



HAL
open science

Tropical Mathematics and the Lambda-Calculus I: Metric and Differential Analysis of Effectful Programs

Davide Barbarossa, Paolo Pistone

► **To cite this version:**

Davide Barbarossa, Paolo Pistone. Tropical Mathematics and the Lambda-Calculus I: Metric and Differential Analysis of Effectful Programs. Computer Science Logic 2024 (CSL 2024), Feb 2024, Napoli, France. 10.4230/LIPIcs.CSL.2024.14 . hal-04777395

HAL Id: hal-04777395

<https://hal.science/hal-04777395v1>

Submitted on 10 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.




L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Tropical Mathematics and the Lambda-Calculus I

Metric and Differential Analysis of Effectful Programs

Daive Barbarossa   

Università di Bologna, Italy

Paolo Pistone   

Università di Bologna, Italy

Abstract

We study the interpretation of the lambda-calculus in a framework based on tropical mathematics, and we show that it provides a unifying framework for two well-developed quantitative approaches to program semantics: on the one hand program metrics, based on the analysis of program sensitivity via Lipschitz conditions, on the other hand resource analysis, based on linear logic and higher-order program differentiation. To do that, we focus on the semantics arising from the relational model weighted over the tropical semiring, and we discuss its application to the study of “best case” program behavior for languages with probabilistic and non-deterministic effects. Finally, we show that a general foundation for this approach is provided by an abstract correspondence between tropical algebra and Lawvere’s theory of generalized metric spaces.

2012 ACM Subject Classification Theory of computation → Lambda calculus; Theory of computation → Categorical semantics; Theory of computation → Linear logic

Keywords and phrases Relational semantics, Differential lambda-calculus, Tropical semiring, Program metrics, Lawvere quantale

Digital Object Identifier 10.4230/LIPIcs.CSL.2024.14

Related Version *Extended Version*: <https://arxiv.org/abs/2311.15704>

Funding This work has been supported by the European Research Council through the project ERC CoG 818616 DIAPASoN.

1 Introduction

In recent years, more and more interest in the programming language community has been directed towards the study of *quantitative* properties of programs like, e.g., the number of computation steps or the probability of convergence, as opposed to purely *qualitative* properties like termination or program equivalence. Notably, a significant effort has been made to extend, or adapt, well-established qualitative methods, like type systems, relational logics or denotational semantics, to account for quantitative properties. We can mention, for example, intersection type systems aimed at capturing time or space resources [3, 28] or convergence probabilities [9, 19], relational logics to account for probabilistic properties like, e.g., differential privacy [13] or metric preservation [24, 68], as well as the study of denotational models for probabilistic [33, 46] or differential [36] extensions of the λ -calculus. The main reason to look for methods relying on (quantitative extensions of) type-theory or denotational semantics is that these approaches yield *modular* and *compositional* techniques, that is, allow one to deduce properties of complex programs from the properties of their constituent parts.

Two approaches to quantitative semantics. Among such quantitative approaches, two have received considerable attention.



© Davide Barbarossa and Paolo Pistone;

licensed under Creative Commons License CC-BY 4.0

32nd EACSL Annual Conference on Computer Science Logic (CSL 2024).

Editors: Aniello Murano and Alexandra Silva; Article No. 14; pp. 14:1–14:23

Leibniz International Proceedings in Informatics



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

On the one hand, there is the approach of *program metrics* [10, 11, 68] and *quantitative equational theories* [58]: when considering probabilistic or approximate computation, rather than asking whether two programs compute *the same* function, it makes more sense to ask whether they compute functions which do not differ *too much*. This has motivated the study of denotational frameworks in which types are endowed with a metric, measuring similarity of behavior; this approach has found applications in, e.g., differential privacy [68] and coinductive methods [12], and was recently extended to account for the full λ -calculus [25, 42, 66].

On the other hand, there is the approach based on *differential* [36] or *resource-aware* [18] extensions of the λ -calculus, which is well-connected to the so-called *relational semantics* [30, 50, 56] and has a syntactic counterpart in the study of *non-idempotent* intersection types [28, 59]. This family of approaches have been exploited to account for higher-order program differentiation [36], to establish reasonable *cost-models* for the λ -calculus [2], and have also been shown suitable to the probabilistic setting [9, 19, 50].

