



A Framework for Approximate Generalization in Quantitative Theories

Temur Kutsia^(✉)  and Cleo Pau

RISC, Johannes Kepler University Linz, Linz, Austria
{kutsia, ipau}@risc.jku.at

Abstract. Anti-unification aims at computing generalizations for given terms, retaining their common structure and abstracting differences by variables. We study quantitative anti-unification where the notion of the common structure is relaxed into “proximal” up to the given degree with respect to the given fuzzy proximity relation. Proximal symbols may have different names and arities. We develop a generic set of rules for computing minimal complete sets of approximate generalizations and study their properties. Depending on the characterizations of proximities between symbols and the desired forms of solutions, these rules give rise to different versions of concrete algorithms.

Keywords: Generalization · Anti-unification · Quantitative theories · Fuzzy proximity relations

1 Introduction

Generalization problems play an important role in various areas of mathematics, computer science, and artificial intelligence. Anti-unification [12, 14] is a logic-based method for computing generalizations. Being originally used for inductive and analogical reasoning, some recent applications include recursion scheme detection in functional programs [4], programming by examples in domain-specific languages [13], learning bug-fixing from software code repositories [3, 15], automatic program repair [7], preventing bugs and misconfiguration in services [11], linguistic structure learning for chatbots [6], to name just a few.

In most of the existing theories where anti-unification is studied, the background knowledge is assumed to be precise. Therefore, those techniques are not suitable for reasoning with incomplete, imprecise information (which is very common in real-world communication), where the exact equality is replaced by its (quantitative) approximation. Fuzzy proximity and similarity relations are notable examples of such extensions. These kinds of quantitative theories have many useful applications, some most recent ones being related to artificial intelligence, program verification, probabilistic programming, or natural language processing. Many tasks arising in these areas require reasoning methods and computational tools that deal with quantitative information. For instance, approximate

inductive reasoning, reasoning and programming by analogy, similarity detection in programming language statements or in natural language texts could benefit from solving approximate generalization constraints, which is a theoretically interesting and challenging task. Investigations in this direction have been started only recently. In [1], the authors proposed an anti-unification algorithm for fuzzy similarity (reflexive, symmetric, min-transitive) relations, where mismatches are allowed not only in symbol names, but also in their arities (fully fuzzy signatures). The algorithm from [9] is designed for fuzzy proximity (i.e., reflexive and symmetric) relations with mismatches only in symbol names.

In this paper, we study approximate anti-unification from a more general perspective. The considered relations are fuzzy proximity relations. Proximal symbols may have different names and arities. We consider four different variants of relating arguments between different proximal symbols: unrestricted relations/functions, and correspondence (i.e. left- and right-total) relations/functions. A generic set of rules for computing minimal complete sets of generalizations is introduced and its termination, soundness and completeness properties are proved. From these rules, we obtain concrete algorithms that deal with different kinds of argument relations. We also show how the existing approximate anti-unification algorithms and their generalizations fit into this framework.

Organization: In Sect. 2 we introduce the notation and definitions. Section 3 is devoted to a technical notion of term set consistency and to an algorithm for computing elements of consistent sets of terms. It is used later in the main set of anti-unification rules, which are introduced and characterized in Sect. 4. The concrete algorithms obtained from those rules are also described in this section. In Sect. 5, we discuss complexity. Section 6 offers a high-level picture of the studied problems and concludes.

An extended version of this work can be found in the technical report [8].

2 Preliminaries

Proximity Relations. Given a set S , a mapping \mathcal{R} from $S \times S$ to the real interval $[0, 1]$ is called a binary *fuzzy relation* on S . By fixing a number λ , $0 \leq \lambda \leq 1$, we can define the crisp (i.e., two-valued) counterpart of \mathcal{R} , named the λ -cut of \mathcal{R} , as $\mathcal{R}_\lambda := \{(s_1, s_2) \mid \mathcal{R}(s_1, s_2) \geq \lambda\}$. A fuzzy relation \mathcal{R} on a set S is called a *proximity relation* if it is reflexive ($\mathcal{R}(s, s) = 1$ for all $s \in S$) and symmetric ($\mathcal{R}(s_1, s_2) = \mathcal{R}(s_2, s_1)$ for all $s_1, s_2 \in S$). A T-norm \wedge is an associative, commutative, non-decreasing binary operation on $[0, 1]$ with 1 as the unit element. We take minimum in the role of T-norm.

Terms and Substitutions. We consider a first-order alphabet consisting of a set of fixed arity function symbols \mathcal{F} and a set of variables \mathcal{V} , which includes a special symbol $_$ (the anonymous variable). The set of *named* (i.e., non-anonymous) variables $\mathcal{V} \setminus \{_ \}$ is denoted by \mathcal{V}^N . When the set of variables is not explicitly

specified, we mean \mathcal{V} . The set of terms $\mathcal{T}(\mathcal{F}, \mathcal{V})$ over \mathcal{F} and \mathcal{V} is defined in the standard way: $t \in \mathcal{T}(\mathcal{F}, \mathcal{V})$ iff t is defined by the grammar $t := x \mid f(t_1, \dots, t_n)$, where $x \in \mathcal{V}$ and $f \in \mathcal{F}$ is an n -ary symbol with $n \geq 0$. Terms over $\mathcal{T}(\mathcal{F}, \mathcal{V}^{\mathbb{N}})$ are defined similarly except that all variables are taken from $\mathcal{V}^{\mathbb{N}}$.

We denote arbitrary function symbols by f, g, h , constants by a, b, c , variables by x, y, z, v , and terms by s, t, r . The *head* of a term is defined as $\text{head}(x) := x$ and $\text{head}(f(t_1, \dots, t_n)) := f$. For a term t , we denote with $\mathcal{V}(t)$ (resp. by $\mathcal{V}^{\mathbb{N}}(t)$) the set of all variables (resp. all named variables) appearing in t . A term is called *linear* if no named variable occurs in it more than once.

The deanonymization operation *deanon* replaces each occurrence of the anonymous variable in a term by a fresh variable. For instance, $\text{deanon}(f(_, x, g(_))) = f(y', x, g(y''))$, where y' and y'' are fresh. Hence, $\text{deanon}(t) \in \mathcal{T}(\mathcal{F}, \mathcal{V}^{\mathbb{N}})$ is unique up to variable renaming for all $t \in \mathcal{T}(\mathcal{F}, \mathcal{V})$. $\text{deanon}(t)$ is linear iff t is linear.

The notions of *term depth*, *term size* and a *position in a term* are defined in the standard way, see, e.g. [2]. By $t|_p$ we denote the subterm of t at position p and by $t[s]_p$ a term that is obtained from t by replacing the subterm at position p by the term s .

A *substitution* is a mapping from $\mathcal{V}^{\mathbb{N}}$ to $\mathcal{T}(\mathcal{F}, \mathcal{V}^{\mathbb{N}})$ (i.e., without anonymous variables), which is the identity almost everywhere. We use the Greek letters $\sigma, \vartheta, \varphi$ to denote substitutions, except for the identity substitution which is written as *Id*. We represent substitutions with the usual set notation. *Application* of a substitution σ to a term t , denoted by $t\sigma$, is defined as $_ \sigma := _$, $x\sigma := \sigma(x)$, $f(t_1, \dots, t_n)\sigma := f(t_1\sigma, \dots, t_n\sigma)$. Substitution *composition* is defined as a composition of mappings. We write $\sigma\vartheta$ for the composition of σ with ϑ .

Argument Relations and Mappings. Given two sets $N = \{1, \dots, n\}$ and $M = \{1, \dots, m\}$, a binary *argument relation* over $N \times M$ is a (possibly empty) subset of $N \times M$. We denote argument relations by ρ . An argument relation $\rho \subseteq N \times M$ is (i) *left-total* if for all $i \in N$ there exists $j \in M$ such that $(i, j) \in \rho$; (ii) *right-total* if for all $j \in M$ there exists $i \in N$ such that $(i, j) \in \rho$. *Correspondence relations* are those that are both left- and right-total.

An *argument mapping* is an argument relation that is a partial injective function. In other words, an argument mapping π from $N = \{1, \dots, n\}$ to $M = \{1, \dots, m\}$ is a function $\pi : I_n \mapsto I_m$, where $I_n \subseteq N$, $I_m \subseteq M$ and $|I_n| = |I_m|$. Note that it can be also the empty mapping: $\pi : \emptyset \mapsto \emptyset$. The inverse of an argument mapping is again an argument mapping.

Given a proximity relation \mathcal{R} over \mathcal{F} , we assume that for each pair of function symbols f and g with $\mathcal{R}(f, g) = \alpha > 0$, where f is n -ary and g is m -ary, there is also given an argument relation ρ over $\{1, \dots, n\} \times \{1, \dots, m\}$. We use the notation $f \sim_{\mathcal{R}, \alpha}^{\rho} g$. These argument relations should satisfy the following conditions: ρ is the empty relation if f or g is a constant; ρ is the identity if $f = g$; $f \sim_{\mathcal{R}, \alpha}^{\rho} g$ iff $g \sim_{\mathcal{R}, \alpha}^{\rho^{-1}} f$, where ρ^{-1} is the inverse of ρ .

