

# Gradual Sensitivity Typing

Damián Arquez

PLEIAD Laboratory

Computer Science Department (DCC)

University of Chile &amp; IMFD

Santiago, Chile

darquez@dcc.uchile.cl

Matías Toro

PLEIAD Laboratory

Computer Science Department (DCC)

University of Chile &amp; IMFD

Santiago, Chile

mtoro@dcc.uchile.cl

Éric Tanter

PLEIAD Laboratory

Computer Science Department (DCC)

University of Chile &amp; IMFD

Santiago, Chile

etanter@dcc.uchile.cl

**Abstract**—Reasoning about the sensitivity of functions with respect to their inputs has interesting applications in various areas, such as differential privacy. Sensitivity type systems support the modular checking and enforcement of sensitivity, but can be overly conservative for certain useful programming patterns. Hence, there is value in bringing the benefits of gradual typing to these disciplines in order to ease their adoption and applicability. In this work, we motivate, formalize, and prototype gradual sensitivity typing. The language GSoul supports both the unrestricted *unknown sensitivity* and bounded imprecision in the form of *sensitivity intervals*. Gradual sensitivity typing allows programmers to smoothly evolve typed programs without any static sensitivity information towards hardened programs with a mix of static and dynamic sensitivity checking. Additionally, gradual sensitivity typing seamlessly enables precise runtime sensitivity checks whenever fully static checking would be overly conservative. We formalize a core of the language GSoul and prove that it satisfies both the gradual guarantees and sensitivity type soundness, known as *metric preservation*. We establish that, in general, gradual metric preservation is termination insensitive, and that one can achieve termination-sensitive gradual metric preservation by hardening specifications to bounded imprecision. Furthermore, we show how the latter can be used to formally reason about the differential privacy guarantees of gradually-typed programs. We implement a prototype of GSoul, which provides an interactive test bed for further exploring the potential of gradual sensitivity typing.

**Index Terms**—gradual typing, differential privacy, function sensitivity

## I. INTRODUCTION

Function sensitivity, also called Lipschitz continuity, is an upper bound of how much the output may change given an input perturbation. More formally, a function  $f$  is  $s$ -sensitive iff for all  $x$  and  $y$ ,  $|f(x) - f(y)| \leq s|x - y|$ . Sensitivity plays an important role in different areas such as control theory [1], dynamic systems [2], program analysis [3] and differential privacy [4]. For example, a function  $f$  can be made differentially private by adding random noise to its output. The added noise has to be sufficient to guarantee privacy but also tight enough to avoid compromising the utility of the private result. This is usually achieved by calibrating the amount of noise using the sensitivity of the function.

Many approaches have been proposed to reason about sensitivity, either statically [5, 6, 7, 8, 9, 10, 11, 12], or

dynamically [13]. Handling sensitivity statically as an effect in the sense of type-and-effects systems [14] has the advantage of providing strong guarantees statically. However, type-and-effect systems can make certain useful programming patterns tedious or overly conservative. Combining the advantages of static and dynamic typechecking is a very active area with a long history, but, to the best of our knowledge, sensitivity has not been explored under this perspective. One prominent approach for combining static and dynamic typechecking is gradual typing [15]. Gradual typing supports the smooth transition between dynamic and static checking within the same language by introducing imprecise types and consistent relations. Imprecision is handled optimistically during static typechecking, and is backed by runtime checks in order to detect potential violations of static assumptions. Hence, the programmer is able to choose at a fine-grained level which portions of the program are dynamically checked and which ones are statically checked. Gradual typing has been studied in many settings such as subtyping [16, 17], references [18, 19], ownership [20], information-flow security typing [21, 22, 23] and refinement types [24], among (many) others.

The fact that sensitivity is a quantity makes static reasoning particularly challenging, and forces simply-typed approaches to over-approximate sensitivity in problematic ways. For instance, unless one considers some form of dependent typing [6], the sensitivity of a function that recurses over its argument would be deemed infinite [5]. Also, the potential for divergence in a language introduces the possibility of different interpretations of sensitivity type soundness, known as *metric preservation* [5]. Metric preservation is a hyperproperty that captures the bound on how much the result of two similar computations may change given an input variation. In the presence of possible divergence, the situation is similar to that of information flow security, where both termination-sensitive and termination-insensitive notions of noninterference have been studied [25, 26, 27]. For metric preservation, a *termination-insensitive* interpretation says that, if the function terminates on both inputs, then the output differences is bounded [13, 28]. A *termination-sensitive* interpretation says that, if the function terminates on the first input, then it also terminates on the second, and the output differences is bounded [5, 6]. Azevedo de Amorim et al. [28] study a termination-insensitive characterization of metric preservation

This work was partially funded by the Millennium Science Initiative Program: code ICN17\_002 and ANID FONDECYT Project 11250054.

as well as the necessary conditions to establish when two programs behave the same in terms of termination.

This work studies the integration of gradual typing with sensitivity typing, with the following contributions:

- We introduce and motivate GSoul, a functional programming language with gradual sensitivity typing (§III), featuring a latent type-and-effect discipline with explicit sensitivity polymorphism. Notably, GSoul not only supports the *unknown* sensitivity—which stands for any sensitivity, possibly infinite—but also a *bounded* form of imprecision in the form of *sensitivity intervals*. Bounded imprecision provides programmers with a fine-grained mechanism to balance flexibility and static checking.
- After reviewing the key concepts of semantic sensitivity typing with a core language and a static syntactic typing discipline (§III), we introduce GS, a core subset of GSoul that captures the key elements of gradual sensitivities (§IV).
- We describe the semantic model of GS types, formalizing the novel notion of gradual metric preservation, both termination-insensitive and termination-sensitive (§V). This presentation leaves the exact runtime monitoring mechanism abstract.
- Following the traditional approach of gradual languages, we describe an internal language GSCheck that features a type-driven runtime monitoring mechanism based on evidences [16] (§VI). We then describe syntactic typing for GS via its elaboration to GSCheck, and establish its metatheory (§VII): type safety, gradual guarantees [29], and most importantly, sensitivity type soundness. Syntactically well-typed GS programs satisfy gradual metric preservation, and furthermore, if their sensitivity type imprecision is *bounded*, they satisfy termination-sensitive gradual metric preservation.
- We discuss the importance of termination-sensitive gradual metric preservation for differential privacy [4], describing two variants of core differentially-private mechanisms and algorithms in GSoul (§VIII).

We discuss related work in §IX and conclude in §X. The supplementary material provides proofs of all the formal results presented here. A prototype implementation of GSoul is available at <https://github.com/pleiad/g soul-lang>.

## II. GRADUAL SENSITIVITY TYPING IN ACTION

We illustrate the limitations of the conservative nature of static sensitivity typing both for recursive functions, and for a practical scenario based on a standard differential privacy algorithm, Above Threshold [4]. We then show how gradual sensitivity typing can address these issues, by supporting the whole spectrum between fully dynamic and fully static sensitivity checking.

*a) Limitations of static sensitivity typing:* We start by exploring how recursive functions are dealt with in a static sensitivity discipline, using FUZZ [5] as a reference. In this language, 1-sensitive functions are given the special type  $\tau_1 \multimap \tau_2$ . Dually, functions of unbounded sensitivity are written as

usual:  $\tau_1 \rightarrow \tau_2$ . For a minimal example, consider the following definition of a recursive `scale` function (using concrete FUZZ syntax), which receives a factor and a value to be scaled:

```
scale 0 v = 0
scale n v = v + scale (n-1) v
```

The only possible FUZZ type for this function is  $\mathbb{N} \rightarrow \mathbb{R} \rightarrow \mathbb{R}$ , which characterizes it as infinitely sensitive in both arguments (as expressed by the use of arrows  $\rightarrow$ , instead of lollipop  $\multimap$ ). However, notice that the sensitivity on the scaled value—the second argument—actually depends on the value of the factor argument. Because this value is unknown statically, the type system over-approximates the sensitivity of the function as infinite. DFUZZ [6] tackles this issue via *lightweight dependent types*: `scale` can be given the type  $\forall i. \mathbb{N}[i] \rightarrow \mathbb{R} \multimap_i \mathbb{R}$ , denoting that the function is  $i$ -sensitive in the second argument. However, such dependent types make the type system more complex for users to harness, and do not always capture more precise sensitivities than FUZZ (e.g. if at the call site of `scale` the first argument value is not known statically, infinity remains the only sound option).

Another problematic restriction of static sensitivity typing manifests for recursive function definitions that capture free variables from their surrounding scope. Consider the FUZZ typing rule for fixpoints [5]:

$$\frac{\Gamma, f :_{\infty} \tau_1 \multimap \tau_2 \vdash e : \tau_1 \multimap \tau_2}{\infty \Gamma \vdash \text{fix } f.e : \tau_1 \multimap \tau_2}$$

This rule reports the fixpoint as infinitely sensitive in captured variables, as observable in the conclusion of the rule where  $\Gamma$  is scaled by  $\infty$ . Indeed, there is no way to know statically how many recursive calls will be performed, and therefore how many times the captured variable will be accessed. Azevedo de Amorim et al. [28] also identify this issue and design a fixpoint rule that provides a better bound for a particular class of recursive constructs, such as exponentially decaying lists, but their rule does not cover the general case.

Over-approximation in static sensitivity typing can also be problematic in many other practical scenarios. For instance, consider the Above Threshold (AT) algorithm from the differential privacy literature [4]. AT receives a list of 1-sensitive queries and a threshold, and approximately returns the index of the first query whose result is above the threshold. Suppose we have a list of functions  $f_1, \dots, f_n$  written in FUZZ, each representing a query over a database. Additionally, let us assume that each  $f_i$  is *semantically* 1-sensitive. Since a type system is but a syntactic conservative *approximation* of the semantics, we can expect each function  $f_i$  to have type  $\text{DB} \multimap_{s_i} \mathbb{R}$ , where  $1 \leq s_i$ . Building a list of such functions would necessarily give the list a type that accounts for the *worst* sensitivity of the given functions:  $(\text{DB} \multimap_{s_{\max}} \mathbb{R})$  list, where  $s_{\max} = \max(s_1, \dots, s_n)$ . If the sensitivity of a single function in the list is over-approximated, the whole list is impossible to use with AT given that the algorithm expects a list of strictly 1-sensitive queries.

<sup>1</sup>We leave the DB type abstract, but the reader can think of it as a set of records, where each record is a tuple of values.

Instead of a list, one could use a tuple of functions of type  $\text{DB} \multimap_{s_1} \mathbb{R} \times \dots \times \text{DB} \multimap_{s_n} \mathbb{R}$ , thereby retaining the precise sensitivity information of each function in the tuple. However, if we need to project out the tuple given a statically-unknown index, we are back to the same issue: the minimal sound return type for  $\lambda(i : \mathbb{N}).\lambda(Q : \text{DB} \multimap_{s_1} \mathbb{R} \times \dots \times \text{DB} \multimap_{s'_n} \mathbb{R}).Q[i]$  is  $\text{DB} \multimap_{s_{\max}} \mathbb{R}$ , so any  $Q[i]$  is deemed  $s_{\max}$ -sensitive by the type system, despite the fact that, when evaluated it is actually a  $s_i$ -sensitive function. Note that this scenario is outside the reach of DFUZZ as well.

b) *Gradual sensitivities*: Gradual sensitivity typing as provided in GSoul can elegantly address the limitations of static sensitivity typing, by selectively relying on runtime checking whenever desired or needed. In contrast to runtime sensitivity monitoring [13], which completely foregoes static checking, gradual sensitivity typing does allow static guarantees to be enforced where possible. To support gradual sensitivity typing, sensitivities in GSoul are extended with an unknown sensitivity  $?$ , akin to the unknown type of gradual typing [15]. Intuitively,  $?$  represents *any* sensitivity and allows the programmer to introduce imprecision in the sensitivity information, relying on optimistic static checking, backed by runtime checks.

Gradual sensitivity ordering  $\lesssim$  is defined by plausibility: an unknown sensitivity  $?$  is plausibly both smaller and greater than any other sensitivity, *i.e.*  $? \lesssim s$  and  $s \lesssim ?$ , for any sensitivity  $s$ . Notice that  $\lesssim$  is not transitive:  $2 \lesssim ?$  and  $? \lesssim 1$ , but  $2 \not\lesssim 1$ . Plausible ordering induces a notion of *consistent subtyping*: given a resource  $r$  in scope, the type  $\text{Number}[r] \rightarrow \text{Number}[?r]$  is a consistent subtype of  $\text{Number}[r] \rightarrow \text{Number}[r]$  because  $? \lesssim 1$ . This notion of consistent subtyping is akin to that studied for other gradual typing disciplines [17, 30, 24, 16], but here focused on the sensitivity information conveyed in types.

The flexibility afforded by consistent subtyping is backed by runtime checks at the boundaries between types of different precision, ensuring at runtime that no static assumptions are silently violated. For instance, revisiting the AT example in GSoul, we can give the list of queries the *imprecise* type  $\text{List}\langle [r](\text{DB}[r] \rightarrow \text{Number}[?r]) \rangle$ , *i.e.* a list of functions with unknown sensitivity<sup>2</sup>. As  $?$  is plausibly smaller than any other gradual sensitivity, the functions in the list can optimistically be passed to AT, without the need to statically comply with the 1-sensitivity restriction. If one of the functions in the list happens to violate the 1-sensitivity assumption, AT will simply halt with an error. In §VIII we define a gradual version of AT that actually avoids erroring—even in the presence of queries that are more than 1-sensitive—and prove its differential privacy guarantees.

