisation for Economic Co-operation and Development
BNS	BasicNeedsStress
ES	EconomicStress
SDS	StructuralDevelopmentStress
BSES	BroadSocioEconomicStress
GTR	GreenTransitionReadiness
SEP	SustainableEconomicPerformance
JG	JoblessGrowth (analyst concept)
Rec.	Recovery (analyst concept)

References

1. W3C OWL Working Group. OWL 2 Web Ontology Language Document Overview (Second Edition). 2012. Available online: <https://www.w3.org/TR/2012/REC-owl2-overview-20121211/> (accessed on 14 September 2026).
2. Cuenca Grau, B.; Horrocks, I.; Motik, B.; Parsia, B.; Patel-Schneider, P.; Sattler, U. OWL 2: The next step for OWL. *J. Web Semant.* **2008**, *6*, 309–322. [CrossRef]
3. Baader, F.; McGuinness, D.L.; Nardi, D.; Patel-Schneider, P.F. (Eds.) *The Description Logic Handbook: Theory, Implementation and Applications*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2007. [CrossRef]
4. Laskey, K.; Laskey, K. Uncertainty Reasoning for the World Wide Web: Report on the URW3-XG Incubator Group. 2008. Available online: <https://www.w3.org/2005/Incubator/urw3/XGR-urw3-20080331> (accessed on 14 September 2026).
5. Zadeh, L.A. Fuzzy sets. *Inf. Control* **1965**, *8*, 338–353. [CrossRef]
6. Yen, J. Generalizing term subsumption languages to fuzzy logic. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI 1991)*; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 1991; Volume 1, pp. 472–477. Available online: <https://dl.acm.org/doi/10.5555/1631171.1631242> (accessed on 14 September 2026).
7. Straccia, U. *Foundations of Fuzzy Logic and Semantic Web Languages*, 1st ed.; Chapman and Hall/CRC: New York, NY, USA, 2014; p. 386. [CrossRef]

8. Borgwardt, S.; Peñaloza, R. Fuzzy Description Logics—A Survey. In *Proceedings of the 11th International Conference on Scalable Uncertainty Management (SUM 2017)*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 31–45. [\[CrossRef\]](#)
9. Cross, V.; Chen, S. Fuzzy Ontologies: State of the Art Revisited. In *Proceedings of the 37th Conference of the North American Fuzzy Information Processing Society (NAFIPS 2018)*; Communications in Computer and Information Science; Springer: Berlin/Heidelberg, Germany, 2018; Volume 831, pp. 230–242. [\[CrossRef\]](#)
10. Bobillo, F.; Straccia, U. Fuzzy ontology representation using OWL 2. *Int. J. Approx. Reason.* **2011**, *52*, 1073–1094. [\[CrossRef\]](#)
11. Bobillo, F.; Straccia, U. The fuzzy ontology reasoner fuzzyDL. *Knowl.-Based Syst.* **2016**, *95*, 12–34. [\[CrossRef\]](#)
12. Bobillo, F.; Filippone, G.; La Rosa, G.; Straccia, U.; Tabacchi, M.E. Fuzzy OWL 2 Reasoning: A Re-Engineered Python Framework. *arXiv* **2026**, arXiv:2606.19392. [\[CrossRef\]](#)
13. SDF Team. Fuzzy DL OWL 2. GitHub Repository. 2025. Available online: https://github.com/SDF-Unipa/fuzzy_dl_owl2 (accessed on 14 September 2026).
14. SDF Team. Fuzzy DL OWL 2. ReadTheDocs. 2025. Available online: <https://fuzzy-dl-owl2.readthedocs.io/en/latest> (accessed on 14 September 2026).
15. SDF Team. PyOWL2. PiPy. Available online: <https://pypi.org/project/pyowl2/> (accessed on 14 September 2026).
16. Bobillo, F.; Delgado, M.; Gómez-Romero, J. DeLorean: A reasoner for fuzzy OWL 2. *Expert Syst. Appl.* **2012**, *39*, 258–272. [\[CrossRef\]](#)
17. Stoilos, G.; Simou, N.; Stamou, G.; Kollias, S. Uncertainty and the Semantic Web. *IEEE Intell. Syst.* **2006**, *21*, 84–87. [\[CrossRef\]](#)
18. Stoilos, G.; Stamou, G.; Pan, J. Fuzzy extensions of OWL: Logical properties and reduction to Fuzzy Description Logics. *Int. J. Approx. Reason.* **2010**, *51*, 656–679. [\[CrossRef\]](#)
19. Stoilos, G.; Stamou, G.B. Extending Fuzzy Description Logics for the Semantic Web. In *Proceedings of the 3rd International Workshop on OWL: Experience and Directions (OWLED 2007)*, Innsbruck, Austria, 6–7 June 2007; CEUR Workshop Proceedings; Volume 258. Available online: https://www.researchgate.net/publication/221218549_Extending_Fuzzy_Description_Logics_for_the_Semantic_Web (accessed on 14 September 2026).