Example 1. Assume that we have four different versions of defining the notion of author (e.g., originated from four different knowledge bases) $author_1$ (*first-name, middle-initial, last-name*), $author_2$ (*first-name, last-name*), $author_3$ (*last-name, first-name, middle-initial*), and $author_4$ (*full-name*). One could define the argument relations/mappings between these function symbols e.g., as follows:

$$\begin{aligned} author_1 &\sim_{\mathcal{R},0.7}^{\{(1,1),(3,2)\}} author_2, & author_1 &\sim_{\mathcal{R},0.9}^{\{(3,1),(1,2),(2,3)\}} author_3, \\ author_1 &\sim_{\mathcal{R},0.5}^{\{(1,1),(3,1)\}} author_4, & author_2 &\sim_{\mathcal{R},0.7}^{\{(1,2),(2,1)\}} author_3, \\ author_2 &\sim_{\mathcal{R},0.5}^{\{(1,1),(2,1)\}} author_4, & author_3 &\sim_{\mathcal{R},0.5}^{\{(1,1),(2,1)\}} author_4. \end{aligned}$$

Proximity Relations over Terms. Each proximity relation \mathcal{R} in this paper is defined on $\mathcal{F} \cup \mathcal{V}$ such that $\mathcal{R}(f, x) = 0$ for all $f \in \mathcal{F}$ and $x \in \mathcal{V}$, and $\mathcal{R}(x, y) = 0$ for all $x \neq y, x, y \in \mathcal{V}$. We assume that \mathcal{R} is *strict*: for all $w_1, w_2 \in \mathcal{F} \cup \mathcal{V}$, if $\mathcal{R}(w_1, w_2) = 1$, then $w_1 = w_2$. Yet another assumption is that for each $f \in \mathcal{F}$, its (\mathcal{R}, λ) -proximity class $\{g \mid \mathcal{R}(f, g) \geq \lambda\}$ is *finite* for any \mathcal{R} and λ .

We extend such an \mathcal{R} to terms from $\mathcal{T}(\mathcal{F}, \mathcal{V})$ as follows:

- (a) $\mathcal{R}(t, s) := 0$ if $\mathcal{R}(\text{head}(s), \text{head}(t)) = 0$;
- (b) $\mathcal{R}(t, s) := 1$ if $t = s$ and $t, s \in \mathcal{V}$;
- (c) $\mathcal{R}(t, s) := \mathcal{R}(f, g) \wedge \mathcal{R}(t_{i_1}, s_{j_1}) \wedge \dots \wedge \mathcal{R}(t_{i_k}, s_{j_k})$, if $t = f(t_1, \dots, t_n), s = g(s_1, \dots, s_m), f \sim_{\mathcal{R}, \lambda}^\rho g$, and $\rho = \{(i_1, j_1), \dots, (i_k, j_k)\}$.

If $\mathcal{R}(t, s) \geq \lambda$, we write $t \simeq_{\mathcal{R}, \lambda} s$. When $\lambda = 1$, the relation $\simeq_{\mathcal{R}, \lambda}$ does not depend on \mathcal{R} due to strictness of the latter and is just the syntactic equality $=$.

The (\mathcal{R}, λ) -proximity class of a term t is $\mathbf{pc}_{\mathcal{R}, \lambda}(t) := \{s \mid s \simeq_{\mathcal{R}, \lambda} t\}$.

Generalizations. Given \mathcal{R} and λ , a term r is an (\mathcal{R}, λ) -generalization of (alternatively, (\mathcal{R}, λ) -more general than) a term t , written as $r \lesssim_{\mathcal{R}, \lambda} t$, if there exists a substitution σ such that $\text{deanon}(r)\sigma \simeq_{\mathcal{R}, \lambda} \text{deanon}(t)$. The strict part of $\lesssim_{\mathcal{R}, \lambda}$ is denoted by $\prec_{\mathcal{R}, \lambda}$, i.e., $r \prec_{\mathcal{R}, \lambda} t$ if $r \lesssim_{\mathcal{R}, \lambda} t$ and not $t \lesssim_{\mathcal{R}, \lambda} r$.

Example 2. Given a proximity relation \mathcal{R} , a cut value λ , constants $a \sim_{\mathcal{R}, \alpha_1}^\emptyset b$ and $b \sim_{\mathcal{R}, \alpha_2}^\emptyset c$, binary function symbols f and h , and a unary function symbol g such that $h \sim_{\mathcal{R}, \alpha_3}^{\{(1,1),(1,2)\}} f$ and $h \sim_{\mathcal{R}, \alpha_4}^{\{(1,1)\}} g$ with $\alpha_i \geq \lambda, 1 \leq i \leq 4$, we have

- $h(x, _) \lesssim_{\mathcal{R}, \lambda} h(a, x)$, because $h(x, x')\{x \mapsto a, x' \mapsto x\} = h(a, x) \simeq_{\mathcal{R}, \lambda} h(a, x)$.
- $h(x, _) \lesssim_{\mathcal{R}, \lambda} h(_, x)$, because $h(x, x')\{x \mapsto y', x' \mapsto x\} = h(y', x) \simeq_{\mathcal{R}, \lambda} h(y', x)$.
- $h(x, x) \not\lesssim_{\mathcal{R}, \lambda} h(_, x)$, because $h(x, x) \not\lesssim_{\mathcal{R}, \lambda} h(y', x)$.
- $h(x, _) \lesssim_{\mathcal{R}, \lambda} f(a, c)$, because $h(x, x')\{x \mapsto b\} = h(b, x') \simeq_{\mathcal{R}, \lambda} f(a, c)$.
- $h(x, _) \lesssim_{\mathcal{R}, \lambda} g(c)$, because $h(x, x')\{x \mapsto c\} = h(c, x') \simeq_{\mathcal{R}, \lambda} g(c)$.

The notion of *syntactic generalization* of a term is a special case of (\mathcal{R}, λ) -generalization for $\lambda = 1$. We write $r \lesssim t$ to indicate that r is a syntactic generalization of t . Its strict part is denoted by \prec .

Since \mathcal{R} is strict, $r \lesssim t$ is equivalent to $\text{deanon}(r)\sigma = \text{deanon}(t)$ for some σ (note the syntactic equality here).

Theorem 1. *If $r \lesssim t$ and $t \lesssim_{\mathcal{R},\lambda} s$, then $r \lesssim_{\mathcal{R},\lambda} s$.*

Proof. $r \lesssim t$ implies $\text{deanon}(r)\sigma = \text{deanon}(t)$ for some σ , while from $t \lesssim_{\mathcal{R},\lambda} s$ we have $\text{deanon}(t)\vartheta \simeq_{\mathcal{R},\lambda} \text{deanon}(s)$ for some ϑ . Then $\text{deanon}(r)\sigma\vartheta \simeq_{\mathcal{R},\lambda} \text{deanon}(s)$, which implies $r \lesssim_{\mathcal{R},\lambda} s$. □

Note that $r \lesssim_{\mathcal{R},\lambda} t$ and $t \lesssim_{\mathcal{R},\lambda} s$, in general, do not imply $r \lesssim_{\mathcal{R},\lambda} s$ due to non-transitivity of $\simeq_{\mathcal{R},\lambda}$.

Definition 1 (Minimal complete set of (\mathcal{R}, λ) -generalizations). *Given $\mathcal{R}, \lambda, t_1$, and t_2 , a set of terms T is a complete set of (\mathcal{R}, λ) -generalizations of t_1 and t_2 if*

- (a) every $r \in T$ is an (\mathcal{R}, λ) -generalization of t_1 and t_2 ,
- (b) if r' is an (\mathcal{R}, λ) -generalization of t_1 and t_2 , then there exists $r \in T$ such that $r' \lesssim r$ (note that we use syntactic generalization here).

In addition, T is minimal, if it satisfies the following property:

- (c) if $r, r' \in T, r \neq r'$, then neither $r \prec_{\mathcal{R},\lambda} r'$ nor $r' \prec_{\mathcal{R},\lambda} r$.

A minimal complete set of (\mathcal{R}, λ) -generalizations ((\mathcal{R}, λ) -mcs g) of two terms is unique modulo variable renaming. The elements of the (\mathcal{R}, λ) -mcs g of t_1 and t_2 are called least general (\mathcal{R}, λ) -generalizations ((\mathcal{R}, λ) -lg gs) of t_1 and t_2 .

This definition directly extends to generalizations of finitely many terms.

The problem of computing an (\mathcal{R}, λ) -generalization of terms t and s is called the (\mathcal{R}, λ) -anti-unification problem of t and s . In anti-unification, the goal is to compute their least general (\mathcal{R}, λ) -generalization.

The precise formulation of the anti-unification problem would be the following: Given $\mathcal{R}, \lambda, t_1, t_2$, find an (\mathcal{R}, λ) -lg g r of t_1 and t_2 , substitutions σ_1, σ_2 , and the approximation degrees α_1, α_2 such that $\mathcal{R}(r\sigma_1, t_1) = \alpha_1$ and $\mathcal{R}(r\sigma_2, t_2) = \alpha_2$. A minimal complete algorithm to solve this problem would compute exactly the elements of (\mathcal{R}, λ) -mcs g of t_1 and t_2 together with their approximation degrees. However, as we see below, it is problematic to solve the problem in this form. Therefore, we will consider a slightly modified variant, taking into account anonymous variables in generalizations and relaxing bounds on their degrees.

We assume that the terms to be generalized are ground. It is not a restriction because we can treat variables as constants that are close only to themselves.