Gradual sensitivity typing also offers a simple and effective alternative to handle recursive functions in an exact, albeit dynamically-checked, manner. For recursive functions that capture external variables, the programmer can use the

<sup>2</sup>The syntax  $[r](\dots)$ , used inside the list type parameter, creates a resource-polymorphic type. In practice, these binders could be inferred to make the user-syntax more lightweight.

unknown sensitivity to defer the sensitivity check to runtime, where the number of recursive calls can be observed. Likewise for functions whose sensitivity with respect to an argument depends on the number of recursive calls: consider the `scale` function now defined in GSoul:

```
def scale(n: Number, res v: Number): Number[?v] =
  if (n == 0) then 0 else v + scale(n - 1, v);
```

The `res` modifier in the argument  $v$  indicates that it should be tracked as a resource, and thus the return type of `scale` can specify its sensitivity with respect to  $v$ . Since there is no sound static sensitivity type for `scale`, we can use the unknown sensitivity to defer the sensitivity check to runtime, by annotating the return value as  $\text{Number}[?v]$ . Then, assuming  $f$  requires an argument that is at most 10-sensitive, we obtain the following behavior:

```
f(scale(10, x)) // typechecks, runs successfully
f(scale(11, x)) // typechecks, fails at runtime
```

Therefore, gradual sensitivity typing not only allows programmers to choose between static and dynamic checking as they see fit, but it can also accommodate features that are too conservatively handled by the static discipline.

c) *Bounded imprecision*: A key idea from the onset of gradual typing [15] is that, while the unknown type  $?$  stands for any type whatsoever, one can also use *partially-imprecise* type information—for instance using type  $\text{Int} \rightarrow ?$  to denote any function type whose domain is  $\text{Int}$  but codomain is left open. Partially-imprecise types are key to support the smooth refinement of type information in gradual programs [29].

Transposing this idea to the world of gradual sensitivities, we support not only the unknown sensitivity  $?$ , which denotes any sensitivity, but also bounded intervals such as  $[1, 20]$ , which denote any sensitivity *within a bounded range*<sup>3</sup>.

Sensitivity intervals allow programmers to finely trade flexibility for stronger static guarantees, as well as avoiding runtime checks when possible. For instance, assume a number variable  $x$  in scope that is 1-sensitive in a resource  $r$ . Suppose that we build a list  $\mathfrak{l}$  that contains values obtained by scaling  $x$  by different factors, *e.g.*:

```
let l : List<Number[_]> =
  List(scale(1, x), scale(2, x), scale(3, x));
```

The question arises as to which sensitivity to declare in the type of  $\mathfrak{l}$  (the  $_$  above). To illustrate the implications, consider that we apply one of the following functions, which defer in the sensitivity type of their argument:

```
f: Number[0r] -> Unit    g: Number[1r] -> Unit
h: Number[3r] -> Unit
```

to some index of  $\mathfrak{l}$ , say  $0$ , *i.e.* we compute either  $f(\mathfrak{l}[0])$ ,  $g(\mathfrak{l}[0])$  or  $h(\mathfrak{l}[0])$ . We summarize several cases below depending on which sensitivity is declared for  $\mathfrak{l}$  and which function we choose to apply to the first element of  $\mathfrak{l}$ :

<sup>3</sup>Syntactically, we can unify all forms of sensitivities with intervals, considering that a single sensitivity  $s$  is shorthand for  $[s, s]$ , and that the unknown sensitivity  $?$  is shorthand for  $[0, \infty]$ .

sensitivity of $\ulcorner \_ \urcorner$	apply $f$	apply $g$	apply $h$
$3r$	type error	type error	no check
$?r$	runtime error	check ok	check ok
$0..3r$	runtime error	check ok	no check
$1..3r$	type error	check ok	no check

In a static language like DFuzz [6] in which `scale` can be precisely typed, the most specific and sound sensitivity to declare for  $\ulcorner \_ \urcorner$  is the maximum sensitivity of the values put in the list, *i.e.*  $3r$ . In this case, the applications of  $f$  and  $g$  are rejected by the type system—despite the fact that the underlying value ( $\ulcorner 0 \urcorner$ ) is really 1-sensitive in  $r$ . The application of  $h$  is well-typed, and requires no runtime checking.

If we declare the sensitivity to be unknown  $?r$ , then all applications are optimistically accepted by the gradual type system, given that it is *plausible* for the passed value to meet each function’s sensitivity requirement. This static plausibility gives rise to a runtime check, which fails in the case of  $f$ . This is in essence similar to fully dynamic sensitivity checking as in DDuo [13]. In particular, note that a runtime check is necessary including for the application of  $h$ . In contrast, if we refine the upper bound on the sensitivity using  $0..3r$ , then the runtime check for the application of  $h$  is not necessary any longer because the sensitivity of the passed value is *definitely* valid. Likewise, if we choose the even more precise interval  $1..3r$  then applying  $f$  is now statically rejected by the type system, because  $\ulcorner 0 \urcorner$  is *definitely* at least 1-sensitive in  $r$ . This shows that the gradual type system can exploit partially-imprecise information to deduce definite judgments, which are helpful because they can be (a) reported to users (*e.g.* in the IDE) and (b) exploited by the implementation to avoid unnecessary runtime checks—an objective known as “*pay-as-you-go*” [15].

While our formal treatment (and GSOUL prototype) of gradual sensitivity typing does not currently deal with optimizations, we uncover a strong semantic argument in favor of using bounded sensitivity intervals: they make it possible to provide a stronger, termination-sensitive soundness property for gradual sensitivity typing (§V-C), which is crucial for reasoning about differential privacy (§VIII).

### III. BACKGROUND: TYPING AND SENSITIVITY

We first recap essential concepts of typing and semantics using a simply-typed lambda calculus with primitive operations, which lie at the heart of the development of this work. We then complement this background with the standard notion of sensitivity type soundness.

#### A. Semantics of Types for a Core Calculus

The syntax of the core calculus is the following:

$B ::= \mathbb{R} \mid \mathbb{B}$	(base types)
$\tau ::= B \mid \tau \rightarrow \tau$	(types)
$e ::= c \mid \text{op } \bar{e} \mid x \mid \lambda(x : \tau).e \mid ee$	(expressions)

Types  $\tau$  are either base types (real numbers or booleans) or function types. Expressions  $e$  are either constants  $c$  *i.e.* literal values of base types (such as `true` or `false` for type  $\mathbb{B}$ ), primitive operations `op` applied to a list of expressions  $\bar{e}$ , variables  $x$ , lambdas  $\lambda(x : \tau).e$ , or applications  $ee$ . We write  $e \Downarrow v$  to

denote that expression  $e$  reduces to a value  $v$ ; the call-by-value operational semantics of the calculus are standard [31].

Semantically, we say that a closed expression  $e$  has type  $\tau$ , written  $\models e : \tau$ , if it evaluates to a value of type  $\tau$ , *i.e.* a canonical form of the type  $\tau$ . We write  $\llbracket \tau \rrbracket$  to denote the set of canonical forms of type  $\tau$ , *e.g.*  $\llbracket \mathbb{B} \rrbracket = \{\text{true}, \text{false}\}$ . Formally,<sup>4</sup>

$$\models e : \tau \stackrel{\text{def}}{=} e \Downarrow v \wedge v \in \llbracket \tau \rrbracket$$

As standard in typed calculi, this notion can be extended to open expressions by introducing type environments  $\Gamma$ , which map variables to types. Then, we may use substitutions  $\gamma$ —also known as value environments—to close an expression by providing values for its free variables. We write  $\gamma : \Gamma$  to denote that the substitution  $\gamma$  is well-typed with respect to  $\Gamma$ :

$$\gamma : \Gamma \stackrel{\text{def}}{=} \forall x \in \text{dom}(\Gamma). \models \gamma(x) : \Gamma(x)$$

Using substitutions, we define semantic typing for open expressions: an open expression  $e$  has type  $\tau$  under a type environment  $\Gamma$  if, given a well-typed substitution  $\gamma$ , the closed expression  $\gamma(e)$  is of type  $\tau$ :

$$\Gamma \models e : \tau \stackrel{\text{def}}{=} \forall \gamma \text{ s.t. } \gamma : \Gamma. \models \gamma(e) : \tau$$

Notice that this *semantic* account of typing is independent of any *syntactic* type system—*i.e.* a decidable typechecking procedure—traditionally noted with  $\vdash$ , instead of  $\models$ . One can consider the standard typing rules for this core calculus [31], and prove the *soundness* of the syntactic type system: if an expression is syntactically well-typed, then it is semantically well-typed, *i.e.*  $\vdash e : \tau \implies \models e : \tau$ .

#### B. Semantics of Sensitivity

Let us explore how to reason about the sensitivity guarantee of expressions in this simple calculus. We start by revisiting the fundamentals of function sensitivity, and then extend this notion to open expressions in our calculus, where free variables can vary independently.

The sensitivity of a function is a measure of how much the output of the function can vary with respect to variations in the input. For instance, a function  $f : \mathbb{R} \rightarrow \mathbb{R}$  is said to be  $s$ -sensitive if and only if for any  $x, y \in \mathbb{R}$ ,  $|f(x) - f(y)| \leq s \cdot |x - y|$ . This notion can be further generalized to functions between arbitrary metric spaces, *i.e.* types  $T$  equipped with a metric  $d_T$ :

**Definition 1** (Function Sensitivity). *A function  $f : T_1 \rightarrow T_2$  is  $s$ -sensitive iff for any  $d \in \mathbb{R}$  and  $a_1, a_2 \in T_1$  such that  $d_{T_1}(a_1, a_2) \leq d$ , we have:*

$$f(a_1) = v_1 \wedge f(a_2) = v_2 \wedge d_{T_2}(v_1, v_2) \leq s \cdot d$$

Following the approach of Reed and Pierce [5], instead of reasoning about the sensitivity of functions, we can reason about the sensitivity of open expressions. In this setting,  $f$

<sup>4</sup>Note that this definition is a *termination-sensitive* characterization of semantic typing, *i.e.* in which only terminating expressions are valid inhabitants of a type. The termination-insensitive alternative has the form  $e \Downarrow v \implies v \in \llbracket \tau \rrbracket$ , making any non-terminating expression a valid inhabitant of any type. We come back to termination issues later in this article.

$$\begin{aligned}
(v_1, v_2) \in \mathcal{V}_\Delta \llbracket B \cdot \Sigma \rrbracket &\stackrel{\text{def}}{=} v_i \in \llbracket B \rrbracket \wedge d_B(v_1, v_2) \leq \Sigma \cdot \Delta \\
(v_1, v_2) \in \mathcal{V}_\Delta \llbracket ((x : \tau_1) \rightarrow T_2) \cdot \Sigma \rrbracket &\stackrel{\text{def}}{=} \forall \Sigma_1, v'_1, v'_2 \text{ s.t. } (v'_1, v'_2) \in \mathcal{V}_\Delta \llbracket \tau_1 \cdot \Sigma_1 \rrbracket. (v_1 v'_1, v_2 v'_2) \in \mathcal{E}_\Delta \llbracket T_2[\Sigma_1/x] \rrbracket (+) \Sigma \\
(e_1, e_2) \in \mathcal{E}_\Delta \llbracket T \rrbracket &\stackrel{\text{def}}{=} e_i \Downarrow v_i \wedge (v_1, v_2) \in \mathcal{V}_\Delta \llbracket T \rrbracket \\
(\gamma_1, \gamma_2) \in \mathcal{S}_\Delta \llbracket \Gamma \rrbracket &\stackrel{\text{def}}{=} \forall x \in \text{dom}(\Gamma). (\gamma_1(x), \gamma_2(x)) \in \mathcal{V}_\Delta \llbracket \Gamma(x) \cdot x \rrbracket
\end{aligned}$$

Fig. 1: Logical relations for semantic sensitivity typing of  $\lambda_S$

(from Definition 1) is analogous to an open expression  $e$  with (possibly several) free variables. As such, the sensitivity is captured as a vector that maps each free variable to a sensitivity, denoted  $\Sigma$ , and written as  $s_1x_1 + \dots + s_nx_n$ . The input distance, previously captured by  $d$ , is now captured by a map from free variables to distances, denoted  $\Delta$ , also written as  $d_1x_1 + \dots + d_nx_n$ . The predicted output distance, previously computed as  $s \cdot d$ , is now computed as the point-wise product of the sensitivity environment and the distance environment, denoted as  $\Sigma \cdot \Delta$ , and defined as  $s_1 \cdot d_1 + \dots + s_n \cdot d_n$ . Finally, the actual inputs (previously  $a_1$  and  $a_2$ ) are now captured by two substitutions  $\gamma_1$  and  $\gamma_2$  that close the expression, and the actual outputs are captured by the results of applying the substitutions, written  $\gamma_1(e)$  and  $\gamma_2(e)$ . Proximity between substitutions is defined as follows:

$$d_\Gamma(\gamma_1, \gamma_2) \leq \Delta \stackrel{\text{def}}{=} \forall x \in \text{dom}(\Gamma). d_{\Gamma(x)}(\gamma_1(x), \gamma_2(x)) \leq \Delta(x)$$

With everything in place, we now define metric preservation, the sensitivity property for our lambda calculus:

**Definition 2 (Metric Preservation).** *Let  $\Gamma = x_1 : B_1, \dots, x_n : B_n$ . An open expression  $e$  satisfies metric preservation, written  $\text{MP}(\Gamma, e, B, \Sigma)$ , iff for all  $\Delta$  and  $\gamma_1, \gamma_2$ , such that  $d_\Gamma(\gamma_1, \gamma_2) \leq \Delta$ :*

$$\gamma_1(e) \Downarrow v_1 \wedge \gamma_2(e) \Downarrow v_2 \wedge d_B(v_1, v_2) \leq \Sigma \cdot \Delta$$

Observe the direct resemblance between Definitions 1 and 2.