In both approaches the notion of *linearity*, in the sense of linear logic [43] (i.e. of using inputs exactly once), plays a crucial role. In metric semantics, linear programs correspond to *non-expansive* maps, i.e. to functions that do not increase distances, and the possibility of duplicating inputs leads to interpret programs with a fixed duplication bound as *Lipschitz-continuous* maps [10]. By contrast, in the standard semantics of the differential λ -calculus, linear programs correspond to linear maps, in the usual algebraic sense, while the possibility of duplicating inputs leads to consider functions defined as *power series*.

A natural question, at this point, is whether these two apparently unrelated ways of interpreting linearity and duplication can be somehow reconciled. At a first glance, there seems to be a “logarithmic” gap between the two approaches: in metric models a program duplicating an input n times yields a linear (hence *Lipschitz*) function nx , whereas in differential models it would lead to a polynomial function x^n , thus not Lipschitz. The fundamental idea behind this work is the observation that this gap is naturally overcome once we interpret these functions in the framework of tropical mathematics, where, as we will see, the monomial x^n precisely reads as the linear function nx .

Tropical mathematics and program semantics. Tropical mathematics was introduced in the seventies by the Brazilian mathematician Imre Simon [72] as an alternative approach to algebra and geometry where the usual ring structure of numbers based on addition and multiplication is replaced by the semiring structure given, respectively, by “min” and “+”. For instance, the polynomial $p(x, y) = x^2 + xy^2 + y^3$, when interpreted over the tropical semiring, translates as the piecewise linear function $\varphi(x, y) = \min\{2x, x + 2y, 3y\}$. In the last decades, tropical geometry evolved into a vast and rich research domain, providing a combinatorial counterpart of usual algebraic geometry, with important connections with optimisation theory [55]. Computationally speaking, working with min and + is generally easier than working with standard addition and multiplication; for instance, the fundamental (and generally intractable) problem of finding the roots of a polynomial admits a *linear time* algorithm in the tropical case (and, moreover, the tropical roots can be used to approximate the actual roots [62]). The combinatorial nature of several methods in tropical mathematics explains why these are so widely applied in computer science, notably for convex analysis and machine learning (see [57] for a recent survey).

Coming back to our discussion on program semantics, tropical mathematics seems to be just what we look for, as it turns polynomial functions like x^n into Lipschitz maps like nx . At this point, it is worth mentioning that a tropical variant of the usual relational semantics of linear logic and the λ -calculus has already been considered [50], and shown

capable of capturing *best-case* quantitative properties, but has not yet been studied in detail. Furthermore, connections between tropical linear algebra and metric spaces have also been observed [38] within the abstract setting of *quantale-enriched* categories [47, 74]. However, a thorough investigation of the interpretation of the λ -calculus within tropical mathematics and of the potentialities of its applications has not yet been undertaken.

In this paper we make a first step in such direction, by demonstrating that the relational interpretation of the λ -calculus based on tropical mathematics does indeed provide the desired bridge between differential and metric semantics, and suggests new combinatorial methods to study probabilistic and non-deterministic programs.

Contributions and outline of the paper. Our contributions in this paper are the following:

- We first show that tropical polynomials naturally arise in the best-case analysis of probabilistic and non-deterministic programs, turning the study of quantitative program behavior into a purely combinatorial problem. This is in Sections 3 and 4.
- We study the relational model over the tropical semiring, which provides a semantics of effectful extensions of the simply typed λ -calculus (STLC in the following) and PCF [67]. Notably, we show that higher-order programs are interpreted by a generalizations of tropical power series [61], and we show that these functions are locally Lipschitz-continuous, thus yielding a full-scale metric semantics. This is in Sections 5 and 6.
- We exploit the differential structure of the relational model to study the *tropical Taylor expansion* of a λ -term, which can be seen as an approximation of the term by way of Lipschitz-continuous maps, and we show that it can be used to compute approximated Lipschitz-constants for higher-order programs. This is in Section 7.
- We conclude by framing the connection between the tropical, differential and metric viewpoints at a more abstract level. We recall a well-known correspondence between Lawvere’s *generalized metric spaces* [51, 74] and modules over the tropical semi-ring [69] and we show that it yields a model of the differential λ -calculus which extends the tropical relational model. This is in Section 8.