Recall that the proximity class of any alphabet symbol is finite. Also, the symbols are related to each other by finitely many argument relations. One may think that it leads to finite proximity classes of terms, but this is not the case. Consider, e.g., \mathcal{R} and λ , where $h \simeq_{\mathcal{R},\lambda}^{\{(1,1)\}} f$ with binary h and unary f . Then the (\mathcal{R}, λ) -proximity class of $f(a)$ is infinite: $\{f(a)\} \cup \{h(a, t) \mid t \in \mathcal{T}(\mathcal{F}, \mathcal{V})\}$. Also, the (\mathcal{R}, λ) -mcs g for $f(a)$ and $f(b)$ is infinite: $\{f(x)\} \cup \{h(x, t) \mid t \in \mathcal{T}(\mathcal{F}, \emptyset)\}$.

Definition 2. *Given the terms $t_1, \dots, t_n, n \geq 1$, a position p in a term r is called irrelevant for (\mathcal{R}, λ) -generalizing (resp. for (\mathcal{R}, λ) -proximity to) t_1, \dots, t_n if $r[s]_p \lesssim_{\mathcal{R},\lambda} t_i$ (resp. $r[s]_p \simeq_{\mathcal{R},\lambda} t_i$) for all $1 \leq i \leq n$ and for all terms s .*

We say that r is a relevant (\mathcal{R}, λ) -generalization (resp. relevant (\mathcal{R}, λ) -proximal term) of t_1, \dots, t_n if $r \lesssim_{\mathcal{R}, \lambda} t_i$ (resp. $r \simeq_{\mathcal{R}, \lambda} t_i$) for all $1 \leq i \leq n$ and $r|_p = _$ for all positions p in r that is irrelevant for generalizing (resp. for proximity to) t_1, \dots, t_n . The (\mathcal{R}, λ) -relevant proximity class of t is

$$\mathbf{rpc}_{\mathcal{R}, \lambda}(t) := \{s \mid s \text{ is a relevant } (\mathcal{R}, \lambda)\text{-proximal term of } t\}.$$

In the example above, position 2 in $h(x, t)$ is irrelevant for generalizing $f(a)$ and $f(b)$, and $h(x, _)$ is one of their relevant generalizations. Note that $f(x)$ is also a relevant generalization of $f(a)$ and $f(b)$, since it contains no irrelevant positions. More general generalizations like, e.g., x , are relevant as well. Similarly, position 2 in $h(a, t)$ is irrelevant for proximity to $f(a)$ and $\mathbf{rpc}_{\mathcal{R}, \lambda}(f(a)) = \{f(a), h(a, _)\}$. Generally, $\mathbf{rpc}_{\mathcal{R}, \lambda}(t)$ is finite for any t due to the finiteness of proximity classes of symbols and argument relations mentioned above.

Definition 3 (Minimal complete set of relevant (\mathcal{R}, λ) -generalizations). Given \mathcal{R} , λ , t_1 , and t_2 , a set of terms T is a complete set of relevant (\mathcal{R}, λ) -generalizations of t_1 and t_2 if

- (a) every element of T is a relevant (\mathcal{R}, λ) -generalization of t_1 and t_2 , and
- (b) if r is a relevant (\mathcal{R}, λ) -generalization of t_1 and t_2 , then there exists $r' \in T$ such that $r \lesssim r'$.

The minimality property is defined as in Definition 1.

This definition directly extends to relevant generalizations of finitely many terms. We use (\mathcal{R}, λ) -mcsrg as an abbreviation for minimal complete set of relevant (\mathcal{R}, λ) -generalization. Like relevant proximity classes, mcsrg's are also finite.

Lemma 1. For given \mathcal{R} and λ , if all argument relations are correspondence relations, then (\mathcal{R}, λ) -mcsrg's and (\mathcal{R}, λ) -proximity classes for all terms are finite.

Proof. Under correspondence relations no term contains an irrelevant position for generalization or for proximity. □

Hence, for correspondence relations the notions of mcsrg and mcsrg coincide, as well as the notions of proximity class and relevant proximity class.

For a term r , we define its *linearized version* $\text{lin}(r)$ as a term obtained from r by replacing each occurrence of a named variable in r by a fresh one. For instance, $\text{lin}(f(x, _, g(y, x, a), b)) = f(x', _, g(y', x'', a), b)$, where x', x'', y' are fresh variables. Linearized versions of terms are unique modulo variable renaming.

Definition 4 (Generalization degree upper bound). Given two terms r and t , a proximity relation \mathcal{R} , and a λ -cut, the (\mathcal{R}, λ) -generalization degree upper bound of r and t , denoted by $\text{gdub}_{\mathcal{R}, \lambda}(r, t)$, is defined as follows:

Let $\alpha := \max\{\mathcal{R}(\text{lin}(r)\sigma, t) \mid \sigma \text{ is a substitution}\}$. Then $\text{gdub}_{\mathcal{R}, \lambda}(r, t)$ is α if $\alpha \geq \lambda$, and 0 otherwise.

Intuitively, $\text{gdub}_{\mathcal{R},\lambda}(r, t) = \alpha$ means that no instance of r can get closer than α to t in \mathcal{R} . From the definition it follows that if $r \lesssim_{\mathcal{R},\lambda} t$, then $0 < \lambda \leq \text{gdub}_{\mathcal{R},\lambda}(r, t) \leq 1$ and if $r \not\lesssim_{\mathcal{R},\lambda} t$, then $\text{gdub}_{\mathcal{R},\lambda}(r, t) = 0$.

The upper bound computed by gdub is more relaxed than it would be if the linearization function were not used, but this is what we will be able to compute in our algorithms later.

Example 3. Let $\mathcal{R}(a, b) = 0.6$, $\mathcal{R}(b, c) = 0.7$, and $\lambda = 0.5$. Then $\text{gdub}_{\mathcal{R},\lambda}(f(x, b), f(a, c)) = 0.7$ and $\text{gdub}_{\mathcal{R},\lambda}(f(x, x), f(a, c)) = \text{gdub}_{\mathcal{R},\lambda}(f(x, y), f(a, c)) = 1$.

It is not difficult to see that if $r\sigma \simeq_{\mathcal{R},\lambda} t$, then $\mathcal{R}(r\sigma, t) \leq \text{gdub}_{\mathcal{R},\lambda}(r, t)$. In Example 3, for $\sigma = \{x \mapsto b\}$ we have $\mathcal{R}(f(x, x)\sigma, f(a, c)) = \mathcal{R}(f(b, b), f(a, c)) = 0.6 < \text{gdub}_{\mathcal{R},\lambda}(f(x, x), f(a, c)) = 1$.

We compute $\text{gdub}_{\mathcal{R},\lambda}(r, t)$ as follows: If r is a variable, then $\text{gdub}_{\mathcal{R},\lambda}(r, t) = 1$. Otherwise, if $\text{head}(r) \sim_{\mathcal{R},\beta}^{\rho} \text{head}(t)$, then $\text{gdub}_{\mathcal{R},\lambda}(r, t) = \beta \wedge \bigwedge_{(i,j) \in \rho} \text{gdub}_{\mathcal{R},\lambda}(r|_i, t|_j)$. Otherwise, $\text{gdub}_{\mathcal{R},\lambda}(r, t) = 0$.

3 Term Set Consistency

The notion of term set consistency plays an important role in the computation of proximal generalizations. Intuitively, a set of terms is (\mathcal{R}, λ) -consistent if all the terms in the set have a common (\mathcal{R}, λ) -proximal term. In this section, we discuss this notion and the corresponding algorithms.

Definition 5 (Consistent set of terms). A finite set of terms T is (\mathcal{R}, λ) -consistent if there exists a term s such that $s \simeq_{\mathcal{R},\lambda} t$ for all $t \in T$.

(\mathcal{R}, λ) -consistency of a finite term set T is equivalent to $\bigcap_{t \in T} \mathbf{pc}_{\mathcal{R},\lambda}(t) \neq \emptyset$, but we cannot use this property to decide consistency, since proximity classes of terms can be infinite (when the argument relations are not restricted). For this reason, we introduce the operation \sqcap on terms as follows: (i) $t \sqcap _ = _ \sqcap t = t$, (ii) $f(t_1, \dots, t_n) \sqcap f(s_1, \dots, s_n) = f(t_1 \sqcap s_1, \dots, t_n \sqcap s_n)$, $n \geq 0$. Obviously, \sqcap is associative (A), commutative (C), idempotent (I), and has $_$ as its unit element (U). It can be extended to sets of terms: $T_1 \sqcap T_2 := \{t_1 \sqcap t_2 \mid t_1 \in T_1, t_2 \in T_2\}$. It is easy to see that \sqcap on sets also satisfies the ACIU properties with the set $\{_ \}$ playing the role of the unit element.

Lemma 2. A finite set of terms T is (\mathcal{R}, λ) -consistent iff $\bigcap_{t \in T} \mathbf{rpc}_{\mathcal{R},\lambda}(t) \neq \emptyset$.

Proof. (\Rightarrow) If $s \simeq_{\mathcal{R},\lambda} t$ for all $t \in T$, then $s_t \in \mathbf{rpc}_{\mathcal{R},\lambda}(t)$, where s_t is obtained from s by replacing all subterms that are irrelevant for its (\mathcal{R}, λ) -proximity to t by $_$. Assume $T = \{t_1, \dots, t_n\}$. Then $s_{t_1} \sqcap \dots \sqcap s_{t_n} \in \bigcap_{t \in T} \mathbf{rpc}_{\mathcal{R},\lambda}(t)$.