### C. Semantics of Sensitivity Typing

As the first step for designing a sound sensitivity type system, we enrich the types of the core calculus with sensitivity information. We call this new calculus  $\lambda_S$ . We start by indexing the types with a sensitivity environment, written  $\tau \cdot \Sigma$ . The goal is to prove that if an open expression  $e$  has type  $\tau \cdot \Sigma$ , then it is metric-preserving with respect to  $\Sigma$ . The syntax for our *sensitivity types* is then<sup>5</sup>:

$$\tau ::= B \mid (x : \tau) \rightarrow T \quad T ::= \tau \cdot \Sigma$$

Notice that in an arrow  $(x : \tau) \rightarrow T$ , the codomain  $T$  may use the variable  $x$  in its top-level or any nested sensitivity environment. For example, a 2-sensitive function type on numbers is written  $(x : \mathbb{R}) \rightarrow \mathbb{R} \cdot 2x$ .

We now turn to the semantic model of sensitivity types. Like prior work, we rely on binary logical relations [5, 13, 10]; the

<sup>5</sup>A sensitivity type  $\tau \cdot \Sigma$  may also be seen as a type-and-effect, in the sense of [4], where the sensitivity environment  $\Sigma$  is the effect part.

binary nature comes from the fact that metric preservation characterizes the behavior of two runs of an expression. The relation (Figure 1) is defined using two mutually-defined interpretations: one for values  $\mathcal{V}_\Delta \llbracket T \rrbracket$ , and one for expressions  $\mathcal{E}_\Delta \llbracket T \rrbracket$ . These interpretations are indexed by a distance environment  $\Delta$ , which models the distance between inputs across two different executions, mirroring the definition of MP.

The interpretation of base types  $B \cdot \Sigma$  is straightforward: the values must be valid canonical forms of the base type  $B$  (denoted by  $\llbracket B \rrbracket$ ) and their distance must be bounded by the sensitivity environment  $\Sigma$ . Notice that this definition directly resembles the conclusion part of MP.

Function types  $((x : \tau_1) \rightarrow T_2) \cdot \Sigma$  are interpreted inductively: a pair of functions is related if for any pair of related inputs, the outputs are related. Remarkably, the outputs must be related at the type resulting of:

- 1) Replacing  $x$  in the codomain  $T_2$  by the actual sensitivity of the arguments,  $\Sigma_1$ . A sensitivity environment substitution,  $\Sigma[\Sigma_1/x]$ , replaces all occurrences of  $x$  in  $\Sigma$  by  $\Sigma_1$ . For example,  $(3x + 4y)[2y + z/x] = 3(2y + z) + 4y = 10y + 3z$ . Substitution on types is defined recursively.
- 2) Adding the function sensitivity  $\Sigma$ , which accounts for the variation between the functions themselves. The  $(+)$  meta-operator is defined as  $(\tau \cdot \Sigma_1) (+) \Sigma_2 = \tau \cdot \Sigma_1 + \Sigma_2$ . To illustrate the necessity of this operation, consider the expression  $x$  of type  $((y : \mathbb{R}) \rightarrow \mathbb{R} \cdot 0x) \cdot 1x$ , closed by  $\gamma_1 = \{x \mapsto \lambda y. 0\}$  and  $\gamma_2 = \{x \mapsto \lambda y. 1\}$ . Note how, if the sensitivity environment  $1x$  is not considered, the resulting type of applying  $\gamma_i(x)$  to any pair of arguments will be  $\mathbb{R} \cdot \emptyset$ , but  $d_{\mathbb{R}}(0, 1) = 1 \not\leq \emptyset \cdot \Delta = 0$ .

Two closed expressions are related if, when evaluated, their results are related according to the value interpretation of the same type. Finally, two substitutions are related if they provide related values for each variable in the type environment.

We define *semantic sensitivity typing*,  $\Gamma \models_S e : T$ , to denote that expression  $e$  reduces to close enough—as described by  $T$ —values, when closed by two neighboring substitutions:

**Definition 3 (Semantic Sensitivity Typing).**

$$\Gamma \models_S e : T \stackrel{\text{def}}{=} \forall \Delta, \gamma_1, \gamma_2 \text{ s.t. } (\gamma_1, \gamma_2) \in \mathcal{S}_\Delta \llbracket \Gamma \rrbracket. (\gamma_1(e), \gamma_2(e)) \in \mathcal{E}_\Delta \llbracket T \rrbracket$$

We can then prove that semantic sensitivity typing implies metric preservation:

$$\begin{array}{c}
\text{(SCONST)} \\
\frac{c \in \llbracket B \rrbracket}{\Gamma \vdash_S c : B \cdot \emptyset} \\
\\
\text{(SVAR)} \\
\frac{x : \tau \in \Gamma}{\Gamma \vdash_S x : \tau \cdot x} \\
\\
\text{(SAPP)} \\
\frac{\Gamma \vdash_S e_1 : ((x : \tau_1) \rightarrow T_2) \cdot \Sigma \quad \Gamma \vdash_S e_2 : \tau_1 \cdot \Sigma_1}{\Gamma \vdash_S e_1 e_2 : T_2[\Sigma_1/x] \text{ (+)} \Sigma} \\
\\
\text{(SOP)} \\
\frac{\Gamma \vdash_S e_i : T_i \quad \delta_{\text{op}}(\overline{T_i}) = T}{\Gamma \vdash_S \text{op } \overline{e_i} : T} \\
\\
\text{(SLAM)} \\
\frac{\Gamma, x : \tau_1 \vdash_S e : T_2}{\Gamma \vdash_S \lambda(x : \tau_1).e : ((x : \tau_1) \rightarrow T_2) \cdot \emptyset} \\
\\
\text{(SSUB)} \\
\frac{\Gamma \vdash_S e : T \quad T <: T'}{\Gamma \vdash_S e : T'}
\end{array}$$

$$\begin{aligned}
\delta_+(\mathbb{R} \cdot \Sigma_1, \mathbb{R} \cdot \Sigma_2) &= \mathbb{R} \cdot \Sigma_1 + \Sigma_2 \\
\delta_*(\mathbb{R} \cdot \Sigma_1, \mathbb{R} \cdot \Sigma_2) &= \mathbb{R} \cdot \infty(\Sigma_1 + \Sigma_2) \\
\delta_{\leq}(\mathbb{R} \cdot \Sigma_1, \mathbb{R} \cdot \Sigma_2) &= \mathbb{B} \cdot \infty(\Sigma_1 + \Sigma_2)
\end{aligned}$$

Fig. 2: Syntactic Sensitivity Typing for  $\lambda_S$

**Lemma 1** (Soundness of Semantic Sensitivity Typing). *Let  $\Gamma = x_1 : B_1, \dots, x_n : B_n$ .*

$$\Gamma \vDash_S e : B \cdot \Sigma \implies \text{MP}(\Gamma, e, B, \Sigma)$$

As a final step, Figure 2 presents a *syntactic* sensitivity type system, which is sound with respect to metric preservation. Constants are typed with an empty sensitivity environment, as they are pure, *i.e.* they do not depend on any variable. Primitive n-ary operations  $\text{op}$  are given meaning by the function  $\delta_{\text{op}}$ , which handles the treatment of sensitivity information. For a few examples, definitions of  $\delta$  are given for addition, multiplication, and comparison. Multiplication scales the resulting sensitivity by infinite, as the operation is not bounded. Comparison is similar, but the resulting type is boolean; a small variation in its operands can change the resulting boolean value, hence the infinite sensitivity. Variables are typed getting their type from the type environment, and their sensitivity is the variable itself, *i.e.* they are 1-sensitive on themselves. As standard, the body of a lambda is typed under an extended type environment and, same as constants, the sensitivity of a lambda is pure. For function applications, we mirror the definition of related functions in the logical relation. Finally, we admit a subsumption rule that allows an expression to be typed with a supertype of its actual type. Subtyping is naturally induced by the order of sensitivities, *i.e.*  $(x : \mathbb{R}) \rightarrow \mathbb{R} \cdot 2x <: (x : \mathbb{R}) \rightarrow \mathbb{R} \cdot 3x$  because  $2 \leq 3$ .

We can prove that syntactic typing implies semantic typing:

**Lemma 2** (Syntactic Typing Implies Semantic Typing).

$$\Gamma \vdash_S e : T \implies \Gamma \vDash_S e : T$$

Finally, combining the Lemmas 1 and 2 we establish that syntactic sensitivity typing implies metric preservation:

**Theorem 3** (Soundness of Syntactic Sensitivity Typing). *Let  $\Gamma = x_1 : B_1, \dots, x_n : B_n$ .*

$$\Gamma \vdash_S e : B \cdot \Sigma \implies \text{MP}(\Gamma, e, B, \Sigma)$$

$$\begin{aligned}
s &::= [s_1, s_2] \quad \text{with: } ? \triangleq [0, \infty], \quad s \triangleq [s, s] \\
\Sigma &::= s x + \dots + s x \\
\tau &::= B \mid T \rightarrow T \mid \forall x. T \\
T &::= \tau \cdot \Sigma \\
e &::= c \mid \text{op } \overline{e} \mid x \mid \lambda(x : T).e \mid e e \mid \Lambda x. e \mid e[\Sigma] \mid e :: T
\end{aligned}$$

Fig. 3: Syntax of GS

#### D. Metric Preservation and Non-Termination

So far we have worked under the assumption that expressions are terminating. While this may be true for the  $\lambda_S$ , it is not the case for more expressive languages; certainly not for gradual languages, such as GSoul, where the possibility of errors is first-class. To conclude this section, we adapt MP to a weaker, albeit more general, notion of metric preservation that accounts for non-termination, EMP:

**Definition 4** (Exceptional Metric Preservation). *Let  $\Gamma = x_1 : B_1, \dots, x_n : B_n$ . An open expression  $e$  satisfies exceptional metric preservation, written  $\text{EMP}(\Gamma, e, \tau, \Sigma)$ , iff for all  $\Delta$  and  $\gamma_1, \gamma_2$ , such that  $d_\Gamma(\gamma_1, \gamma_2) \leq \Delta$ :*

$$(\gamma_1(e) \Downarrow v_1 \wedge \gamma_2(e) \Downarrow v_2) \implies d_\tau(v_1, v_2) \leq \Sigma \cdot \Delta$$

Notice that this definition is termination-insensitive, as it does not require that the expressions share the same termination behavior. Naturally, a termination-sensitive variant could be defined by moving  $\gamma_2(e) \Downarrow v_2$  to the conclusion of the implication, *i.e.*  $\gamma_1(e) \Downarrow v_1 \implies (\gamma_2(e) \Downarrow v_2 \wedge d_\tau(v_1, v_2) \leq \Sigma \cdot \Delta)$ . However, this requires a more subtle analysis of the language semantics. We come back to this issue in §V-C.

This concludes the presentation of the blueprint we follow in the remainder of this article to establish the semantics and metatheory of *gradual* sensitivity typing.

#### IV. GS: A GRADUAL SENSITIVITY LANGUAGE

In this section, we introduce GS, the core calculus that serves as the formal foundation of GSoul. Figure 3 shows the syntax of types and expressions in GS, which presents 3 remarkable differences from  $\lambda_S$ : (1) we extend the syntax of types with a new type constructor  $\forall x. T$ , (2) we introduce gradual sensitivities  $s$  and (3) we extend expressions with ascriptions  $e :: T$ . We discuss each of these in turn.

##### A. Distance Variables for Metric Preservation

Across §III we explored the semantics of sensitivity modeling variations through free variables in expressions. While this is convenient for reasoning about semantics, it may not be practical for a programming language, as it treats all variables as resources worth tracking, which consequently adds a significant amount of information to types. GSoul instead allows the programmer to specify which variables should be tracked as resources, *e.g.* consider a simple function:

```
def foo(a: Number, res b: Number): Number[2b] =
  a + b + b;
```

The **res** modifier in the function declaration specifies that the variable  $b$  should be tracked as a resource. Consequently, the return type can be annotated with the sensitivity of the body with respect to  $b$ . Observe that the first argument  $a$  is not tracked for sensitivity, because it is not marked as a resource.

In order to model this specification, let us revisit the syntax of function types for  $\lambda_S, (x : \tau_1) \rightarrow T_2$ . Here, the sensitivity information conveyed in the codomain of the function,  $T_2$ , is dependent on the name of the argument,  $x$ , hence the need of specifying the name of the variable in the domain—as in  $(x : \tau_1)$ . For GS, instead of having all functions treat their argument as resources, we introduce a more general construct that decouples the role of  $(x : \_)$  from function types:  $\forall x.T$ , which abstracts a sensitivity type  $T$  over a *distance variable*  $x$ . Distance variables  $x$  (written in blue roman typestyle) differ from lambda-bound variables  $x$  in that they live at the type level, and are used solely to describe the sensitivity of computations. Then, the type  $(x : \tau_1) \rightarrow T_2$  can be now written as  $\forall x.\tau_1.x \rightarrow T_2$ , noting that the syntax of function types is now simply  $T_1 \rightarrow T_2$ . Accordingly, we introduce the distance abstraction term  $\lambda x.e$ . A distance abstraction is eliminated by instantiating the distance variable with a concrete sensitivity environment, using the distance application form  $e[\Sigma]$ . This is akin to explicit parametric polymorphism [32].