20. Gao, M.; Liu, C. Extending OWL by fuzzy description logic. In *Proceedings of the 17th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'05)*, Hong Kong, China, 14–16 November 2005; Volume 2005, pp. 562–567. [\[CrossRef\]](#)
21. Bobillo, F.; Straccia, U. Optimising Fuzzy Description Logic Reasoners with General Concept Inclusions Absorption. *Fuzzy Sets Syst.* **2016**, *292*, 98–129. [\[CrossRef\]](#)
22. Grabisch, M.; Marichal, J.L.; Mesiar, R.; Pap, E. Aggregation functions: Means. *Inf. Sci.* **2011**, *181*, 1–22. [\[CrossRef\]](#)
23. Marichal, J.L. On Sugeno integral as an aggregation function. *Fuzzy Sets Syst.* **2000**, *114*, 347–365. [\[CrossRef\]](#)
24. Protégé Team. Protégé. Available online: <https://protege.stanford.edu/> (accessed on 14 September 2026).
25. Musen, M.A. The Protégé project: A look back and a look forward. *AI Matters* **2015**, *1*, 4–12. [\[CrossRef\]](#) [\[PubMed\]](#)
26. Huitzil, I.; Bobillo, F.; Gomez-Romero, J.; Straccia, U. Fudge: Fuzzy Ontology Building with Consensuated Fuzzy Datatypes. *Fuzzy Sets Syst.* **2020**, *401*, 91–112. [\[CrossRef\]](#)
27. Cardillo, F.A.; Debole, F.; Straccia, U. PN-OWL: A Two Stage Algorithm to Learn Fuzzy Concept Inclusions from OWL Ontologies. *Fuzzy Sets Syst.* **2024**, *490*, 109048. [\[CrossRef\]](#)
28. Reina, V. fuzzy_dl_owl2_testing: Test Suite. GitHub Repository. Available online: https://github.com/renvins/fuzzy_dl_owl2_testing (accessed on 14 September 2026).
29. Gurobi Optimization, LLC. Gurobi Optimizer. Available online: <https://www.gurobi.com/product> (accessed on 14 September 2026).
30. Gurobi Optimization, LLC. Python Interface to Gurobi: Gurobipy v13.0.2. PiPy. Available online: <https://pypi.org/project/gurobipy/> (accessed on 14 September 2026).
31. Hall, J.; Galabova, I.; Gottwald, L.; Feldmeier, M. HiGHS—High Performance Software for Linear Optimization. Available online: <https://highs.dev/> (accessed on 14 September 2026).
32. Forrest, J.; Lougee-Heimer, R. CBC User Guide. COIN-OR. Documentation. Available online: <https://www.coin-or.org/Cbc/cbcuserguide.html> (accessed on 14 September 2026).
33. Free Software Foundation. GLPK (GNU Linear Programming Kit). Available online: <https://www.gnu.org/software/glpk/> (accessed on 14 September 2026).
34. Toffolo, T.A.M. Python Tools for Modeling and Solving Mixed-Integer Linear Programs (MIPs): mip v1.16rc0. PiPy. Available online: <https://pypi.org/project/mip/> (accessed on 14 September 2026).
35. Roy, J.S. PuLP is an LP Modeler Written in Python. PuLP Can Generate MPS or LP Files and Call GLPK, COIN CLP/CBC, CPLEX, and GUROBI to Solve Linear Problems: PuLP v3.2.1. PiPy. Available online: <https://pypi.org/project/PuLP/> (accessed on 14 September 2026).
36. Bumbulis, P. re2c—Regular Expressions to Code. Documentation. 1993. Available online: <https://re2c.org/> (accessed on 14 September 2026).

37. Paxson, V. Fast Lexical Analyzer Generator. Documentation. 1987. Available online: <https://github.com/westes/flex> (accessed on 14 September 2026).
38. Behnel, S.; Bradshaw, R.; Woods, D.; Valo, M.; Dalcín, L. Cython—C-Extensions for Python. Documentation. 2007. Available online: <https://cython.org/> (accessed on 14 September 2026).