(\Leftarrow) Obvious, since $s \simeq_{\mathcal{R},\lambda} t$ for $s \in \bigcap_{t \in T} \mathbf{rpc}_{\mathcal{R},\lambda}(t)$ and for all $t \in T$. □

Now we design an algorithm \mathcal{C} that computes $\bigcap_{t \in T} \mathbf{rpc}_{\mathcal{R},\lambda}(t)$ without actually computing $\mathbf{rpc}_{\mathcal{R},\lambda}(t)$ for each $t \in T$. A special version of the algorithm can be used to decide the (\mathcal{R}, λ) -consistency of T .

The algorithm is rule-based. The rules work on states, that are pairs $\mathbf{I}; s$, where s is a term and \mathbf{I} is a finite set of expressions of the form x in T , where T is a finite set of terms. \mathcal{R} and λ are given. There are two rules (\uplus stands for disjoint union):

Rem: Removing the empty set

$$\{x \text{ in } \emptyset\} \uplus \mathbf{I}; s \Longrightarrow \mathbf{I}; s\{x \mapsto _ \}.$$

Red: Reduce a set to new sets

$\{x \text{ in } \{t_1, \dots, t_m\}\} \uplus \mathbf{I}; s \Longrightarrow \{y_1 \text{ in } T_1, \dots, y_n \text{ in } T_n\} \cup \mathbf{I}; s\{x \mapsto h(y_1, \dots, y_n)\}$, where $m \geq 1$, h is an n -ary function symbol such that $h \sim_{\mathcal{R}, \gamma_k}^{\rho_k} \text{head}(t_k)$ with $\gamma_k \geq \lambda$ for all $1 \leq k \leq m$, and $T_i := \{t_k|_j \mid (i, j) \in \rho_k, 1 \leq k \leq m\}$, $1 \leq i \leq n$, is the set of all those arguments of the terms t_1, \dots, t_m that are supposed to be (\mathcal{R}, λ) -proximal to the i 's argument of h .

To compute $\prod_{t \in T} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$, \mathfrak{C} starts with $\{x \text{ in } T\}; x$ and applies the rules as long as possible. **Red** causes branching. A state of the form $\emptyset; s$ is called a success state. A failure state has the form $\mathbf{I}; s$, to which no rule applies and $\mathbf{I} \neq \emptyset$. In the full derivation tree, each leaf is a either success or a failure state.

Example 4. Assume a, b, c are constants, g, f, h are function symbols with the arities respectively 1, 2, and 3. Let λ be given and \mathcal{R} be defined so that $\mathcal{R}(a, b) \geq \lambda$, $\mathcal{R}(b, c) \geq \lambda$, $h \sim_{\mathcal{R}, \beta}^{\{(1,1), (1,2)\}}$ f , $h \sim_{\mathcal{R}, \gamma}^{\{(2,1)\}}$ g with $\beta \geq \lambda$ and $\gamma \geq \lambda$. Then

$$\begin{aligned} \mathbf{rpc}_{\mathcal{R}, \lambda}(f(a, c)) &= \{f(a, c), f(b, c), f(a, b), f(b, b), h(b, _, _)\}, \\ \mathbf{rpc}_{\mathcal{R}, \lambda}(g(a)) &= \{g(a), g(b), h(_, a, _), h(_, b, _)\}, \end{aligned}$$

and $\mathbf{rpc}_{\mathcal{R}, \lambda}(f(a, c)) \sqcap \mathbf{rpc}_{\mathcal{R}, \lambda}(g(a)) = \{h(b, a, _), h(b, b, _)\}$. We show how to compute this set with \mathfrak{C} : $\{x \text{ in } \{f(a, c), g(a)\}\}; x \Longrightarrow_{\text{Red}} \{y_1 \text{ in } \{a, c\}, y_2 : \{a\}, y_3 \text{ in } \emptyset\}; h(y_1, y_2, y_3) \Longrightarrow_{\text{Rem}} \{y_1 \text{ in } \{a, c\}, y_2 : \{a\}\}; h(y_1, y_2, _) \Longrightarrow_{\text{Red}} \{y_2 \text{ in } \{a\}\}; h(b, y_2, _)$. Here we have two ways to apply **Red** to the last state, leading to two elements of $\mathbf{rpc}_{\mathcal{R}, \lambda}(f(a, c)) \sqcap \mathbf{rpc}_{\mathcal{R}, \lambda}(g(a))$: $h(b, a, _)$ and $h(b, b, _)$.

Theorem 2. *Given a finite set of terms T , the algorithm \mathfrak{C} always terminates starting from the state $\{x \text{ in } T\}; x$ (where x is a fresh variable). If S is the set of success states produced at the end, we have $\{s \mid \emptyset; s \in S\} = \prod_{t \in T} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$.*

Proof. Termination: Associate to each state $\{x_1 \text{ in } T_1, \dots, x_n \text{ in } T_n\}; s$ the multi-set $\{d_1, \dots, d_n\}$, where d_i is the maximum depth of terms occurring in T_i . $d_i = 0$ if $T_i = \emptyset$. Compare these multisets by the Dershowitz-Manna ordering [5]. Each rule strictly reduces them, which implies termination.

By the definitions of $\mathbf{rpc}_{\mathcal{R}, \lambda}$ and \sqcap , $h(s_1, \dots, s_n) \in \prod_{t \in \{t_1, \dots, t_m\}} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$ iff $h \sim_{\mathcal{R}, \gamma_k}^{\rho_k} \text{head}(t_k)$ with $\gamma_k \geq \lambda$ for all $1 \leq k \leq m$ and $s_i \in \prod_{t \in T_i} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$, where $T_i = \{t_k|_j \mid (i, j) \in \rho_k, 1 \leq k \leq m\}$, $1 \leq i \leq n$. Therefore, in the **Rem** rule, the instance of x (which is $h(y_1, \dots, y_n)$) is in $\prod_{t \in \{t_1, \dots, t_m\}} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$ iff for each $1 \leq i \leq n$ we can find an instance of y_i in $\prod_{t \in T_i} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$. If T_i is empty, it means that the i 's argument of h is irrelevant for terms in $\{t_1, \dots, t_m\}$ and can be replaced by $_$. (**Rem** does it in a subsequent step.) Hence, in each success branch of the derivation tree, the algorithm \mathfrak{C} computes one element of $\prod_{t \in T} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$. Branching at **Red** helps produce all elements of $\prod_{t \in T} \mathbf{rpc}_{\mathcal{R}, \lambda}(t)$. \square

It is easy to see how to use \mathfrak{C} to decide the (\mathcal{R}, λ) -consistency of T : it is enough to find one successful branch in the \mathfrak{C} -derivation tree for $\{x \text{ in } T\}; x$. If there is no such branch, then T is not (\mathcal{R}, λ) -consistent. In fact, during the derivation we can even ignore the second component of the states.

4 Solving Generalization Problems

Now we can reformulate the anti-unification problem that will be solved in the remaining part of the paper. \mathcal{R} is a proximity relation and λ is a cut value.

Given: \mathcal{R} , λ , and the ground terms $t_1, \dots, t_n, n \geq 2$.

Find: a set S of tuples $(r, \sigma_1, \dots, \sigma_n, \alpha_1, \dots, \alpha_n)$ such that

- $\{r \mid (r, \dots) \in S\}$ is an (\mathcal{R}, λ) -mcsrg of t_1, \dots, t_n ,
- $r\sigma_i \simeq_{\mathcal{R}, \lambda} t_i$ and $\alpha_i = \text{gdub}_{\mathcal{R}, \lambda}(r, t_i), 1 \leq i \leq n$, for each $(r, \sigma_1, \dots, \sigma_n, \alpha_1, \dots, \alpha_n) \in S$.

(When $n = 1$, this is a problem of computing a relevant proximity class of a term.) Below we give a set of rules, from which one can obtain algorithms to solve the anti-unification problem for four versions of argument relations:

1. The most general (unrestricted) case; see algorithm \mathfrak{A}_1 below, the computed set of generalizations is an mcsrg;
2. Correspondence relations: using the same algorithm \mathfrak{A}_1 , the computed set of generalizations is an mcsrg;
3. Mappings: using a dedicated algorithm \mathfrak{A}_2 , the computed set of generalizations is an mcsrg;
4. Correspondence mappings (bijections): using the same algorithm \mathfrak{A}_2 , the computed set of generalizations is an mcsrg.

Each of them has also the corresponding linear variant, computing minimal complete sets of (relevant) linear (\mathcal{R}, λ) -generalizations. They are denoted by adding the superscript *lin* to the corresponding algorithm name: $\mathfrak{A}_1^{\text{lin}}$ and $\mathfrak{A}_2^{\text{lin}}$.

For simplicity, we formulate the algorithms for the case $n = 2$. They can be extended for arbitrary n straightforwardly.

The main data structure in these algorithms is an anti-unification triple (AUT) $x : T_1 \triangleq T_2$, where T_1 and T_2 are finite *consistent* sets of ground terms. The idea is that x is a common generalization of all terms in $T_1 \cup T_2$. A configuration is a tuple $A; S; r; \alpha_1; \alpha_2$, where A is a set of AUTs to be solved, S is a set of solved AUTs (the store), r is the generalization computed so far, and the α 's are the current approximations of generalization degree upper bounds of r for the input terms.