Note that the **res** modifier can be implemented simply as syntactic sugar for a distance abstraction. For instance, the function  $\lambda x.\lambda(n : \mathbb{R}.x).n + n$  is the same as the following function—using concrete syntax and the **res** modifier:

```
def double(res n: Number): Number[2n] = n + n;
```

In fact, the **res** modifier is implemented as a desugaring step in the parser, which transforms the function into a distance abstraction as follows:

```
def double[x](n: Number[1x]): Number[2x] = n + n;
```

Here, the distance variable  $x$  is introduced and used to parameterize the distance of the input variable  $n$ . Furthermore, when applying a function that was declared using the **res** modifier, GSoul does not require to explicitly eliminate the distance variable using a concrete sensitivity environment, as it is done automatically by using the inferred sensitivity environment from the actual argument (similarly to how a practical polymorphic language like Haskell dispenses the programmer from explicitly providing type applications).

## B. Imprecise Sensitivity Information

*Gradual sensitivities* represent *imprecise* sensitivity information that is optimistically handled by the type system. A gradual sensitivity  $s$  is defined as a valid interval of two static sensitivities, capturing the plausibility of a sensitivity being any number within the range (Figure 3).<sup>6</sup> Consequently, the unknown sensitivity  $?$  is simply sugar for the interval  $[0, \infty]$ , and a *fully precise* sensitivity  $s$  is sugar for the interval  $[s, s]$ . Precision in the context of gradual sensitivities corresponds

<sup>6</sup>By extension, types, expressions and type environments may be gradual, though we retain the same meta-variables from the syntax presented earlier.

to interval inclusion. Formally,  $[s_1, s_2]$  is more precise than  $[s_3, s_4]$ , written  $[s_1, s_2] \sqsubseteq [s_3, s_4]$ , if and only if  $s_1 \geq s_3$  and  $s_2 \leq s_4$ . Then, as one would expect, the unknown sensitivity is the least precise sensitivity and, conversely, a sensitivity  $[s, s]$  is said to be fully precise, as no other sensitivity is more precise than itself but the same one.

Sensitivity precision naturally induces a notion of precision on types, e.g.  $\mathbb{R}.[5, 10]x \rightarrow \mathbb{R}.[10, 10]x \sqsubseteq \mathbb{R}.[0, 10]x \rightarrow \mathbb{R}.[5, 15]x$ , because  $[5, 10] \sqsubseteq [0, 10]$  and  $[10, 10] \sqsubseteq [5, 15]$ . Type precision captures the idea of a type containing more precise sensitivity information than another type, hence it corresponds to checking the point-wise precision of the sensitivities in the involved types. Consequently, type precision in gradual typing is covariant in the domain of arrows [16, 29].

Plausibility for gradual sensitivities arises from relaxing the order of sensitivities, which yields an optimistic relation for gradual sensitivities, denoted by  $\lesssim$ . Intuitively, a gradual sensitivity  $s_1$  is less than or equal to another,  $s_2$ , if there *exists* a sensitivity within  $s_1$  that is less or equal to another one within  $s_2$ . This is easily checkable by comparing the lower bound of  $s_1$  with the upper bound of  $s_2$ , which corresponds to the best-case scenario. Formally, we define:

**Definition 5** (Consistent Sensitivity Ordering).  $[s_1, s_2] \lesssim [s_3, s_4]$  if and only if  $s_1 \leq s_4$ .

To illustrate the plausible nature of this relation, note that  $5 \lesssim [0, 10]$  holds, even though there are sensitivities within  $[0, 10]$  that are not less than or equal to 5. Conversely, we still have that  $10 \not\lesssim [0, 5]$ , because there is no sensitivity within  $[0, 5]$  that is greater than or equal to 10. Lastly, notice that this definition satisfies the behavior described in §III for the unknown sensitivity:  $? \lesssim s$  and  $s \lesssim ?$  for any  $s$ .

One crucial aspect of consistent sensitivity order is that, in general, the transitivity of two consistent judgments is not guaranteed to hold.<sup>7</sup> For instance, although  $10 \lesssim ?$  and  $? \lesssim 5$ , the transitive judgement of both,  $10 \lesssim 5$ , does not hold. Nevertheless, in some cases the transitivity of two consistent judgments is indeed satisfied, e.g. with  $5 \lesssim ?$  and  $? \lesssim 10$ , we indeed have that  $5 \lesssim 10$ . Just like with static sensitivities, the order of gradual sensitivities induces an order on sensitivity environments and types. We use  $\lesssim$  to denote *consistent subtyping*, which is the gradual counterpart of subtyping for gradual sensitivities. For example,  $\mathbb{R}.?x_1 \rightarrow \mathbb{R}.10x_2 \lesssim \mathbb{R}.5x_1 \rightarrow \mathbb{R}.?x_2$ , because  $5 \lesssim ?$  and  $10 \lesssim ?$ . Consistent subtyping is defined in Figure 4.

## C. Ascriptions

Lastly, we extend expressions with type ascriptions  $e :: T$ , which represent a syntax-directed subsumption assertion. Well-typedness (either semantic or syntactic) of an ascription,

<sup>7</sup>Note that the consistent lifting of a transitive relation is not generally transitive [16]. Therefore, the consistent lifting of a (pre)order is not itself a (pre)order. Yet the convention for consistent liftings is to preserve the name of the original lifted relation, qualifying it with “consistent” (such as *consistent subtyping* [16, 17]). So, *consistent sensitivity ordering* is not a preorder, just like *consistent label ordering* in gradual security typing [23] is not.

$$\begin{array}{ccc}
\text{(TSUB-B)} & \text{(TSUB-ARROW)} & \text{(TSUB-FORALL)} \\
\frac{}{B \lesssim B} & \frac{T_{21} \lesssim T_{11} \quad T_{12} \lesssim T_{22}}{T_{11} \rightarrow T_{12} \lesssim T_{21} \rightarrow T_{22}} & \frac{T_1 \lesssim T_2}{\forall x. T_1 \lesssim \forall x. T_2}
\end{array}$$

Fig. 4: Consistent subtyping

in a gradual setting, is given by the consistent subtyping relation, e.g. if  $e$  has type  $T_1$  and  $T_1 \lesssim T_2$ , then  $e :: T_2$  should be well-typed.

Given that consistent subtyping is induced by the consistent sensitivity order, transitivity of two (consistent) subtyping judgments is also not guaranteed, i.e. if  $T_1 \lesssim T_2$  and  $T_2 \lesssim T_3$ , then  $T_1 \lesssim T_3$  may not hold. This behavior is key to allow for plausibility in a gradual language. Conversely, this lack of general transitivity for consistent subtyping judgments, implies that ascription must be dealt with carefully in the reduction semantics. Whereas in a static typing setting, type ascriptions only serve the purpose of weakening the type of an expression—thus being a no-op in the reduction semantics—, in a gradual setting, ascriptions may be used to optimistically lower the sensitivity of an expression, which if necessary, in order to avoid violating the sensitivity guarantees of the language, should produce an error during runtime. To illustrate, suppose  $e$  is semantically 2-sensitive on some  $x$  and, consequently, has type  $\mathbb{R} \cdot 2x$ . Then,  $e' = e :: \mathbb{R} \cdot ?x :: \mathbb{R} \cdot 1x$  should be well-typed, since  $2 \lesssim ?$  and  $? \lesssim 1$ . Nonetheless,  $e'$  is clearly not 1-sensitive, and thus should produce an error during runtime. Otherwise, two runs of  $e'$  would eventually produce results further than what is described by the sensitivity type. We present a mechanism to perform such checks in §VI.

## V. SEMANTICS OF GRADUAL SENSITIVITY TYPING

We now define the semantic model for GS types for which, just like in §III, we use a binary logical relation. Before we delve into the details, we write  $e \Downarrow v$  and  $e \Downarrow \mathbf{error}$  to denote that  $e$  reduces to the value  $v$  or an error, respectively.

### A. Distance Reasoning for Gradual Sensitivities

Let us start by defining  $\mathcal{V}_\Delta \llbracket B \cdot \Sigma \rrbracket$ . Whereas in a static setting, the output distance is directly computable (from  $\Delta$  and  $\Sigma$ ), in a gradual setting, imprecision makes it less obvious how to compare two values. Moreover, given the optimism introduced by gradual sensitivities, the actual sensitivity of an expression can only be observed after performing the operations that lead to a value. For instance, the constant 2 obtained from  $x + 1$  is different to that obtained from  $x + x$ , where  $x = 1$ ; the former is 1-sensitive on  $x$ , while the latter is 2-sensitive. Nevertheless, both expressions could be ascribed to  $\mathbb{R} \cdot ?x$ , hiding the actual sensitivity of the expressions in the type. Strictly speaking, the information conveyed in  $\Sigma$  does not provide a precise enough sensitivity to compare two values; it has not been refined by the actual operations performed. The key insight for defining a logical relation for gradual sensitivity typing, is the necessity of monitoring the sensitivity

of expressions during evaluation, and using this information to compare values.

*A monitor for gradual sensitivities:* Based on the operations performed to reach a value  $v$  from an initial expression  $e$ ,  $v$  should contain the *monitored* sensitivity information of  $e$ . We do not prescribe here a specific representation of monitors, but instead only specify its interface. We provide concrete definitions in the next section. First of all, it should be possible to obtain the monitor of a value, with the metafunction  $\text{mon}$ . This monitor should contain (1) any necessary information to perform runtime sensitivity checks, and (2) the actual monitored sensitivity of the computation that produced the value. Because of the second point, another metafunction  $\text{ms}$  can be used to extract a fully-precise sensitivity environment from a monitor, which should correspond to the least conservative approximation of the actual sensitivity of the initial expression, based on the operations performed to reach such value. Moreover, we require monitors to be adequate with respect to the initial sensitivity typing information of expressions, in the sense that if  $e$  has type  $\tau \cdot \Sigma$  and  $e \Downarrow v$ , then  $\text{ms}(\text{mon}(v)) \sqsubseteq \Sigma$ .

### B. Logical relations for Gradual Sensitivity Typing

Figure 5 shows the definition of the logical relation for gradual sensitivity typing. As described before, for base types, we obtain the monitored sensitivity environments from the values, computing  $\Sigma_i = \text{ms}(\text{mon}(v_i))$ . Since the two values may report different sensitivities, we need to join these sensitivities to compute a conservative distance between the values. Hence, the condition over the distance of values is defined as  $d_B(v_1, v_2) \leq (\Sigma_1 \Upsilon \Sigma_2) \cdot \Delta$ .

Since we have re-located the responsibility of introducing resources from lambdas to distance abstractions, the case for related functions is simpler than its  $\lambda_S$  counterpart (Figure 1); related arguments produce related results, without the need to eliminate any variable from the types. In turn, the case for related type abstractions considers the sensitivity environment argument, and eliminates the distance variable from the type. Finally, related computations now consider the possibility of errors during evaluation: if the two expressions reduce to values, then they should be related according to the definition of related base values, otherwise the relation holds vacuously, similar to the definition of EMP. All other aspects are analogous to the ones presented in §III.

We define semantic gradual sensitivity typing just as before:

**Definition 6** (Semantic Gradual Sensitivity Typing).

$$\Gamma \vDash e : T \stackrel{\text{def}}{=} \forall \Delta, \gamma_1, \gamma_2 \text{ s.t. } (\gamma_1, \gamma_2) \in \mathcal{S}_\Delta \llbracket \Gamma \rrbracket. \\
(\gamma_1(e), \gamma_2(e)) \in \mathcal{E}_\Delta \llbracket T \rrbracket$$

Given that in the gradual setting we type expressions under plausibility-based constraints, soundness of semantic typing cannot be established for all possible sensitivities within the ranges of a gradual sensitivity environment  $\Sigma$ . Nevertheless, the observed output distance should still be bounded by  $\Sigma' \cdot \Delta$ , for some fully-precise  $\Sigma'$  that is more precise than  $\Sigma$ . We define GMP, EMP for gradual sensitivities, as follows:

$$\begin{aligned}
(v_1, v_2) \in \mathcal{V}_\Delta \llbracket B \cdot \Sigma \rrbracket &\stackrel{\text{def}}{=} v_i \in \llbracket B \rrbracket \wedge d_B(v_1, v_2) \leq (\Sigma_1 \vee \Sigma_2) \cdot \Delta \quad \text{where } \Sigma_i = \text{ms}(\text{mon}(v_i)) \\
(v_1, v_2) \in \mathcal{V}_\Delta \llbracket (T_1 \rightarrow T_2) \cdot \Sigma \rrbracket &\stackrel{\text{def}}{=} \forall v'_1, v'_2, (v'_1, v'_2) \in \mathcal{V}_\Delta \llbracket T_1 \rrbracket. (v_1 v'_1, v_2 v'_2) \in \mathcal{E}_\Delta \llbracket T_2 (+) \Sigma \rrbracket \\
(v_1, v_2) \in \mathcal{V}_\Delta \llbracket (\forall x. T) \cdot \Sigma \rrbracket &\stackrel{\text{def}}{=} \forall \Sigma'' . (v_1 \llbracket \Sigma'' \rrbracket, v_2 \llbracket \Sigma'' \rrbracket) \in \mathcal{E}_\Delta \llbracket [\Sigma''/x]T (+) \Sigma \rrbracket \\
(e_1, e_2) \in \mathcal{E}_\Delta \llbracket T \rrbracket &\stackrel{\text{def}}{=} (e_1 \Downarrow v_1 \wedge e_2 \Downarrow v_2) \implies (v_1, v_2) \in \mathcal{V}_\Delta \llbracket T \rrbracket \\
(\gamma_1, \gamma_2) \in \mathcal{S}_\Delta \llbracket \Gamma \rrbracket &\stackrel{\text{def}}{=} \forall x \in \text{dom}(\Gamma). (\gamma_1(x), \gamma_2(x)) \in \mathcal{V}_\Delta \llbracket \Gamma(x) \rrbracket
\end{aligned}$$