39. Castronovo, L.; Filippone, G.; Galici, M.; La Rosa, G.; Pavone, A.M.; Tabacchi, M.E. SDF FuzzIA: A Fuzzy-Based, AI Support System for Decision-Making Frameworks. In *Decision Sciences*; Pavone, M., Coello Coello, C.A., Cerulli, R., Greco, S., Talbi, E.G., Eds.; Springer: Cham, Switzerland, 2026; pp. 395–405. [\[CrossRef\]](#)
40. Filippone, G.; La Rosa, G.; Tabacchi, M.E. SDF-FuzzIA: A Fuzzy-Ontology Based Plug-in for the Intelligent Analysis of Geo-Thematic Data. In *Scalable Uncertainty Management*; Destercke, S., Martinez, M.V., Sanfilippo, G., Eds.; Springer: Cham, Switzerland, 2025; pp. 163–169. [\[CrossRef\]](#)
41. Calegari, S.; Sanchez, E. Object-fuzzy concept network: An enrichment of ontologies in semantic information retrieval. *J. Am. Soc. Inf. Sci. Technol.* **2008**, *59*, 2171–2185. [\[CrossRef\]](#)
42. Knappe, R.; Bulskov, H.; Andreasen, T. Perspectives on ontology-based querying. *Int. J. Intell. Syst.* **2007**, *22*, 739–761. [\[CrossRef\]](#)
43. Simou, N.; Athanasiadis, T.; Stoilos, G.; Kollias, S. Image indexing and retrieval using expressive fuzzy description logics. *Signal Image Video Process.* **2008**, *2*, 321–335. [\[CrossRef\]](#)
44. da Costa, P.C.G.; Laskey, K.B.; Lukasiewicz, T. Uncertainty Representation and Reasoning in the Semantic Web. In *Semantic Web Engineering in the Knowledge Society*; Information Science Reference; IGI Global Scientific Publishing: Hershey, PA, USA, 2008; pp. 315–340. [\[CrossRef\]](#)
45. Sanchez, E. (Ed.) *Fuzzy Logic and the Semantic Web; Capturing Intelligence*; Elsevier: Amsterdam, The Netherlands, 2006; Volume 1. [\[CrossRef\]](#)
46. Tho, Q.; Hui, S.; Fong, A.; Cao, T.H. Automatic fuzzy ontology generation for semantic Web. *IEEE Trans. Knowl. Data Eng.* **2006**, *18*, 842–856. [\[CrossRef\]](#)
47. Baader, F.; Horrocks, I.; Lutz, C.; Sattler, U. *An Introduction to Description Logic*; Cambridge University Press: Cambridge, UK, 2017. [\[CrossRef\]](#)
48. Hájek, P. What does mathematical fuzzy logic offer to description logic? In *Fuzzy Logic and the Semantic Web*; Sanchez, E., Ed.; Capturing Intelligence; Elsevier: Amsterdam, The Netherlands, 2006; Chapter 5, pp. 91–100.
49. Atanassov, K.T. Intuitionistic fuzzy sets. *Fuzzy Sets Syst.* **1986**, *20*, 87–96. [\[CrossRef\]](#)
50. Zadeh, L.A. The concept of a linguistic variable and its application to approximate reasoning—I. *Inf. Sci.* **1975**, *8*, 199–249. [\[CrossRef\]](#)
51. Atanassov, K.T.; Gargov, G. Interval valued intuitionistic fuzzy sets. *Fuzzy Sets Syst.* **1989**, *31*, 343–349. [\[CrossRef\]](#)
52. Mendel, J.M.; John, R.I.B. Type-2 fuzzy sets made simple. *IEEE Trans. Fuzzy Syst.* **2002**, *10*, 117–127. [\[CrossRef\]](#)
53. Bustince, H.; Barrenechea, E.; Pagola, M.; Fernandez, J.; Xu, Z.; Bedregal, B.; Montero, J.; Hągras, H.; Herrera, F.; De Baets, B. A Historical Account of Types of Fuzzy Sets and Their Relationships. *IEEE Trans. Fuzzy Syst.* **2016**, *24*, 179–194. [\[CrossRef\]](#)
54. Dubois, D.; Prade, H. Gradualness, uncertainty and bipolarity: Making sense of fuzzy sets. *Fuzzy Sets Syst.* **2012**, *192*, 3–24. [\[CrossRef\]](#)
55. Dubois, D.; Prade, H. Interval-valued Fuzzy Sets, Possibility Theory and Imprecise Probability. In Proceedings of the Joint 4th Conference of the European Society for Fuzzy Logic and Technology and the 11th Rencontres Francophones sur la Logique Floue et ses Applications (EUSFLAT-LFA 2005), Barcelona, Spain, 7–9 September 2005; pp. 314–319.
56. Straccia, U. Description Logics over Lattices. *Int. J. Uncertain. Fuzziness -Knowl.-Based Syst.* **2006**, *14*, 1–16. [\[CrossRef\]](#)
57. Ruspini, E.H. A new approach to clustering. *Inf. Control* **1969**, *15*, 22–32. [\[CrossRef\]](#)
58. Yager, R.R. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Trans. Syst. Man Cybern.* **1988**, *18*, 183–190. [\[CrossRef\]](#)
59. Bobillo, F.; Straccia, U. Syntax and Semantics of FuzzyDL. Available online: <https://www.umbertostraccia.it/cs/software/fuzzyDL/fuzzyDL.html> (accessed on 14 September 2026).
60. Saisana, M.; Saltelli, A.; Tarantola, S. Uncertainty and Sensitivity Analysis Techniques as Tools for the Quality Assessment of Composite Indicators. *J. R. Stat. Soc. Ser. A Stat. Soc.* **2005**, *168*, 307–323. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.