Before formulating the rules, we discuss one peculiarity of approximate generalizations:

Example 5. For a given \mathcal{R} and λ , assume $\mathcal{R}(a, b) \geq \lambda, \mathcal{R}(b, c) \geq \lambda, h \sim_{\mathcal{R}, \alpha}^{\{(1,1), (1,2)\}}$ f and $h \sim_{\mathcal{R}, \beta}^{\{(1,1)\}}$ g , where f is binary, g, h are unary, $\alpha \geq \lambda$ and $\beta \geq \lambda$. Then

- $h(b)$ is an (\mathcal{R}, λ) -generalization of $f(a, c)$ and $g(a)$.
- x is the only (\mathcal{R}, λ) -generalization of $f(a, d)$ and $g(a)$. One may be tempted to have h as the head of the generalization, e.g., $h(x)$, but x cannot be instantiated by any term that would be (\mathcal{R}, λ) -close to both a and d , since in the given \mathcal{R} , d is (\mathcal{R}, λ) -close only to itself. Hence, there would be no instance of $h(x)$ that is (\mathcal{R}, λ) -close to $f(a, d)$. Since there is no other alternative (except h) for the common neighbor of f and g , the generalization should be a fresh variable x .

This example shows that generalization algorithms should take into account not only the heads of the terms to be generalized, but also should look deeper, to make sure that the arguments grouped together by the given argument relation have a common neighbor. This justifies the requirement of consistency of a set of arguments, the notion introduced in the previous section and used in the decomposition rule below.

4.1 Anti-unification for Unrestricted Argument Relations

Algorithms $\mathfrak{A}_1^{\text{lin}}$ and \mathfrak{A}_1 use the rules below to transform configurations into configurations. Given \mathcal{R} , λ , and the ground terms t_1 and t_2 , we create the initial configuration $\{x : \{t_1\} \triangleq \{t_2\}\}; \emptyset; x; 1; 1$ and apply the rules as long as possible. Note that the rules preserve consistency of AUTs. The process generates a finite complete tree of derivations, whose terminal nodes have configurations with the first component empty. We will show how from these terminal configurations one collects the result as required in the anti-unification problem statement.

Tri: Trivial

$$\{x : \emptyset \triangleq \emptyset\} \uplus A; S; r; \alpha_1; \alpha_2 \implies A; S; r\{x \mapsto _ \}; \alpha_1; \alpha_2.$$

Dec: Decomposition

$$\{x : T_1 \triangleq T_2\} \uplus A; S; r; \alpha_1; \alpha_2 \implies \{y_i : Q_{i1} \triangleq Q_{i2} \mid 1 \leq i \leq n\} \cup A; S; r\{x \mapsto h(y_1, \dots, y_n)\}; \alpha_1 \wedge \beta_1; \alpha_2 \wedge \beta_2,$$

where $T_1 \cup T_2 \neq \emptyset$; h is n -ary with $n \geq 0$; y_1, \dots, y_n are fresh; and for $j = 1, 2$, if $T_j = \{t_1^j, \dots, t_{m_j}^j\}$, then

- $h \sim_{\mathcal{R}, \gamma_k^j}^{\rho_k^j} \text{head}(t_k^j)$ with $\gamma_k^j \geq \lambda$ for all $1 \leq k \leq m_j$ and $\beta_j = \gamma_1^j \wedge \dots \wedge \gamma_{m_j}^j$ (note that $\beta_j = 1$ if $m_j = 0$),
- for all $1 \leq i \leq n$, $Q_{ij} = \cup_{k=1}^{m_j} \{t_k^j \mid q \mid (i, q) \in \rho_k^j\}$ and is (\mathcal{R}, λ) -consistent.

Sol: Solving

$\{x : T_1 \triangleq T_2\} \uplus A; S; r; \alpha_1; \alpha_2 \implies A; \{x : T_1 \triangleq T_2\} \cup S; r; \alpha_1; \alpha_2$, if Tri and Dec rules are not applicable. (It means that at least one $T_i \neq \emptyset$ and either there is no h as it is required in the Dec rule, or at least one Q_{ij} from Dec is not (\mathcal{R}, λ) -consistent.)

Let *expand* be an *expansion operation* defined for sets of AUTs as

$$\text{expand}(S) := \{x : \prod_{t \in T_1} \text{rpc}_{\mathcal{R}, \lambda}(t) \triangleq \prod_{t \in T_2} \text{rpc}_{\mathcal{R}, \lambda}(t) \mid x : T_1 \triangleq T_2 \in S\}.$$

Exhaustive application of the three rules above leads to configurations of the form $\emptyset; S; r; \alpha_1; \alpha_2$, where r is a linear term. These configurations are further postprocessed, replacing S by $\text{expand}(S)$. We will use the letter E for expanded stores. Hence, terminal configurations obtained after the exhaustive rule application and expansion have the form $\emptyset; E; r; \alpha_1; \alpha_2$, where r is a linear term.¹ This is what Algorithm $\mathfrak{A}_1^{\text{lin}}$ stops with.

To an expanded store $E = \{y_1 : Q_{11} \triangleq Q_{12}, \dots, y_n : Q_{n1} \triangleq Q_{n2}\}$ we associate two sets of substitutions $\Sigma_L(E)$ and $\Sigma_R(E)$, defined as follows: $\sigma \in \Sigma_L(E)$ (resp. $\sigma \in \Sigma_R(E)$) iff $\text{dom}(\sigma) = \{y_1, \dots, y_n\}$ and $y_i \sigma \in Q_{i1}$ (resp. $y_i \sigma \in Q_{i2}$) for each $1 \leq i \leq n$. We call them the sets of *witness substitutions*.

Configurations containing expanded stores are called *expanded configurations*. From each expanded configuration $C = \emptyset; E; r; \alpha_1; \alpha_2$, we construct the set $S(C) := \{r, \sigma_1, \sigma_2, \alpha_1, \alpha_2 \mid \sigma_1 \in \Sigma_L(E), \sigma_2 \in \Sigma_R(E)\}$.

Given an anti-unification problem $\mathcal{R}, \lambda, t_1$ and t_2 , the *answer computed by Algorithm $\mathfrak{A}_1^{\text{lin}}$* is the set $S := \cup_{i=1}^m S(C_i)$, where C_1, \dots, C_m are all of the final expanded configurations reached by $\mathfrak{A}_1^{\text{lin}}$ for $\mathcal{R}, \lambda, t_1$, and t_2 .²

Example 6. Assume a, b, c and d are constants with $b \sim_{\mathcal{R}, 0.5}^{\emptyset} c, c \sim_{\mathcal{R}, 0.6}^{\emptyset} d$, and f, g and h are respectively binary, ternary and quaternary function symbols with $h \sim_{\mathcal{R}, 0.7}^{\{(1,1), (3,2), (4,2)\}} f$ and $h \sim_{\mathcal{R}, 0.8}^{\{(1,1), (3,3)\}} g$. For the proximity relation \mathcal{R} given in this way and $\lambda = 0.5$, Algorithm $\mathfrak{A}_1^{\text{lin}}$ performs the following steps to anti-unify $f(a, b)$ and $g(a, c, d)$:

$$\begin{aligned} & \{x : \{f(a, b)\} \triangleq \{g(a, c, d)\}; \emptyset; x; 1; 1 \implies_{\text{Dec}} \\ & \{x_1 : \{a\} \triangleq \{a\}, x_2 : \emptyset \triangleq \emptyset, x_3 : \{b\} \triangleq \{d\}, \\ & \quad x_4 : \{b\} \triangleq \emptyset\}; \emptyset; h(x_1, x_2, x_3, x_4); 0.7; 0.8 \implies_{\text{Dec}} \\ & \{x_2 : \emptyset \triangleq \emptyset, x_3 : \{b\} \triangleq \{d\}, x_4 : \{b\} \triangleq \emptyset\}; \emptyset; h(a, x_2, x_3, x_4); 0.7; 0.8 \implies_{\text{Tri}} \\ & \{x_3 : \{b\} \triangleq \{d\}, x_4 : \{b\} \triangleq \emptyset\}; \emptyset; h(a, _, x_3, x_4); 0.7; 0.8 \implies_{\text{Dec}} \\ & \{x_4 : \{b\} \triangleq \emptyset\}; \emptyset; h(a, _, c, x_4); 0.5; 0.6. \end{aligned}$$

Here *Dec* applies in two different ways, with the substitutions $\{x_4 \mapsto b\}$ and $\{x_4 \mapsto c\}$, leading to two final configurations: $\emptyset; \emptyset; h(a, _, c, b); 0.5; 0.6$ and $\emptyset; \emptyset; h(a, _, c, c); 0.5; 0.6$. The witness substitutions are the identity substitutions. We have $\mathcal{R}(h(a, _, c, b), f(a, b)) = 0.5, \mathcal{R}(h(a, _, c, b), g(a, c, d)) = 0.6, \mathcal{R}(h(a, _, c, c), f(a, b)) = 0.5$, and $\mathcal{R}(h(a, _, c, c), g(a, c, d)) = 0.6$.