Fig. 5: Logical relations for semantic gradual sensitivity typing

**Definition 7** (Gradual Metric Preservation). *Let  $\Gamma = x_1 : B_1 \cdot \Sigma_1, \dots, x_n : B_n \cdot \Sigma_n$ . An open expression  $e$  satisfies gradual metric preservation, written  $\text{GMP}(\Gamma, e, B, \Sigma)$ , iff for all  $\Delta$  and  $\gamma_1, \gamma_2$  such that  $d_\Gamma(\gamma_1, \gamma_2) \leq \Delta$ :*

$$\begin{aligned}
(\gamma_1(e) \Downarrow v_1 \wedge \gamma_2(e) \Downarrow v_2) &\implies \\
\exists \Sigma' \text{ s.t. } \text{static}(\Sigma') \wedge \Sigma' \sqsubseteq \Sigma. d_B(v_1, v_2) &\leq \Sigma' \cdot \Delta
\end{aligned}$$

The predicate  $\text{static}(\Sigma)$  holds when all gradual sensitivities in  $\Sigma$  are fully precise. We then prove that the semantic typing is sound with respect to the gradual metric preservation:

**Lemma 4** (Semantic Gradual Sensitivity Typing is Sound). *Let  $\Gamma = x_1 : B_1 \cdot \Sigma_1, \dots, x_n : B_n \cdot \Sigma_n$ .*

$$\Gamma \vDash e : B \cdot \Sigma \implies \text{GMP}(\Gamma, e, B, \Sigma)$$

### C. Termination-sensitive metric preservation

The notion of gradual metric preservation presented above (Def. 7) is *termination insensitive*: if either of the two executions fail, the property vacuously holds. This property is weaker than the one originally proposed by Reed and Pierce [5] for their static sensitivity type system. We now study how to achieve a stronger, *termination sensitive* characterization of gradual metric preservation, *i.e.* either both executions fail or both reduce to values at a bounded distance. We show that such a stronger property is achievable if the imprecision of sensitivities is bounded (*i.e.* not infinite).

A first observation towards enforcing the same termination behavior across runs is to realize that two substitutions  $\gamma_1$  and  $\gamma_2$  may introduce values with different levels of precision that may cause an error in only one of the executions. To prevent this, we take inspiration on prior work [33] and restrict the monitors of related values to be *equi-precise*, denoted by  $\sqsupseteq$ . Suppose that  $m_i = \text{mon}(v_i)$ , then  $m_1 \sqsupseteq m_2$  if and only if  $m_1 \sqsubseteq m_2$  and  $m_2 \sqsubseteq m_1$ <sup>8</sup>. Second, we draw from D’Antoni et al. [11] and Azevedo de Amorim et al. [28], who show, in their respective calculi, that when restricting the predicted output distance to be finite, the termination behavior of executions for a given expression is uniform. We define the following predicate to check if a distance environment (or a sensitivity environment) is bounded, *i.e.* does not contain  $\infty$ :

$$\text{bounded}(\Delta) \stackrel{\text{def}}{=} \forall x \in \text{dom}(\Delta). \infty \not\sqsubseteq \Delta(x)$$

<sup>8</sup>Notice that the ability to compare monitors under a precision relation is only necessary for the termination-sensitive logical relations.

Figure 6 presents the termination-sensitive logical relations for GS types, denoted  $\mathcal{V}_\Delta^{\text{ts}} \llbracket T \rrbracket$ ,  $\mathcal{E}_\Delta^{\text{ts}} \llbracket T \rrbracket$ ,  $\mathcal{S}_\Delta^{\text{ts}} \llbracket \Gamma \rrbracket$ . The main differences with respect to their termination-insensitive counterparts are: (1) all definitions of related values now include a check for the equi-precision of their monitors, (2) related computations require the same termination behavior, and (3) related substitutions require the inputs distance to be finite. Additionally, the case for related distance abstractions now requires the sensitivity environment argument to be bounded, *i.e.* bounded( $\Sigma''$ ). Then, we define semantic termination-sensitive typing as follows:

**Definition 8** (Semantic (Termination-Sensitive) Gradual Sensitivity Typing).

$$\begin{aligned}
\Gamma \vDash_{\text{ts}} e : \tau \cdot \Sigma &\stackrel{\text{def}}{=} \text{bounded}(\Sigma) \wedge \\
&\forall \Delta, \gamma_1, \gamma_2 \text{ s.t. } (\gamma_1, \gamma_2) \in \mathcal{S}_\Delta^{\text{ts}} \llbracket \Gamma \rrbracket. \\
&(\gamma_1(e), \gamma_2(e)) \in \mathcal{E}_\Delta^{\text{ts}} \llbracket \tau \cdot \Sigma \rrbracket
\end{aligned}$$

Besides relying on the termination-sensitive logical relations, Definition 8 also requires that the top-level sensitivity environment is bounded. This latter requirement, in conjunction with the condition over  $\Delta$  in the definition of related substitutions, ensures that the output distance is not plausibly infinite, obtaining the same termination behavior across runs. Moreover, notice that this criterion is modular, *i.e.*  $e$  may use any form of imprecision internally, as long as the top-level sensitivity environment is bounded. This definition illustrates how the progressive hardening of precision may allow a programmer to obtain stronger guarantees about their programs.

Finally, following the same ideas, we can define a termination-sensitive variant of GMP:

**Definition 9** ((Termination-Sensitive) Gradual Metric Preservation). *Let  $\Gamma = x_1 : B_1 \cdot \Sigma_1, \dots, x_n : B_n \cdot \Sigma_n$ . An open expression  $e$  satisfies termination-sensitive gradual metric preservation, written  $\text{GMP}(\Gamma, e, B, \Sigma)$ , iff for all  $\Delta$  and  $\gamma_1, \gamma_2$  such that  $d_\Gamma^{\text{ts}}(\gamma_1, \gamma_2) \leq \Delta$ :*

$$\begin{aligned}
\gamma_1(e) \Downarrow v_1 &\implies \gamma_2(e) \Downarrow v_2 \wedge \\
\exists \Sigma' \text{ s.t. } \text{static}(\Sigma') \wedge \Sigma' \sqsubseteq \Sigma. d_B(v_1, v_2) &\leq \Sigma' \cdot \Delta
\end{aligned}$$

The new definition of substitutions proximity,  $d_\Gamma^{\text{ts}}(\cdot, \cdot) \leq \Delta$ , analogous to  $\mathcal{S}_\Delta \llbracket \Gamma \rrbracket$ , requires  $\Delta$  to be bounded and related values to have equi-precise monitors. Then, we have the following lemma:

$$\begin{aligned}
(v_1, v_2) \in \mathcal{V}_\Delta^{\text{ts}}[[B \cdot \Sigma]] &\stackrel{\text{def}}{=} \text{mon}(v_1) \sqsupseteq \text{mon}(v_2) \wedge v_i \in [[B]] \wedge d_B(v_1, v_2) \leq (\Sigma_1 \vee \Sigma_2) \cdot \Delta \quad \text{where } \Sigma_i = \text{ms}(\text{mon}(v_i)) \\
(v_1, v_2) \in \mathcal{V}_\Delta^{\text{ts}}[(T_1 \rightarrow T_2) \cdot \Sigma] &\stackrel{\text{def}}{=} \text{mon}(v_1) \sqsupseteq \text{mon}(v_2) \wedge \forall v'_1, v'_2, (v'_1, v'_2) \in \mathcal{V}_\Delta^{\text{ts}}[T_1]. (v_1 v'_1, v_2 v'_2) \in \mathcal{E}_\Delta^{\text{ts}}[T_2 (+) \Sigma] \\
(v_1, v_2) \in \mathcal{V}_\Delta^{\text{ts}}[(\forall x. T) \cdot \Sigma] &\stackrel{\text{def}}{=} \text{mon}(v_1) \sqsupseteq \text{mon}(v_2) \wedge \forall \Sigma'' \text{ s.t. } \text{bounded}(\Sigma''). (v_1 [\Sigma''], v_2 [\Sigma'']) \in \mathcal{E}_\Delta^{\text{ts}}[[\Sigma''/x]T (+) \Sigma] \\
(e_1, e_2) \in \mathcal{E}_\Delta^{\text{ts}}[T] &\stackrel{\text{def}}{=} e_1 \Downarrow v_1 \implies e_2 \Downarrow v_2 \wedge (v_1, v_2) \in \mathcal{V}_\Delta^{\text{ts}}[T] \\
(\gamma_1, \gamma_2) \in \mathcal{S}_\Delta^{\text{ts}}[\Gamma] &\stackrel{\text{def}}{=} \text{bounded}(\Delta) \wedge \forall x \in \text{dom}(\Gamma). (\gamma_1(x), \gamma_2(x)) \in \mathcal{V}_\Delta^{\text{ts}}[\Gamma(x)]
\end{aligned}$$

Fig. 6: Logical relations for semantic termination-sensitive gradual sensitivity typing

**Lemma 5** (Soundness of (Termination-Sensitive) Gradual Semantic Typing). *Let  $\Gamma = x_1 : B_1 \cdot \Sigma_1, \dots, x_n : B_n \cdot \Sigma_n$ .*

$$\Gamma \models_{\text{ts}} e : B \cdot \Sigma \implies \text{GMP}_{\text{ts}}(\Gamma, e, B, \Sigma)$$

To conclude this section, we put our attention on the runtime semantics of GS, which we have left abstract so far. We know that the runtime of GS must be able to monitor the sensitivity of expressions during their evaluation, and use this information to both perform runtime checks and provide a sound approximation of the actual sensitivity of the expressions. In the next sections, we tackle this specification in two steps: in §VI we present an intermediate language, variant of GS, that is sound with respect to the semantic typing defined in this section; and, in §VII, we define a syntax-directed and type-preserving translation from GS to this intermediate language, which also plays the role of a syntactic type system for GS.

## VI. RUNTIME SEMANTICS FOR GS

To define the runtime semantics of GS, we follow the general recipe of gradual languages, where source language expressions are statically typechecked and elaborated to an internal language with explicit runtime checks that account for the optimism that arises from imprecise type information [15, 29, 16]. Specifically, we follow the general approach of Garcia et al. [16], which refines the initial justification of optimistic judgments in expressions as they reduce, yielding *type-driven* monitors for gradual programs. We first describe this approach in our setting of imprecise sensitivity information in §VI-A, and then define the internal language GSCHECK in §VI-B. §VII then describes the syntactic typechecking and elaboration of GS to GSCHECK, and proves that any syntactically well-typed GS term elaborates to a semantically well-typed GSCHECK term that therefore satisfies (possibly termination-sensitive) metric preservation, as well as the properties expected from gradual languages [29].

### A. Evidence for Consistent Sensitivity Subtyping

To systematically gradualize typing disciplines, the Abstracting Gradual Typing approach [16] recognizes that, in the presence of consistent judgments—subtyping, in our case—evaluation must be concerned about not only *that* an optimistic judgment holds, but also *why* it holds, locally. During reduction, these local justifications need to be combined to keep justifying the plausibility of subsequent steps, or fail otherwise to report a violation of the optimistic assumptions.

To illustrate, suppose we have  $\cdot \models e : \mathbb{R} \cdot 2x$  and we want to evaluate  $e' = e :: \mathbb{R} \cdot ?x :: \mathbb{R} \cdot 1x$ . Observe that (i) each ascription in  $e'$  is *locally justified* thanks to the use of  $?x$  in the first ascription, and (ii) evaluating  $e'$  should produce an error at runtime, given that the sensitivity of  $e$  is  $2x$ , not  $1x$ .

In order to achieve this behavior, we must refute the transitivity of  $\mathbb{R} \cdot 2x \lesssim \mathbb{R} \cdot ?x$  and  $\mathbb{R} \cdot ?x \lesssim \mathbb{R} \cdot 1x$ . However, given that each judgment holds individually, this is only possible if we know that the actual sensitivity cannot be *simultaneously* greater than 2 and less than 1: we care about the individual proofs of the consistent subtyping judgments and how they can be combined (or not) to justify their transitivity. To this end, Garcia et al. [16] introduce the notion of *evidences*  $\varepsilon$  that represent the proofs of consistent subtyping judgments. We write  $\varepsilon \triangleright T_1 \lesssim T_2$  to denote that the evidence  $\varepsilon$  *justifies* that  $T_1$  is a consistent subtype of  $T_2$ . During reduction, one can compute whether the transitivity of two consistent subtyping judgments is justified using *consistent transitivity* (noted  $\circ$ ) of their respective evidences: given  $\varepsilon_1 \triangleright T_1 \lesssim T_2$  and  $\varepsilon_2 \triangleright T_2 \lesssim T_3$ , if  $\varepsilon_1 \circ \varepsilon_2$  is defined, then  $\varepsilon_1 \circ \varepsilon_2 \triangleright T_1 \lesssim T_3$ . If the combination is undefined, then a runtime error must be raised, as the transitivity has been effectively refuted.