2 A Bridge between Metric and Differential Aspects

In this section, we discuss in some more detail the two approaches to quantitative semantics we mentioned in the Introduction, at the same time providing an overview of how we aim at bridging them using tropical mathematics.

Metric Approach: Bounded λ -Calculus. In many situations (e.g. when dealing with computationally difficult problems) one does not look for algorithms to compute a function *exactly*, but rather to approximate it (in an efficient way) within some error bound. In other common situations (e.g. in differential privacy [8, 68]) one needs to verify that an algorithm is not *too sensitive* to errors, that is, that a small error in the input will produce a comparably small error in the output. In all these cases, it is common to consider forms of denotational semantics in which types are endowed with a *behavioral metric*, that is, a metric on programs which accounts for differences in behavior. A fundamental insight coming from this line of work is that, if one can somehow *bound* the number of times that a program may duplicate its input, the resulting program will be *Lipschitz-continuous*: if M may duplicate at most L times, then an error ϵ between two inputs will result in an error less or equal to $L \cdot \epsilon$ in the corresponding outputs [10, 68] (yet, this property may fail in a concurrent setting, see e.g. [41]). For instance, the higher-order program $M = \lambda f. \lambda x. f(f(x))$, which

duplicates the functional input f , yields a 2-Lipschitz map between the metric space $\mathbb{R} \multimap \mathbb{R}$ of non-expansive real functions and itself: if f, g are two non-expansive maps differing by at most ϵ (i.e. for which $|f(x) - g(x)| \leq \epsilon$ holds for all $x \in \mathbb{R}$), then the application of M to f and g will produce two maps differing by at most 2ϵ .

These observations have led to the study of λ -calculi with *graded* exponentials, like Fuzz [68], inspired from Girard’s Bounded Linear Logic [44], which have been applied to the study of differential privacy [10, 39]. The types of such systems are defined by combining linear constructors with a *graded linear exponential comonad* $!_r(-)$ [49].


Yet, what about the good old, “unbounded”, simply typed λ -calculus? Actually, by using unbounded duplications, one might lose the Lipschitz property. For instance, while the functions $M_k = \lambda x.k \cdot x : \mathbb{R} \rightarrow \mathbb{R}$ are all Lipschitz-continuous, with Lipschitz constant k , the function $M = \lambda x.x^2$ obtained by “duplicating” x is not Lipschitz anymore: M is, so to say, *too* sensitive to errors. More abstractly, it is well-known that the category \mathbf{Met} of metric spaces and non-expansive maps, is *not* cartesian closed, so it is not a model of STLC (yet, several cartesian closed *sub*-categories of \mathbf{Met} do exist, see e.g. [21, 25]). Still, one might observe that the program M above is actually Lipschitz-continuous, if not globally, at least *locally* (i.e. over any compact set). Indeed, some cartesian closed categories of locally Lipschitz maps have been produced in the literature [33, 66], and a new example will be exhibited in this paper.

Resource Approach: the Differential λ -Calculus. A different family of approaches to linearity and duplication arises from the study of the *differential λ -calculus* [36] (and differential linear logic [30]) and its categorical models. The key ingredient is a *differential constructor* $D[_, _]$, added to the usual syntax of the λ -calculus. The intuition is that, given M of type $A \rightarrow B$ and N of type A , the program $D[M, N]$, still of type $A \rightarrow B$, corresponds to the *linear application* of M to N : this means that N is passed to M so that the latter may use it exactly once. This is also why $D[M, N]$ still has type $A \rightarrow B$, since M might need *other* copies of an input of type A . In particular, the application of $D[M, N]$ to an “error term” 0 ensures that M will use N exactly once (we say *linearly*).