If we had $h \sim_{\mathcal{R}, 0.7}^{\{(1,1), (1,2), (4,2)\}} f$, then the algorithm would perform only the *Sol* step, because in the attempt to apply *Dec* to the initial configuration, the set

¹ Note that no side of the AUTs in E in those configurations is empty due to the condition at the **Decomposition** rule requiring the Q_{ij} 's to be (\mathcal{R}, λ) -consistent.

² If we are interested only in linear generalizations *without witness substitutions*, there is no need in computing expanded configurations in $\mathfrak{A}_1^{\text{lin}}$.

$Q_{11} = \{a, b\}$ is inconsistent: $\mathbf{rpc}_{\mathcal{R},\lambda}(a) = \{a\}$, $\mathbf{rpc}_{\mathcal{R},\lambda}(b) = \{b, c\}$, and, hence, $\mathbf{rpc}_{\mathcal{R},\lambda}(a) \sqcap \mathbf{rpc}_{\mathcal{R},\lambda}(b) = \emptyset$.

Algorithm \mathfrak{A}_1 is obtained by further transforming the expanded configurations produced by $\mathfrak{A}_1^{\text{lin}}$. This transformation is performed by applying the **Merge** rule below as long as possible. Intuitively, its purpose is to make the linear generalization obtained by $\mathfrak{A}_1^{\text{lin}}$ less general by merging some variables.

Mer: Merge

$$\emptyset; \{x_1 : R_{11} \triangleq R_{12}, x_2 : R_{21} \triangleq R_{22}\} \uplus E; r; \alpha_1; \alpha_2 \implies \\ \emptyset; \{y : Q_1 \triangleq Q_2\} \cup E; r\sigma; \alpha_1; \alpha_2,$$

where $Q_i = (R_{1i} \sqcap R_{2i}) \neq \emptyset$, $i = 1, 2$, y is fresh, and $\sigma = \{x_1 \mapsto y, x_2 \mapsto y\}$.

The answer computed by \mathfrak{A}_1 is defined similarly to the answer computed by $\mathfrak{A}_1^{\text{lin}}$.

Example 7. Assume a, b are constants, f_1, f_2, g_1 , and g_2 are unary function symbols, p is a binary function symbol, and h_1 and h_2 are ternary function symbols. Let λ be a cut value and \mathcal{R} be defined as $f_i \sim_{\mathcal{R}, \alpha_i}^{\{(1,1)\}} h_i$ and $g_i \sim_{\mathcal{R}, \beta_i}^{\{(1,2)\}} h_i$ with $\alpha_i \geq \lambda$, $\beta_i \geq \lambda$, $i = 1, 2$. To generalize $p(f_1(a), g_1(b))$ and $p(f_2(a), g_2(b))$, we use \mathfrak{A}_1 . The derivation starts as

$$\{x : \{p(f_1(a), g_1(b))\} \triangleq \{p(f_2(a), g_2(b))\}\}; \emptyset; x; 1; 1 \implies_{\text{Dec}} \\ \{y_1 : \{f_1(a)\} \triangleq \{f_2(a)\}, y_2 : \{g_1(b)\} \triangleq \{g_2(b)\}\}; \emptyset; p(y_1, y_2); 1; 1 \implies_{\text{Sol}}^2 \\ \emptyset; \{y_1 : \{f_1(a)\} \triangleq \{f_2(a)\}, y_2 : \{g_1(b)\} \triangleq \{g_2(b)\}\}; p(y_1, y_2); 1; 1.$$

At this stage, we expand the store, obtaining

$$\emptyset; \{y_1 : \{f_1(a), h_1(a, _, _)\} \triangleq \{f_2(a), h_2(a, _, _)\}, \\ y_2 : \{g_1(b), h_1(_, b, _)\} \triangleq \{g_2(b), h_2(_, b, _)\}\}; p(y_1, y_2); 1; 1.$$

If we had the standard intersection \cap in the **Mer** rule, we would not be able to merge y_1 and y_2 , because the obtained sets in the corresponding AUTs are disjoint. However, **Mer** uses \sqcap : we have $\{f_i(a), h_i(a, _, _)\} \sqcap \{g_i(b), h_i(_, b, _)\} = \{h_i(a, b, _)\}$, $i = 1, 2$ and, therefore, can make the step

$$\emptyset; \{y_1 : \{f_1(a), h_1(a, _, _)\} \triangleq \{f_2(a), h_2(a, _, _)\}, \\ y_2 : \{g_1(b), h_1(_, b, _)\} \triangleq \{g_2(b), h_2(_, b, _)\}\}; p(y_1, y_2); 1; 1 \implies_{\text{Mer}} \\ \emptyset; \{z : \{h_1(a, b, _)\} \triangleq \{h_2(a, b, _)\}\}; p(z, z); 1; 1.$$

Indeed, if we take the witness substitutions $\sigma_i = \{z \mapsto h_i(a, b, _)\}$, $i = 1, 2$, and apply them to the obtained generalization, we get

$$p(z, z)\sigma_1 = p(h_1(a, b, _), h_1(a, b, _)) \simeq_{\mathcal{R}, \lambda} p(f_1(a), g_1(b)), \\ p(z, z)\sigma_2 = p(h_2(a, b, _), h_2(a, b, _)) \simeq_{\mathcal{R}, \lambda} p(f_2(a), g_2(b)).$$

Theorem 3. *Given \mathcal{R} , λ , and the ground terms t_1 and t_2 , Algorithm \mathfrak{A}_1 terminates for $\{x : \{t_1\} \triangleq \{t_2\}\}; \emptyset; x; 1; 1$ and computes an answer set \mathbf{S} such that*

1. the set $\{r \mid (r, \sigma_1, \sigma_2, \alpha_1, \alpha_2) \in S\}$ is an (\mathcal{R}, λ) -mcsrg of t_1 and t_2 ,
2. for each $(r, \sigma_1, \sigma_2, \alpha_1, \alpha_2) \in S$ we have $\mathcal{R}(r\sigma_i, t_i) \leq \alpha_i = \mathbf{gdub}_{\mathcal{R}, \lambda}(r, t_i)$, $i = 1, 2$.

Proof. Termination: Define the depth of an AUT $x : \{t_1, \dots, t_m\} \triangleq \{s_1, \dots, s_n\}$ as the depth of the term $f(g(t_1, \dots, t_m), h(s_1, \dots, s_n))$. The rules **Tri**, **Dec**, and **Sol** strictly reduce the multiset of depths of AUTs in the first component of the configurations. **Mer** strictly reduces the number of distinct variables in generalizations. Hence, these rules cannot be applied infinitely often and \mathfrak{A}_1 terminates.

In order to prove (1), we need to verify three properties:

- Soundness: If $(r, \sigma_1, \sigma_2, \alpha_1, \alpha_2) \in S$, then r is a relevant (\mathcal{R}, λ) -generalization of t_1 and t_2 .
- Completeness: If r' is a relevant (\mathcal{R}, λ) -generalization of t_1 and t_2 , then there exists $(r, \sigma_1, \sigma_2, \alpha_1, \alpha_2) \in S$ such that $r' \lesssim r$.
- Minimality: If r and r' belong to two tuples from S such that $r \neq r'$, then neither $r \prec_{\mathcal{R}, \lambda} r'$ nor $r' \prec_{\mathcal{R}, \lambda} r$.

Soundness: We show that each rule transforms an (\mathcal{R}, λ) -generalization into an (\mathcal{R}, λ) -generalization. Since we start from a most general (\mathcal{R}, λ) -generalization of t_1 and t_2 (a fresh variable x), at the end of the algorithm we will get an (\mathcal{R}, λ) -generalization of t_1 and t_2 . We also show that in this process all irrelevant positions are abstracted by anonymous variables, to guarantee that each computed generalization is relevant.

Dec: The computed h is (\mathcal{R}, λ) -close to the head of each term in $T_1 \cup T_2$. Q_{ij} 's correspond to argument relations between h and those heads, and each Q_{ij} is (\mathcal{R}, λ) -consistent, i.e., there exists a term that is (\mathcal{R}, λ) -close to each term in Q_{ij} . It implies that $x\sigma = h(y_1, \dots, y_n)$ (\mathcal{R}, λ) -generalizes all the terms from $T_1 \cup T_2$. Note that at this stage, $h(y_1, \dots, y_n)$ might not yet be a relevant (\mathcal{R}, λ) -generalization of T_1 and T_2 : if there exists an irrelevant position $1 \leq i \leq n$ for the (\mathcal{R}, λ) -generalization of T_1 and T_2 , then in the new configuration we will have an AUT $y_i : \emptyset \triangleq \emptyset$.

Tri: When **Dec** generates $y : \emptyset \triangleq \emptyset$, the **Tri** rule replaces y by $_$ in the computed generalization, making it relevant.

Sol does not change generalizations.

Mer merges AUTs whose terms have *nonempty* intersection of **rpc**'s. Hence, we can reuse the same variable in the corresponding positions in generalizations, i.e., **Mer** transforms a generalization computed so far into a less general one.

Completeness: We prove a slightly more general statement. Given two finite consistent sets of ground terms T_1 and T_2 , if r' is a relevant (\mathcal{R}, λ) -generalization for all $t_1 \in T_1$ and $t_2 \in T_2$, then starting from $\{x : T_1 \triangleq T_2\}; \emptyset; x; 1; 1$, Algorithm \mathfrak{A}_1 computes a $(r, \sigma_1, \sigma_2, \alpha_1, \alpha_2)$ such that $r' \lesssim r$.