We now formally define evidences and their operations.

a) *Evidences*: In the presence of consistent subtyping [16, 23], an evidence  $\varepsilon$  is typically represented as a pair of sensitivity types, written  $\langle T_1, T_2 \rangle$ , each of which being at least as precise as the types involved in the consistent subtyping relation. Intuitively, the types in the evidence can be seen as refined versions of the types in the consistent subtyping judgment, each of them a proof that the judgment holds.

b) *Interior*: Evidence is initially computed from a consistent subtyping judgment using the *interior* operator, which essentially produces a refined pair of gradual types, based solely on the knowledge that the judgment holds. Naturally, if the judgment does not hold, the interior operator is undefined.

The interior for sensitivity subtyping is naturally defined as the lifting of the interior for consistent sensitivity ordering:

**Definition 10** (Interior).

$$\mathcal{I}([s_1, s_2], [s_3, s_4]) = \langle [s_1, s_2 \wedge s_4], [s_1 \vee s_3, s_4] \rangle$$

If  $s_1 \leq s_2 \wedge s_4$  and  $s_1 \vee s_3 \leq s_4$ . Otherwise, the interior is undefined.

To illustrate, suppose we have  $T_1 = \mathbb{R} \cdot ?x$  and  $T_2 = \mathbb{R} \cdot 10x$ , noting that  $T_1 \lesssim T_2$ . We compute the initial evidence for that

judgment as  $\mathcal{I}(T_1, T_2) = (\mathbb{R} \cdot [0, 10]x, \mathbb{R} \cdot [10, 10]x)$ . Notice how the evidence is (component-wise) more precise than the types involved in the consistent subtyping judgment. In particular, the range  $[0, \infty]$  (corresponding to  $?$ ) gets refined to  $[0, 10]$ , based on the knowledge that it is consistently inferior to  $[10, 10]$ . The plausibility induced by  $?$  has been made more precise by excluding definitely invalid candidate sensitivities.

For a judgment  $\varepsilon \triangleright T_1 \lesssim T_2$  to be valid, the evidence must be at least as precise as the interior, *i.e.*  $\varepsilon \sqsubseteq^2 \mathcal{I}(T_1, T_2)$ —where  $\sqsubseteq^2$  is the precision relation checked component-wise, *i.e.*  $\langle T_1, T_2 \rangle \sqsubseteq^2 \langle T'_1, T'_2 \rangle \iff T_1 \sqsubseteq T'_1 \wedge T_2 \sqsubseteq T'_2$ .

c) *Consistent transitivity*: Finally, the consistent transitivity operator for consistent sensitivity ordering is defined as:

**Definition 11** (Consistent Transitivity).

$$\begin{aligned} & \langle [s_{11}, s_{12}], [s_{13}, s_{14}] \rangle \circ \langle [s_{21}, s_{22}], [s_{23}, s_{24}] \rangle = \\ & \langle [s_{11}, s_{12} \wedge s_{14} \wedge s_{22}], [s_{13} \vee s_{21} \vee s_{23}, s_{24}] \rangle \end{aligned}$$

If  $s_{11} \leq (s_{12} \wedge s_{14} \wedge s_{22})$  and  $(s_{13} \vee s_{21} \vee s_{23}) \leq s_{24}$ . Otherwise, the consistent transitivity operation is undefined.

Same as with the interior, consistent transitivity is naturally defined for types, and is in the supplementary material.

## B. Evidence-Based Reduction

We now present GSCHECK, the internal language for GS, which it enriches with evidences. Figure 7 shows the syntax and notions of reduction (denoted by the use of  $\longrightarrow$ ) for GSCHECK. Syntactically, an ascription is augmented with evidence that justifies its intrinsic consistent subtyping judgment. Values  $v$  correspond to ascribed simple values  $u$ , and introduction forms are also ascribed (this syntactic uniformity minimizes the cases in reduction rules).

a) *Reduction*: Rule (EVR-OP) relies on the metafunction  $[\text{op}]$  and the resulting type is computed using the operator  $\delta_{\text{op}}$ . This operator is generalized to work with evidences, and the resulting evidence is computed using the  $\delta_{\text{op}}^2$  metafunction, which is defined as  $\delta_{\text{op}}^2(\langle \overline{T_{i1}}, \overline{T_{i2}} \rangle) = \langle \delta_{\text{op}}(\overline{T_{i1}}), \delta_{\text{op}}(\overline{T_{i2}}) \rangle$ .

The application rule (EVR-APP) (1) ascribes the argument to the expected type  $T'_1$ , computing a proper evidence by combining the original evidence with the domain part of the evidence of the lambda; and, (2) ascribes the body of the lambda to the expected codomain type, using the codomain part of the evidence of the lambda. Inversion metafunctions  $i\text{dom}$  and  $i\text{cod}$  are used to extract the domain and codomain parts of the evidence, respectively. If the combination of the evidences is not defined, the rule halts with an error. This treatment guarantees that all expected sensitivity types are respected during runtime.

Rule (EVR-INST) reduces the application of a reduced resource abstraction simply by substituting the distance variable  $x$  by the provided sensitivity environment  $\Sigma$ . The body of the abstraction, analogous to the application rule, is ascribed to the instantiated type, which is computed using the inst meta-

function. The body of a distance abstraction  $\Lambda x.e$  is reduced before the substitution is performed.<sup>9</sup>

Finally, rule (EVR-ASCR) merges ascriptions by keeping only the outer type  $T$  and using consistent transitivity to attempt to justify that the type of  $u$  is a consistent subtype of  $T$ . If consistent transitivity is undefined, the program halts with an error.

The reduction of GSCHECK, denoted  $e \mapsto e'$ , is call-by-value, defined with evaluation contexts  $E$  and given by the following rules:

$$\text{(EVR}E) \frac{e_1 \longrightarrow e_2}{E[e_1] \mapsto E[e_2]} \quad \text{(EVR}E\text{ERR)} \frac{e_1 \longrightarrow \mathbf{error}}{E[e_1] \mapsto \mathbf{error}}$$

An evaluation context  $E$  may be seen as a meta-expression with a hole, where the hole can be filled by an expression  $e$ , written as  $E[e]$ . The use of evaluation contexts allows us to decouple the congruence rules, *e.g.* the order of reduction of operands, from the actual notions of reduction, making the presentation more concise. Finally, reduction  $e \Downarrow v$ —as used in Figure 5—is defined as  $e \mapsto^* v$ .

b) *Evidences as Runtime Monitors*: In order to connect GSCHECK and the general treatment of monitors used in the logical relations defined in §V, we need to instantiate the abstract operators  $\text{mon}$  and  $\text{ms}$ —thereby precisely establishing that evidences are the type-driven monitoring mechanism of GSCHECK. If an expression successfully reduces to an ascribed value  $\varepsilon u :: T$ , then the evidence  $\varepsilon$  is guaranteed to convey a sound conservative approximation of the sensitivity of the computation that produced  $u$ , in particular more precise than that contained in  $T$ . Indeed, observe that consistent transitivity is monotone with respect to precision [16], tightening precision as reduction proceeds. Formally, if  $e \mapsto^* \langle \_, T' \rangle u :: T$ , then  $T' \sqsubseteq T$ . To illustrate, suppose that we have  $e \mapsto^* \langle \_, \mathbb{R} \cdot [a, b]x \rangle c :: \mathbb{R} \cdot ?x$  (we omit the left-hand side of the evidence because we want to focus on the evidence counterpart of the ascription type). By construction, evidence operations and reduction rules ensure that the interval  $[a, b]$  is a conservative approximation of the runtime sensitivity of  $e$ , *i.e.* any sensitivity within the interval is a valid sensitivity bound for  $e$ . In particular, the best *monitored sensitivity* of the expression is obtained by taking the lower bound of the right-hand side of the evidence, *i.e.*  $a$ . Given this, we instantiate the abstract operators  $\text{mon}$  and  $\text{ms}$  as follows:

$$\begin{aligned} \text{mon}(\varepsilon u :: T) &= \varepsilon & \text{ms}(\langle \_, \tau \cdot \Sigma \rangle) &= \text{lower}(\Sigma) \\ \text{where } \text{lower}([s_1, \_ ]x_1 + \dots + [s_n, \_ ]x_n) &= s_1x_1 + \dots + s_nx_n \end{aligned}$$

The function  $\text{lower}$  returns a static sensitivity environment obtained from the lower bounds of each gradual sensitivity in its argument.

## VII. SYNTACTIC TYPING AND ELABORATION

Armed with the definition of GSCHECK, we can now establish the main results of this work: a syntactic type system

<sup>9</sup>If we simply substitute  $x$  in  $e$ , we may lose information necessary to refute optimistic static assumptions, failing to produce an error at runtime. A minimal example is the expression  $(\Lambda x_2.c :: x_1 :: ?x_1 :: x_2) [x_1]$ , where  $c$  is a constant and  $x_1$  is in scope. This is reminiscent of a similar issue in gradual parametricity [34, 33].

$$\begin{array}{l}
u ::= c \mid \lambda(x : \tau).e \mid \Lambda x.v \\
v ::= \varepsilon u :: T
\end{array}
\qquad
\begin{array}{l}
e ::= v \mid \text{op } \bar{e} \mid x \mid ee \mid \varepsilon(\Lambda x.e) :: T \mid e[\Sigma] \mid \varepsilon e :: T \\
E ::= \square \mid \text{op } (\bar{v}_i, E, \bar{e}_j) \mid Ee \mid vE \mid E[\Sigma] \mid \varepsilon(\Lambda x.E) :: T \mid \varepsilon E :: T
\end{array}$$

$$\begin{array}{l}
(\text{EVR-OP}) \quad \text{op } (\overline{\varepsilon_i c_i :: T_i}) \longrightarrow \delta_{\text{op}}^2(\overline{\varepsilon_i})([\text{op}] \bar{c}_i) :: \delta_{\text{op}}(\overline{T_i}) \\
(\text{EVR-APP}) \quad (\varepsilon(\lambda(x : T'_1).e) :: T) (\varepsilon_1 u :: T_1) \longrightarrow \begin{cases} \text{icod}(\varepsilon)(e[(\varepsilon_1 \circ \text{idom}(\varepsilon))u :: T'_1/x]) :: \text{cod}(T) \\ \mathbf{error} & \text{if not defined} \end{cases} \\
(\text{EVR-INST}) \quad (\varepsilon(\Lambda x.v) :: T) [\Sigma] \longrightarrow \text{inst}(\varepsilon, \Sigma)(v[\Sigma/x]) :: \text{inst}(T, \Sigma) \\
(\text{EVR-ASCR}) \quad \varepsilon(\varepsilon' u :: T') :: T \longrightarrow \begin{cases} (\varepsilon' \circ \varepsilon)u :: T \\ \mathbf{error} & \text{if not defined} \end{cases}
\end{array}$$

$$\begin{array}{l}
\text{dom}((T_1 \rightarrow T_2) \cdot \Sigma) = T_1 \qquad \text{cod}((T_1 \rightarrow T_2) \cdot \Sigma) = T_2 \text{ (+) } \Sigma \qquad \text{inst}((\Lambda x.T_2) \cdot \Sigma, \Sigma_1) = T_2[\Sigma_1/x] \text{ (+) } \Sigma \\
\text{idom}(\langle T_1, T_2 \rangle) = \langle \text{dom}(T_2), \text{dom}(T_1) \rangle \qquad \text{icod}(\langle T_1, T_2 \rangle) = \langle \text{cod}(T_1), \text{cod}(T_2) \rangle \qquad \text{iinst}(\langle T_1, T_2 \rangle, \Sigma) = \langle \text{inst}(T_1, \Sigma), \text{inst}(T_2, \Sigma) \rangle
\end{array}$$

Fig. 7: Dynamic semantics of GSCHECK

$$\begin{array}{l}
(\text{ELCONST}) \quad \frac{c \in \llbracket B \rrbracket \quad T = B \cdot \emptyset \quad \varepsilon = \mathcal{I}(T, T)}{\Gamma \vdash c \rightsquigarrow \varepsilon c :: T : T} \qquad (\text{ELOP}) \quad \frac{\Gamma \vdash e_i \rightsquigarrow e'_i : T_i \quad \delta_{\text{op}}(\overline{T_i}) = T}{\Gamma \vdash \text{op } \bar{e}_i \rightsquigarrow \text{op } e'_i : T} \qquad (\text{ELVAR}) \quad \frac{\Gamma(x) = T}{\Gamma \vdash x \rightsquigarrow x : T} \\
(\text{ELLAM}) \quad \frac{\Gamma, x : T_1 \vdash e \rightsquigarrow e' : T_2 \quad T = (T_1 \rightarrow T_2) \cdot \emptyset \quad \varepsilon = \mathcal{I}(T, T)}{\Gamma \vdash \lambda(x : T_1).e \rightsquigarrow \varepsilon(\lambda(x : T_1).e') :: T : T} \qquad (\text{ELAPP}) \quad \frac{\Gamma \vdash e_i \rightsquigarrow e'_i : T_i \quad \varepsilon_2 = \mathcal{I}(T_2, \text{dom}(T_1))}{\Gamma \vdash e_1 e_2 \rightsquigarrow e'_1 (\varepsilon_2 e'_2 :: T_1) : \text{cod}(T_1)} \\
(\text{ELA}) \quad \frac{\Gamma \vdash e \rightsquigarrow e' : T \quad T = (\forall x.T) \cdot \emptyset \quad \varepsilon = \mathcal{I}(T, T)}{\Gamma \vdash \Lambda x.e \rightsquigarrow \varepsilon \Lambda x.e' :: T : T} \qquad (\text{ELINST}) \quad \frac{\Gamma \vdash e \rightsquigarrow e' : T}{\Gamma \vdash e[\Sigma] \rightsquigarrow e'[\Sigma] : \text{inst}(T, \Sigma)} \qquad (\text{ELASCR}) \quad \frac{\Gamma \vdash e \rightsquigarrow e' : T \quad \varepsilon = \mathcal{I}(T, T')}{\Gamma \vdash e :: T' \rightsquigarrow \varepsilon e' :: T' : T'}
\end{array}$$

Fig. 8: Syntactic Elaboration for GS

for GS that ensures that well-typed expressions satisfy gradual metric preservation. Technically, we define a syntactic type-directed elaboration from GS to GSCHECK (§VII-A). This elaboration plays two roles: it is an algorithmic syntactic typechecking procedure for GS, and it produces terms from the internal language GSCHECK that can then be executed. We prove that if a GS expression is successfully elaborated—*i.e.* that it is syntactically well-typed—then the produced GSCHECK expression is semantically well-typed. Finally, we prove that GS satisfies standard type safety as well as the gradual guarantees [29], and crucially, gradual sensitivity type soundness (§VII-B).

#### A. Type-Directed Elaboration

Figure 8 shows the type-directed elaboration from GS to GSCHECK. Rules (ELCONST), (ELLAM) and (ELA) generate evidences for ascribing introduction forms using the interior operator,  $\mathcal{I}$ . Notice that the interior is always defined in these rules, since the types are equal, hence in the consistent subtyping relation. For (ELLAM) and (ELA), the bodies of lambdas and distance abstractions are elaborated inductively.

Rule (ELOP) elaborates the arguments of an operator application and computes the type using the operator  $\delta_{\text{op}}$  (similar to

the typing rule presented in Figure 2). Rule (ELVAR) simply retrieves the type of the variable from the typing environment.

Rule (ELAPP) inductively elaborates both inner expressions, but also ascribes the argument of the application to the exact type expected by the callee. Consequently, an evidence for such ascription is produced using the interior operator. Note however that the interior may not be defined for the involved types, thereby revealing a syntactic type error.

Rule (ELINST) simply elaborates the callee and computes the type using the  $\text{inst}$  function. Finally, rule (ELASCR) is straightforward, as it simply elaborates the inner expression and ascribes it to the specified type. If the interior is not defined, a syntactic type error is exhibited.

Crucially, if elaboration succeeds, *i.e.* the source expression is syntactically well-typed, then the elaborated expression is semantically well-typed—possibly termination-sensitive if its gradual sensitivity is bounded.

**Lemma 6** (Syntactic Elaboration implies Semantic Typing). *If  $\Gamma \vdash e \rightsquigarrow e' : \tau \cdot \Sigma$  then:*

- 1)  $\Gamma \models e' : \tau \cdot \Sigma$ , and
- 2) if bounded( $\Sigma$ ) then  $\Gamma \models_{\text{ts}} e' : \tau \cdot \Sigma$

Finally, by means of the elaboration, we can define a syntactic typechecking judgment and reduction for GS expressions:

$$\begin{aligned} \Gamma \vdash e : T &\stackrel{\text{def}}{=} \Gamma \vdash e \rightsquigarrow e' : T \text{ for some } e' \\ e \Downarrow r &\stackrel{\text{def}}{=} \vdash e \rightsquigarrow e' : T \wedge e' \mapsto^* r \end{aligned}$$

In the evaluation judgment  $e \Downarrow r$ ,  $e$  is a GS expression and the result  $r$  is either a GSCHECK value  $v$  or **error**.

### B. Metatheory of GS

We now establish that GS is type safe, satisfies the gradual guarantees, and is type sound in that syntactically well-typed terms satisfy gradual metric preservation. Note that some of these results follow from the results for the internal evidence-based language combined with the fact that syntactic typing for GS is directly defined as the successful elaboration to GSCHECK.

First, GS is type safe: closed expressions do not get stuck.

**Proposition 7** (Type Safety). *Let  $e$  be a closed GS expression. If  $\vdash e : T$ , then  $e \Downarrow r$ .*

GS satisfies the static and dynamic gradual guarantees [29]. The static guarantee says that typeability is monotone with respect to imprecision.<sup>10</sup>

**Proposition 8** (Static Gradual Guarantee). *Let  $e_1$  and  $e_2$  be two closed GS expressions such that  $e_1 \sqsubseteq e_2$  and  $\vdash e_1 : T_1$ . Then,  $\vdash e_2 : T_2$  and  $T_1 \sqsubseteq T_2$ .*

GS also satisfies the dynamic gradual guarantee: any program that reduces without error will continue to do so if imprecision is increased.

**Proposition 9** (Dynamic Gradual Guarantee). *Let  $e_1$  and  $e_2$  be two well-typed closed GS expressions such that  $e_1 \sqsubseteq e_2$ . If  $e_1 \Downarrow v_1$  then  $e_2 \Downarrow v_2$ , where  $v_1 \sqsubseteq v_2$ .*

Any syntactically well-typed GS term satisfies gradual metric preservation. More specifically, a well-typed term satisfies termination-insensitive gradual metric preservation, and if its type has bounded imprecision, then it also satisfies termination-sensitive gradual metric preservation.

**Theorem 10** (Sensitivity Type Soundness). *Let  $\Gamma = x_1 : B_1 \cdot \Sigma_1, \dots, x_n : B_n \cdot \Sigma_n$ . If  $\Gamma \vdash e \rightsquigarrow e' : B \cdot \Sigma$  then:*

- 1)  $\text{GMP}(\Gamma, e', B, \Sigma)$ , and
- 2) if bounded( $\Sigma$ ) then  $\text{GMP}_{\text{ts}}(\Gamma, e', B, \Sigma)$

## VIII. DIFFERENTIAL PRIVACY REASONING

Differential privacy [4] is a strong formal property that ensures that the output of a computation (or query) does not reveal information about any single element in the input data. Formally, a program  $M$  is  $\epsilon$ -differentially private if for all neighboring inputs  $x$  and  $x'$ , and for all possible outputs  $r$  we have  $\Pr[M(x) = r] \leq e^\epsilon \Pr[M(x') = r]$ . In general, differential privacy algorithms assume pure (total) functions,

<sup>10</sup>Precision between expressions is the standard, natural syntactical lifting of type precision to terms (provided in the supplementary material).

and hence the possibility of leaks happening through abnormal termination behavior like errors is usually not considered.

Because sensitivity analysis is often used to build differentially-private programs, it is important to address the question of differential privacy built on top of a gradual sensitivity language like GSoul. In this section, we first illustrate why *termination-insensitive* gradual metric preservation falls short of achieving privacy. We then demonstrate that *termination-sensitive* gradual metric preservation can be used in differential privacy proofs to ensure not only that sensitivity specifications are respected, but also that a program does not leak information through termination behavior. Specifically, we present some gradual variants of core differentially-private mechanisms and algorithms in GSoul that leverage termination-sensitive metric preservation, which can serve as fundamental components in larger differentially-private programs.

### A. Termination Insensitivity and Differential Privacy

We start by considering expressions that only satisfy termination-insensitive metric preservation (Definition 7). In such a scenario, the differential privacy definition can be violated: it is possible to find a program  $M$ , such that given two different inputs  $x$  and  $x'$  at distance 1, then  $\forall r. \Pr[M(x) = r] \not\leq e^\epsilon \Pr[M(x') = r]$ . This occurs because  $M(x')$  could result in an error (or diverge), reducing the probability  $\Pr[M(x') = r]$  to 0. To illustrate this, consider the following typing judgment:

$$\begin{aligned} x : \mathbb{R} \cdot x, f : (\mathbb{R} \cdot x \rightarrow \mathbb{B} \cdot ?x) \cdot \emptyset \\ \vdash \text{if } (f x) \text{ then } 1 \text{ else } (x :: \mathbb{R} \cdot ?x :: \mathbb{R} \cdot 0x) : \mathbb{R} \cdot ?x \end{aligned}$$

The resulting type for this expression is computed by joining the types of the two branches of the conditional,  $\mathbb{R} \cdot 0x$  and  $\mathbb{R} \cdot ?x$ , with the sensitivity environment of the condition,  $?x$ . Note that this well-typed term fails at runtime for any  $x$  and  $f$  such that  $f x$  evaluates to false. Then, in order for two executions to exhibit different termination behavior, it suffices to pick environments  $\Delta = 1x$ , and:

$$\gamma_1 = \{f \mapsto \lambda x. x > 0, x \mapsto 0\} \quad \gamma_2 = \{f \mapsto \lambda x. x > 0, x \mapsto 1\}$$

Importantly, this expression is not semantically well-typed in a termination-sensitive regime, as  $?x$  is not bounded. If the codomain of  $f$  were bounded,  $x > 0$  could not be admitted as the body of  $f$ , and  $\gamma_1$  and  $\gamma_2$  could not be used to distinguish executions. This observation highlights the importance of strengthening gradual programs so that imprecision is bounded in order to achieve differential privacy.

### B. Gradual Laplace Mechanism

The Laplace mechanism is a fundamental building block in differential privacy, which adds noise drawn from the Laplace distribution to the output of a query. In particular, if a query  $Q$  is 1-sensitive, then  $Q(D) + \mathcal{L}(1/\epsilon)$  is  $\epsilon$ -differentially private, for any database  $D$  [4]. With gradual sensitivities, we define the *Gradual Laplace Mechanism* (GLM) which expects an unknown-sensitive function and performs a runtime check to ensure that the function is 1-sensitive. Figure 9 presents the

```

1 def GLM<T>(
2   res x: T,
3   f: T[1x] -> Number[?x],
4   eps: Number,
5 ) = laplace(f(x) :: Number[1x], 1, eps);

```

Fig. 9: Gradual Laplace Mechanism (GLM)

```

1 def GAT<T>(
2   res db: T,
3   fs: List<T[1db] -> Number[?db]>,
4   thr: Number,
5   eps: Number,
6 ): Number = {
7   let noisyThr = laplace(thr, 1, eps / 2);
8   fs.indexOf(
9     fn (f: T[1db] -> Number[?db]) =>
10    try {
11      let noisyVal = GLM<T>(db, f, eps / 4);
12      noisyVal >= noisyThr;
13    } catch { false; }
14  );
15 };

```

Fig. 10: Gradual Above Threshold (GAT)

definition of GLM in GSOUL, parametric on the type ( $T$ ) of the input  $x$ .

Given a function  $f$ , the expression  $f(x) :: \text{Number}[1x]$  has a bounded sensitivity so it is termination-sensitive semantically well-typed. Therefore, for any two values for  $x$  (which are necessarily at a bounded distance), GLM either (1) always fails for any input  $x$ , or (2) always produces noisy values and is  $\text{eps}$ -differentially private.

### C. Gradual Above Threshold

Recall the Above Threshold (AT) algorithm discussed in §II. This algorithm uses the Laplace mechanism as a subroutine to add noise to query results, which are then compared with a (noisy) threshold [4]. We follow the same approach to define a gradual version of the AT algorithm, which instead uses the gradual Laplace mechanism GLM defined previously. The new *Gradual Above Threshold* (GAT) algorithm expects a list of queries of unknown sensitivities, but instead of failing if the sensitivity of a given function is greater than 1, it simply skips the function. Figure 10 presents the GAT algorithm, written in GSOUL. First, the algorithm computes a noisy version of the threshold (line 7). Then, we use the `indexOf` function to find the index of the first query whose (noisy) result is above the noisy threshold. Notice that, given the use of GLM, executing line 11 may fail, but the `try/catch` surrounding it ensures that in such case the algorithm continues with the next query. In contrast to AT, GAT can receive a list of gradually well-typed functions of varied sensitivities, and the algorithm will only consider the 1-sensitive ones. Hence, the algorithm returns the index of the *first* 1-sensitive query whose result is above  $\text{thr}$ .

The proof that GAT is differentially private follows the same methodology used by Dwork and Roth [4]. The key difference is that, in the case of GAT, the proof first establishes

a termination-sensitive result for the GLM call. Thus, for a particular index, line 11 will always behave the same, either succeeding or failing. This guarantees that the same functions are skipped every time, independently of the database. Finally, given that the discarded functions do not reveal information about the database, the same proof structure of AT can be followed (see supplementary material).

Note that with termination-sensitive metric preservation, runtime consistent transitivity errors only reveal information about the sensitivity of functions, *not about their arguments*, which is sufficient to preserve differential privacy. For instance, for GLM, an error will only reveal that the function passed as argument is not 1-sensitive, independent of the input value. Proving that this holds in general, along with integrating this kind of reasoning in a differential privacy programming language such as FUZZ [5] or JAZZ [10] is left to future work.

## IX. RELATED WORK

a) *Sensitivity and programming languages*: The first type system for reasoning about sensitivity is FUZZ [5], a language for differential privacy using linear types. Several variations have been studied, such as DFUZZ [6], FUZZI [7], and *Adaptive Fuzz* [8].  $\mu$ FUZZ [11] extends the FUZZ compiler to generate nonlinear constraints that are discharged by an SMT solver, resulting in an automatic type-based sensitivity analysis. All of these type systems measure sensitivity and also track and enforce differential privacy. Near et al. [9] tackle differential privacy in DUET with two mutually-defined languages, one dedicated to sensitivity and one to privacy. JAZZ [10] follows the approach of DUET, and includes SAX, a sensitivity language with contextual linear types and delayed sensitivity effects. The starting point of GSOUL, the static language SOUL, is very close to SAX. Whereas FUZZ-like languages track the sensitivity of program variables using linear types, Abua et al. [12] propose SOLO, which tracks a fixed amount of sensitive resources and avoids linear types. Lobo-Vesga et al. [35] develop a haskell library, DPELLA, able to reason about the accuracy of differentially private queries. DDUO [13] is a library for dynamic sensitivity tracking in Python, which tackles the expressiveness issues of static sensitivity type systems. However, as a fully dynamic approach, DDUO does not provide any static guarantee; in contrast, GSOUL allows programmers to balance static and dynamic checking as needed. To the best of our knowledge, this integration of static and dynamic sensitivity analysis within one language is unique to GSOUL.