We may assume w.l.o.g. that r' is a relevant (\mathcal{R}, λ) -lgg. Due to the transitivity of \lesssim , completeness for such an r' will imply it for all terms more general than r' .

We proceed by structural induction on r' . If r' is a (named or anonymous) variable, the statement holds. Assume $r' = h(r'_1, \dots, r'_n)$, $T_1 = \{u_1, \dots, u_m\}$, and $T_2 = \{w_1, \dots, w_l\}$. Then h is such that $h \sim_{\mathcal{R}, \beta_i}^{\rho_i} \text{head}(u_i)$ for all $1 \leq i \leq m$ and $h \sim_{\mathcal{R}, \gamma_j}^{\mu_j} \text{head}(w_j)$ for all $1 \leq j \leq l$. Moreover, each r'_k is a relevant (\mathcal{R}, λ) -generalization of $Q_{k1} = \cup_{i=1}^m \{u_i|_q \mid (k, q) \in \rho_i\}$ and $Q_{k2} = \cup_{j=1}^l \{w_j|_q \mid (k, q) \in \mu_j\}$ and, hence, Q_{k1} and Q_{k2} are (\mathcal{R}, λ) -consistent. Therefore, we can perform a step by Dec, choosing $h(y_1, \dots, y_k)$ as the generalization term and $y_i : Q_{i1} \triangleq Q_{i2}$ as the new AUTs. By the induction hypothesis, for each $1 \leq i \leq n$ we can compute a relevant (\mathcal{R}, λ) -generalization r_i for Q_{i1} and Q_{i2} such that $r'_i \lesssim r_i$.

If r' is linear, then the combination of the current Dec step with the derivations that lead to those r_i 's computes a tuple $(r, \dots) \in \mathbf{S}$, where $r = h(r_1, \dots, r_n)$ and, hence, $r' \lesssim r$.

If r' is non-linear, assume without loss of generality that all occurrences of a shared variable z appear as the direct arguments of h : $z = r'_{k_1} = \dots = r'_{k_p}$ for $1 \leq k_1 < \dots < k_p \leq n$. Since r' is an lgg, Q_{k_i1} and Q_{k_i2} cannot be generalized by a non-variable term, thus, Tri and Dec are not applicable. Therefore, the AUTs $y_i : Q_{k_i1} \triangleq Q_{k_i2}$ would be transformed by Sol. Since all pairs Q_{k_i1} and Q_{k_i2} , $1 \leq i \leq p$, are generalized by the same variable, we have $\sqcap_{t \in Q_j} \mathbf{rpc}_{\mathcal{R}, \lambda}(t) \neq \emptyset$, where $Q_j = \cup_{i=1}^p Q_{k_ij}$, $j = 1, 2$. Additionally, $r'_{k_1}, \dots, r'_{k_p}$ are all occurrences of z in r' . Hence, the condition of Mer is satisfied and we can extend our derivation with $p - 1$ -fold application of this rule, obtaining $r = h(r_1, \dots, r_n)$ with $z = r_{k_1} = \dots = r_{k_p}$, implying $r' \lesssim r$.

Minimality: Alternative generalizations are obtained by branching in Dec or Mer. If the current generalization r is transformed by Dec into two generalizations r_1 and r_2 on two branches, then $r_1 = h_1(y_1, \dots, y_m)$ and $r_2 = h_2(z_1, \dots, z_n)$ for some h 's, and fresh y 's and z 's. It may happen that $r_1 \lesssim_{\mathcal{R}, \lambda} r_2$ or vice versa (if h_1 and h_2 are (\mathcal{R}, λ) -close to each other), but neither $r_1 \prec_{\mathcal{R}, \lambda} r_2$ nor $r_2 \prec_{\mathcal{R}, \lambda} r_1$ holds. Hence, the set of generalizations computed before applying Mer is minimal. Mer groups AUTs together maximally, and different groupings are not comparable. Therefore, variables in generalizations are merged so that distinct generalizations are not $\prec_{\mathcal{R}, \lambda}$ -comparable. Hence, (1) is proven.

As for (2), for $i = 1, 2$, from the construction in Dec follows $\mathcal{R}(r\sigma_i, t_i) \leq \alpha_i$. Mer does not change α_i , thus, $\alpha_i = \mathbf{gdub}_{\mathcal{R}, \lambda}(r, t_i)$ also holds, since the way how α_i is computed corresponds exactly to the computation of $\mathbf{gdub}_{\mathcal{R}, \lambda}(r, t_i)$: $r \lesssim_{\mathcal{R}, \lambda} t_i$ and only the decomposition changes the degree during the computation. \square

The corollary below is proved similarly to Theorem 3:

Corollary 1. *Given \mathcal{R} , λ , and the ground terms t_1 and t_2 , Algorithm $\mathfrak{A}_1^{\text{lin}}$ terminates for $\{x : \{t_1\} \triangleq \{t_2\}\}; \emptyset; x; 1; 1$ and computes an answer set \mathbf{S} such that*

1. *the set $\{r \mid (r, \sigma_1, \sigma_2, \alpha_1, \alpha_2) \in \mathbf{S}\}$ is a minimal complete set of relevant linear (\mathcal{R}, λ) -generalizations of t_1 and t_2 ,*
2. *for each $(r, \sigma_1, \sigma_2, \alpha_1, \alpha_2) \in \mathbf{S}$ we have $\mathcal{R}(r\sigma_i, t_i) \leq \alpha_i = \mathbf{gdub}_{\mathcal{R}, \lambda}(r, t_i)$, $i = 1, 2$.*

4.2 Anti-unification with Correspondence Argument Relations

Correspondence relations make sure that for a pair of proximal symbols, no argument is irrelevant for proximity. Left- and right-totality of those relations guarantee that each argument of a term is close to at least one argument of its proximal term and the inverse relation remains a correspondence relation. Consequently, in the Dec rule of \mathfrak{A}_1 , the sets Q_{ij} never get empty. Therefore, the Tri rule becomes obsolete and no anonymous variable appears in generalizations. As a result, the (\mathcal{R}, λ) -mcsrg and the (\mathcal{R}, λ) -mcs g coincide, and the algorithm computes a solution from which we get an (\mathcal{R}, λ) -mcs g for the given anti-unification problem. The linear version $\mathfrak{A}_1^{\text{lin}}$ works analogously.

4.3 Anti-unification with Argument Mappings

When the argument relations are mappings, we are able to design a more constructive method for computing generalizations and their degree bounds (Recall that our mappings are partial injective functions, which guarantees that their inverses are also mappings.) We denote this algorithm by \mathfrak{A}_2 . The configurations stay the same as in before, but the AUTs in A will contain only empty or singleton sets of terms. In the store, we may still get (after the expansion) AUTs with term sets containing more than one element. Only the Dec rule differs from its previous counterpart, having a simpler condition:

Dec: Decomposition

$$\{x : T_1 \triangleq T_2\} \uplus A; S; r; \alpha_1; \alpha_2 \implies \{y_i : Q_{i1} \triangleq Q_{i2} \mid 1 \leq i \leq n\} \cup A; S; r\{x \mapsto h(y_1, \dots, y_n)\}; \alpha_1 \wedge \beta_1; \alpha_2 \wedge \beta_2,$$

where $T_1 \cup T_2 \neq \emptyset$; h is n -ary with $n \geq 0$; y_1, \dots, y_n are fresh; for $j = 1, 2$ and for all $1 \leq i \leq n$, if $T_j = \{t_j\}$ then $h \sim_{\mathcal{R}, \beta_j}^{\pi_j} \text{head}(t_j)$ and $Q_{ij} = \{t_j | \pi_j(i)\}$, and if $T_j = \emptyset$ then $\beta_j = 1$ and $Q_{ij} = \emptyset$.

This Dec rule is equivalent to the special case of Dec for argument relations where $m_j \leq 1$. The new Q_{ij} 's contain at most one element (due to mappings) and, thus, are always (\mathcal{R}, λ) -consistent. Various choices of h in Dec and alternatives in grouping AUTs in Mer cause branching in the same way as in \mathfrak{A}_1 . It is easy to see that the counterparts of Theorem 3 hold for \mathfrak{A}_2 and $\mathfrak{A}_2^{\text{lin}}$ as well.

A special case of this fragment of anti-unification is anti-unification for similarity relations in fully fuzzy signatures from [1]. Similarity relations are min-transitive proximity relations. The position mappings in [1] can be modeled by our argument mappings, requiring them to be total for symbols of the smaller arity and to satisfy the similarity-specific consistency restrictions from [1].

4.4 Anti-unification with Correspondence Argument Mappings

Correspondence argument mappings are bijections between arguments of function symbols of the same arity. For such mappings, if $h \simeq_{\mathcal{R}, \lambda}^{\pi} f$ and h is n -ary, then f is also n -ary and π is a permutation of $(1, \dots, n)$. Hence, \mathfrak{A}_2 combines

in this case the properties of \mathfrak{A}_1 for correspondence relations (Sect. 4.2) and of \mathfrak{A}_2 for argument mappings (Sect. 4.3): all generalizations are relevant, computed answer gives an mcsg of the input terms, and the algorithm works with term sets of cardinality at most 1.