Regarding termination sensitivity, FUZZ [5] and more recent languages [6, 12] focus on the terminating fragment of their languages. Reed and Pierce [5] discuss the tension between metric preservation and non-termination in their paper, presenting three alternatives: weakening the definition of metric preservation; proving statically that recursive functions terminate, yielding more complex programs; or adding *fuel* to recursive functions, falling back to a default value when exhausted. The latter is adopted in the implementation.  $\mu$ FUZZ [11] also supports recursive types, and termination-sensitive

metric preservation is obtained by reasoning only about finite distances. DUET [9] avoids divergence via terminating looping primitives. DDUO [13] establishes termination-insensitive metric preservation, although they do not explicitly discuss this aspect.

Another line of work for sensitivity verification is based on program logics [36, 37, 38, 39]. These approaches are generally very expressive but less automatic than type systems. Extending gradual verification approaches [40, 41], to account for the aforementioned logics for sensitivity would be an interesting venue for future work.

b) *Gradual typing*: To the best of our knowledge, gradual typing has not been applied to sensitivity typing. It has, however, been applied in languages with one of two particularly interesting properties: languages with type-and-effect disciplines; and languages whose soundness property corresponds to a hyperproperty, such as noninterference [27] or parametricity [32]. Interestingly, gradual security typing has only been explored for termination- (and error-)insensitive characterizations of noninterference [23, 22, 42]. Bañados Schwerter et al. [30, 43] develop a general approach to gradualize type-and-effects, extended by Toro and Tanter [44] to handle effect polymorphism in Scala. Being based on the generic type-and-effect system of Marino and Millstein [45], these approaches cannot handle sensitivities as quantities within a range.

Toro et al. [23] uncovered that the presence of mutable references in a security language yields a gradual language that does not satisfy noninterference; ad-hoc changes to address implicit flows recover noninterference, at the expense of the dynamic gradual guarantee. An interesting perspective is to study an extension of GSOUL with mutable references, investigating if metric preservation and the dynamic gradual guarantee are both satisfied. They also use intervals (of security labels) in their runtime semantics in order to achieve soundness; GSOUL chooses to expose intervals in the source language for convenience. Bañados Schwerter et al. [46] study *forward completeness* of evidence representation as a criterion to ensure that modular type-based semantic invariants are preserved in a gradual language. Proving forward completeness for GSCHECK is an interesting exercise for future work.

## X. CONCLUSION

We have proposed gradual sensitivity typing as a means to seamlessly combine static and dynamic checking of sensitivity specifications, accommodating features and programming patterns that are challenging for fully static approaches. We designed and implemented GSOUL, a functional programming language that can be used to experiment with this approach. We formalized a core of GSOUL and studied its metatheory.

In particular, we have developed a semantic account of gradual metric preservation, taking into consideration its termination aspect, inevitable in a gradual language with runtime errors, and illustrated how to reason about differential privacy with gradual sensitivities.

This work represents an important first step towards gradual differentially-private programming. Future work will focus

on building a privacy layer on top of GSOUL, following the two-language approach of DUET [9] and JAZZ [10]. It is also crucial to further empirically evaluate the benefits and tradeoffs of gradual sensitivity typing, which requires further growing GSOUL to integrate practical libraries, or by implementing gradual sensitivities in a widely-used programming language.

## REFERENCES

- [1] G. Zames, “Input-output feedback stability and robustness, 1959-85,” *IEEE Control Systems Magazine*, vol. 16, no. 3, pp. 61–66, 1996.
- [2] O. Bournez, D. Graça, and E. Hainry, “Robust computations with dynamical systems,” in *MFCS 2010*, ser. Lecture Notes in Computer Science, P. Hlineny and A. Kucera, Eds., vol. 6281. Springer-Verlag, Aug. 2010, pp. 198–208.
- [3] S. Chaudhuri, S. Gulwani, R. Lubliner, and S. Navidpour, “Proving programs robust,” in *Proc. of ESEC/FSE 2011*. ACM, 2011, p. 102–112.
- [4] C. Dwork and A. Roth, “The algorithmic foundations of differential privacy,” *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3–4, p. 211–407, Aug. 2014.
- [5] J. Reed and B. C. Pierce, “Distance makes the types grow stronger: A calculus for differential privacy,” in *Proc. of ICFP 2010*. ACM, Sep. 2010, p. 157–168.
- [6] M. Gaboardi, A. Haeberlen, J. Hsu, A. Narayan, and B. C. Pierce, “Linear dependent types for differential privacy,” in *Proc. of POPL 2013*. ACM, 2013, p. 357–370.
- [7] H. Zhang, E. Roth, A. Haeberlen, B. C. Pierce, and A. Roth, “Fuzzi: A three-level logic for differential privacy,” *PACMPL*, vol. 3, no. ICFP, Jul. 2019.
- [8] D. Winograd-Cort, A. Haeberlen, A. Roth, and B. C. Pierce, “A framework for adaptive differential privacy,” *PACMPL*, vol. 1, no. ICFP, Aug. 2017.
- [9] J. P. Near, D. Darais, C. Abua, T. Stevens, P. Gadamadugu, L. Wang, N. Somani, M. Zhang, N. Sharma, A. Shan, and D. Song, “Duet: An expressive higher-order language and linear type system for statically enforcing differential privacy,” *PACMPL*, vol. 3, no. OOPSLA, Oct. 2019.
- [10] M. Toro, D. Darais, C. Abua, J. Near, D. Áquez, F. Olmedo, and É. Tanter, “Contextual linear types for differential privacy,” *ACM TOPLAS*, vol. 45, no. 2, May 2023.
- [11] L. D’Antoni, M. Gaboardi, E. J. Gallego Arias, A. Haeberlen, and B. Pierce, “Sensitivity analysis using type-based constraints,” in *Proc. of FPDSL 2013*. ACM, 2013, p. 43–50.
- [12] C. Abua, D. Darais, and J. P. Near, “Solo: a lightweight static analysis for differential privacy,” *PACMPL*, vol. 6, no. OOPSLA2, oct 2022.

- [13] C. Abuaah, A. Silence, D. Darais, and J. P. Near, “DDUO: general-purpose dynamic analysis for differential privacy,” in *IEEE CSF 2021*. IEEE, 2021, pp. 1–15.
- [14] D. K. Gifford and J. M. Lucassen, “Integrating functional and imperative programming,” in *Proc. of LFP 1986*, 1986, pp. 28–38.
- [15] J. Siek and W. Taha, “Gradual typing for functional languages,” in *Proc. of the Scheme and Functional Programming Workshop*, Sep. 2006, pp. 81–92.
- [16] R. Garcia, A. M. Clark, and É. Tanter, “Abstracting gradual typing,” in *Proc. of POPL 2016*, R. Bodík and R. Majumdar, Eds. ACM Press, Jan. 2016, pp. 429–442, see erratum: <https://www.cs.ubc.ca/~rxg/agt-erratum.pdf>.
- [17] J. Siek and W. Taha, “Gradual typing for objects,” in *Proc. of ECOOP 2007*, ser. Lect. Notes in Comp. Sci., E. Ernst, Ed., no. 4609. Springer-Verlag, Jul. 2007, pp. 2–27.
- [18] D. Herman, A. Tomb, and C. Flanagan, “Space-efficient gradual typing,” *Higher-Order and Symbolic Computation*, vol. 23, no. 2, pp. 167–189, Jun. 2010.
- [19] J. G. Siek, M. M. Vitousek, M. Cimini, S. Tobin-Hochstadt, and R. Garcia, “Monotonic references for efficient gradual typing,” in *Proc. of ESOP 2015*, ser. Lect. Notes in Comp. Sci., J. Vitek, Ed., vol. 9032. Springer-Verlag, Mar. 2015, pp. 432–456.
- [20] I. Sergey and D. Clarke, “Gradual ownership types,” in *Proc. of ESOP 2012*, ser. Lect. Notes in Comp. Sci., H. Seidl, Ed., vol. 7211. Springer-Verlag, 2012, pp. 579–599.
- [21] T. Disney and C. Flanagan, “Gradual information flow typing,” in *International Workshop on Scripts to Programs*, 2011.
- [22] L. Fennell and P. Thiemann, “Gradual security typing with references,” in *IEEE CSF 2013*, Jun. 2013, pp. 224–239.
- [23] M. Toro, R. Garcia, and É. Tanter, “Type-driven gradual security with references,” *ACM TOPLAS*, vol. 40, no. 4, pp. 16:1–16:55, Nov. 2018.
- [24] N. Lehmann and É. Tanter, “Gradual refinement types,” in *Proc. of POPL 2017*. ACM Press, Jan. 2017, pp. 775–788.
- [25] N. Heintze and J. G. Riecke, “The SLam calculus: Programming with secrecy and integrity,” in *Proc. of POPL 1998*. ACM Press, Jan. 1998, pp. 365–377.
- [26] S. Zdancewic, “Programming languages for information security,” Ph.D. dissertation, Cornell Univ., Aug. 2002.
- [27] J. A. Goguen and J. Meseguer, “Security policies and security models,” in *IEEE SP 1982*, 1982, pp. 11–11.
- [28] A. Azevedo de Amorim, M. Gaboardi, J. Hsu, S.-y. Katsumata, and I. Cherigui, “A semantic account of metric preservation,” in *Proc. of POPL 2017*. ACM, 2017, p. 545–556.
- [29] J. G. Siek, M. M. Vitousek, M. Cimini, and J. T. Boyland, “Refined criteria for gradual typing,” in *SNAPL 2015*, ser. Leibniz International Proceedings in Informatics, vol. 32. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, May 2015, pp. 274–293.
- [30] F. Bañados Schwerter, R. Garcia, and É. Tanter, “A theory of gradual effect systems,” in *Proc. of ICFP 2014*. ACM Press, Sep. 2014, pp. 283–295.
- [31] B. C. Pierce, *Types and programming languages*. Cambridge, MA, USA: MIT Press, 2002.
- [32] J. C. Reynolds, “Types, abstraction, and parametric polymorphism,” in *Information Processing 83*, R. E. A. Mason, Ed. Elsevier, 1983, pp. 513–523.
- [33] M. Toro, E. Labrada, and É. Tanter, “Gradual parametricity, revisited,” *PACMPL*, no. POPL, pp. 17:1–17:30, 2019.
- [34] A. Ahmed, D. Jamner, J. G. Siek, and P. Wadler, “Theorems for free for free: Parametricity, with and without types,” *PACMPL*, vol. 1, no. ICFP, pp. 39:1–39:28, Sep. 2017.
- [35] E. Lobo-Vesga, A. Russo, and M. Gaboardi, “A programming language for data privacy with accuracy estimations,” *ACM TOPLAS*, vol. 43, no. 2, Jun. 2021.
- [36] G. Barthe, B. Köpf, F. Olmedo, and S. Zanella Béguelin, “Probabilistic relational reasoning for differential privacy,” in *Proc. of POPL 2012*. ACM, 2012, p. 97–110.
- [37] G. Barthe, B. Köpf, F. Olmedo, and S. Zanella-Béguelin, “Probabilistic relational reasoning for differential privacy,” *ACM TPLS* 2013, vol. 35, no. 3, Nov. 2013.
- [38] G. Barthe, M. Gaboardi, B. Grégoire, J. Hsu, and P.-Y. Strub, “Proving differential privacy via probabilistic couplings,” in *Proc. of LICS 2016*. ACM, 2016, p. 749–758.
- [39] T. Sato, G. Barthe, M. Gaboardi, J. Hsu, and S.-y. Katsumata, “Approximate span liftings: Compositional semantics for relaxations of differential privacy,” in *Proc. of LICS 2019*, pp. 1–14.
- [40] J. Bader, J. Aldrich, and É. Tanter, “Gradual program verification,” in *Prof. of VMCAI 2018*, I. Dillig and J. Palsberg, Eds., vol. 10747. Springer-Verlag, Jan. 2018, pp. 25–46.
- [41] J. Wise, J. Bader, C. Wong, J. Aldrich, É. Tanter, and J. Sunshine, “Gradual verification of recursive heap data structures,” *Proc. of OOPSLA 2020*, vol. 4, pp. 228:1–228:28, Nov. 2020.
- [42] A. A. de Amorim, M. Fredrikson, and L. Jia, “Reconciling noninterference and gradual typing,” in *Proc. of LICS 2020*. ACM, 2020, p. 116–129.
- [43] F. Bañados Schwerter, R. Garcia, and É. Tanter, “Gradual type-and-effect systems,” *Journal of Functional Programming*, vol. 26, pp. 19:1–19:69, Sep. 2016.
- [44] M. Toro and É. Tanter, “Customizable gradual polymorphic effects for Scala,” in *Proc. of OOPSLA 2015*. ACM Press, Oct. 2015, pp. 935–953.
- [45] D. Marino and T. Millstein, “A generic type-and-effect system,” in *Proc. of TLDI 2009*, pp. 39–50.
- [46] F. Bañados Schwerter, A. M. Clark, K. A. Jafery, and R. Garcia, “Abstracting gradual typing moving forward: precise and space-efficient,” *PACMPL*, vol. 5, no. POPL, Jan. 2021.