5 Remarks About the Complexity

The proximity relation \mathcal{R} can be naturally represented as an undirected graph, where the vertices are function symbols and an edge between them indicates that they are proximal. Graphs induced by proximity relations are usually sparse. Therefore we can represent them by (sorted) adjacency lists. In the adjacency lists, we can also accommodate the argument relations and proximity degrees.

In the rest of this section we use the following notation:

- n : the size of the input (number of symbols) of the corresponding algorithms,
- Δ : the maximum degree of \mathcal{R} considered as a graph,
- \mathfrak{a} : the maximum arity of function symbols that occur in \mathcal{R} .
- $m^{\bullet n}$: a function defined on natural numbers m and n such that $1^{\bullet n} = n$ and $m^{\bullet n} = m^n$ for $m \neq 1$.

We assume that the given anti-unification problem is represented as a completely shared directed acyclic graph (dag). Each node of the dag has a pointer to the adjacency list (with respect to \mathcal{R}) of the symbol in the node.

Theorem 4. *Time complexities of \mathfrak{C} and the linear versions of the generalization algorithms are as follows:*

- \mathfrak{C} for argument relations and $\mathfrak{A}_1^{\text{lin}}$: $O(n \cdot \Delta \cdot \Delta^{\bullet \mathfrak{a}^n})$,
- \mathfrak{C} for argument mappings and $\mathfrak{A}_2^{\text{lin}}$: $O(n \cdot \Delta \cdot \Delta^{\bullet n})$.

Proof (Sketch). In \mathfrak{C} , in the case of argument relations, an application of the Red rule to a state $\mathbf{I}; s$ replaces one element of \mathbf{I} of size m by at most \mathfrak{a} new elements, each of them of size $m - 1$. Hence, one branch in the search tree for \mathfrak{C} , starting from a singleton set \mathbf{I} of size n , will have the length at most $l = \sum_{i=0}^{n-1} \mathfrak{a}^i$. At each node on it there are at most Δ choices of applying Red with different h 's, which gives the total size of the search tree to be at most $\sum_{i=0}^{l-1} \Delta^i$, i.e., the number of steps performed by \mathfrak{C} in the worst case is $O(\Delta^{\bullet \mathfrak{a}^n})$. Those different h 's are obtained by intersecting the proximity classes of the heads of terms $\{t_1, \dots, t_m\}$ in the Red rule. In our graph representation of the proximity relation, proximity classes of symbols are exactly the adjacency lists of those symbols which we assume are sorted. Their maximal length is Δ . Hence, the work to be done at each node of the search tree of \mathfrak{C} is to find the intersection of at most n sorted lists, each containing at most Δ elements. It needs $O(n \cdot \Delta)$ time. It gives the time complexity $O(n \cdot \Delta \cdot \Delta^{\bullet \mathfrak{a}^n})$ of \mathfrak{C} for the relation case.


In the mapping case, an application of the Red rule to a state $\mathbf{I}; s$ replaces one element of \mathbf{I} of size m by at most \mathfrak{a} new elements of the total size $m - 1$. Therefore, the maximal length of a branch is n , the branching factor is Δ , and

the amount of work at each node, like above, is $O(n \cdot \Delta)$. Hence, the number of steps in the worst case is $O(\Delta^{\bullet n})$ and the time complexity of \mathfrak{C} is $O(n \cdot \Delta \cdot \Delta^{\bullet n})$.

The fact that consistency check is incorporated in the Dec rule in $\mathfrak{A}_1^{\text{lin}}$ can be used to guide the application of this rule, using the values memoized by the previous applications of Red. The very first time, the appropriate h in Dec is chosen arbitrarily. In any subsequent application of this rule, h is chosen according to the result of the Red rule that has already been applied to the arguments of the current AUT for their consistency check, as required by the condition of Dec. In this way, the applications of Dec and Sol will correspond to the applications of Red. There is a natural correspondence between the applications of Rem and Tri rules. Therefore, $\mathfrak{A}_1^{\text{lin}}$ will have the search tree analogous to that of \mathfrak{C} . Hence the complexity of $\mathfrak{A}_1^{\text{lin}}$ is $O(n \cdot \Delta \cdot \Delta^{\bullet n})$. $\mathfrak{A}_2^{\text{lin}}$ does not call the consistency check, but does the same work as \mathfrak{C} and, hence, has the same complexity $O(n \cdot \Delta \cdot \Delta^{\bullet n})$. \square

6 Discussion and Conclusion

The diagram below illustrates the connections between different anti-unification problems based on argument relations:



The arrows indicate the direction from more general problems to more specific ones. For the unrestricted cases (left column) we compute mcsrg's. For correspondence relations and correspondence mappings (right column), mcsg's are computed. (In fact, for them, the notions of mcsrg and mcsg coincide). The algorithms for relations (upper row) are more involved than those for mappings (lower row): Those for relations deal with AUTs containing arbitrary sets of terms, while for mappings, those sets have cardinality at most one, thus simplifying the conditions in the rules. Moreover, the two cases in the lower row generalize the existing anti-unification problems:

- the unrestricted mappings case generalizes the problem from [1] by extending similarity to proximity and relaxing the smaller-side-totality restriction;
- the correspondence mappings case generalizes the problem from [9] by allowing permutations between arguments of proximal function symbols.

All our algorithms can be easily turned into anti-unification algorithms for crisp tolerance relations³ by taking lambda-cuts and ignoring the computation of the approximation degrees. Besides, they are modular and can be used to compute only linear generalizations by just skipping the merging rule. We provided complexity estimations for the algorithms that compute linear generalizations (that often are of practical interest).

³ Tolerance: reflexive, symmetric, not necessarily transitive relation. According to Poincaré, a fundamental notion for mathematics applied to the physical world.

In this paper, we did not consider cases when the same pair of symbols is related to each other by more than one argument relation. Our results can be extended to them, that would open a way towards approximate anti-unification modulo background theories specified by shallow collapse-free axioms. Another interesting direction of future work would be extending our results to quantitative algebras [10] that also deal with quantitative extensions of equality.

Acknowledgments. Supported by the Austrian Science Fund, project P 35530.

References

1. Ait-Kaci, H., Pasi, G.: Fuzzy lattice operations on first-order terms over signatures with similar constructors: a constraint-based approach. *Fuzzy Sets Syst.* **391**, 1–46 (2020). <https://doi.org/10.1016/j.fss.2019.03.019>
2. Baader, F., Nipkow, T.: *Term Rewriting and All That*. Cambridge University Press, Cambridge (1998)
3. Bader, J., Scott, A., Pradel, M., Chandra, S.: Getafix: learning to fix bugs automatically. *Proc. ACM Program. Lang.* **3**(OOPSLA), 159:1–159:27 (2019). <https://doi.org/10.1145/3360585>
4. Barwell, A.D., Brown, C., Hammond, K.: Finding parallel functional pearls: Automatic parallel recursion scheme detection in Haskell functions via anti-unification. *Future Gener. Comput. Syst.* **79**, 669–686 (2018). <https://doi.org/10.1016/j.future.2017.07.024>
5. Dershowitz, N., Manna, Z.: Proving termination with multiset orderings. *Commun. ACM* **22**(8), 465–476 (1979). <https://doi.org/10.1145/359138.359142>
6. Galitsky, B.: *Developing Enterprise Chatbots - Learning Linguistic Structures*. Springer, Heidelberg (2019). <https://doi.org/10.1007/978-3-030-04299-8>
7. Kirbas, S., et al.: On the introduction of automatic program repair in Bloomberg. *IEEE Softw.* **38**(4), 43–51 (2021). <https://doi.org/10.1109/MS.2021.3071086>
8. Kutsia, T., Pau, C.: A framework for approximate generalization in quantitative theories. *RISC Report Series 22-04*, Research Institute for Symbolic Computation, Johannes Kepler University Linz (2022). <https://doi.org/10.35011/risc.22-04>
9. Kutsia, T., Pau, C.: Matching and generalization modulo proximity and tolerance relations. In: Özgün, A., Zinova, Y. (eds.) *TbiLLC 2019*. LNCS, vol. 13206, pp. 323–342. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-98479-3_16
10. Mardare, R., Panangaden, P., Plotkin, G.D.: Quantitative algebraic reasoning. In: Grohe, M., Koskinen, E., Shankar, N. (eds.) *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2016*, pp. 700–709. ACM (2016). <https://doi.org/10.1145/2933575.2934518>
11. Mehta, S., et al.: Rex: preventing bugs and misconfiguration in large services using correlated change analysis. In: Bhagwan, R., Porter, G. (eds.) *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2020*, Santa Clara, CA, USA, 25–27 February 2020, pp. 435–448. USENIX Association (2020). <https://www.usenix.org/conference/nsdi20/presentation/mehta>
12. Plotkin, G.D.: A note on inductive generalization. *Mach. Intell.* **5**(1), 153–163 (1970)
13. Raza, M., Gulwani, S., Milic-Frayling, N.: Programming by example using least general generalizations. In: Brodley, C.E., Stone, P. (eds.) *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 27–31 July 2014, Québec City, Québec, Canada, pp. 283–290. AAAI Press (2014)

14. Reynolds, J.C.: Transformational systems and the algebraic structure of atomic formulas. *Mach. Intell.* **5**(1), 135–151 (1970)
15. Rolim, R., Soares, G., Gheyi, R., D’Antoni, L.: Learning quick fixes from code repositories. *CoRR* abs/1803.03806 (2018). <http://arxiv.org/abs/1803.03806>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