The reason why D is called a “differential”, is twofold: semantically, its interpretation is a generalisation of the usual differential form analysis (see Section 5); syntactically, it allows to define the so-called *Taylor expansion* \mathcal{T} of programs: the idea is that one can expand any application MN as an infinite formal sum of *linear* applications $D^k[M, N^k]0$, i.e. where N is linearly passed exactly k times to M ; doing this recursively gives rise to the suggestive Taylor formula $\mathcal{T}(MN) := \sum_{k=0}^{\infty} \frac{1}{k!} \cdot D^k[\mathcal{T}(M), \mathcal{T}(N)^k]0$. In other words, unbounded duplications correspond to some sort of limit of bounded, but arbitrarily large, ones.

Tropical Mathematics: Lipschitz Meets Taylor. At this point, as the Taylor formula decomposes an unbounded application as a limit of bounded ones, one might well ask whether it could be possible to see this formula as interpreting a λ -term as a limit of Lipschitz maps, in some sense, thus bridging the metric and differential approaches. Here, a natural direction to look for is the *weighted relational semantics* [50], due to its strict relations with the Taylor expansion of programs. However, in this semantics, arbitrary terms correspond to power series, and terms with bounded applications correspond to polynomials, hence in any case to functions which are *not* Lipschitz.

Yet, what if such polynomials were tropical ones, i.e. piecewise linear functions? This way, the Taylor formula could really be interpreted as a decomposition of λ -terms via limits (indeed, infs) of Lipschitz maps. In other words, unbounded term application could be seen as a limit of *more and more sensitive* operations.



■ **Figure 1** Tropical polynomials $\varphi_0, \dots, \varphi_4$ (top to bottom), and their limit tps φ (in violet). The points where the slope changes are the tropical roots of φ , i.e. the points $x = 2^{-(i+1)}$, satisfying $ix + 2^{-i} = (i+1)x + 2^{-(i+1)}$.

This viewpoint, that we develop in the following sections, leads to the somehow unexpected discovery of a bridge between the metric and differential study of higher-order programs. This connection not only suggests the application of optimization methods based on tropical mathematics to the study of the λ -calculus and its quantitative extensions, but it scales to a more abstract level, leading to introduce a differential operator for continuous functors between *generalized* metric spaces (in the sense of [51]), as shown in Section 8.

3 Tropical Polynomials and Power Series

At the basis of our approach is the observation that the *tropical semiring* $([0, \infty], \min, +)$, which is at the heart of tropical mathematics, coincides with the *Lawvere quantale* $\mathbb{L} = ([0, \infty], \geq, +)$ [47, 74], the structure at the heart of the categorical study of metric spaces initiated by Lawvere himself [51]. Let us recall that a quantale is a complete lattice endowed with a continuous monoid action. In the case of \mathbb{L} , the lattice is defined by the reverse order \geq on \mathbb{R} , and the monoid action is provided by addition. Notice that the lattice join operation of \mathbb{L} coincides with the idempotent semiring operation \min .

Power series and polynomials over the tropical semiring are defined as follows:

► **Definition 1.** A tropical power series (tps) in k variables is a function $f : \mathbb{L}^k \rightarrow \mathbb{L}$ of shape $f(x) = \inf_{i \in I} \{ix + \hat{f}(i)\}$, where $I \subseteq \mathbb{N}^k$, ix is the scalar product and $\hat{f} \in \mathbb{L}^{\mathbb{N}^k}$ is a vector of coefficients. When I is finite, f is called a tropical polynomial.

Hence, a unary tps is a function $f : \mathbb{L} \rightarrow \mathbb{L}$ of the form $f(x) = \inf_{i \in I} \{ix + a_i\}$, with $I \subseteq \mathbb{N}$ and the $a_i \in \mathbb{L}$. In Section 5 we also consider tps in *infinitely many* variables.

A tropical polynomial is always a piece-wise linear function since, e.g. in one variable, it has shape $f(x) = \min_{0 \leq j \leq n} \{i_j x + c_{i_j}\}$. For example, the polynomials $\varphi_n(x) = \min_{0 \leq j \leq n} \{jx + 2^{-j}\}$ are illustrated in Fig. 1 for $0 \leq n \leq 4$.

A *tropical root* of a tps φ is a point $x \in \mathbb{L}$ where φ is not differentiable (i.e. where the slope of φ changes). When φ is a polynomial, the roots of φ coincide with the points where the minimum defining φ is attained at least twice (see Figure 1). Unlike in standard algebra, tropical roots of tropical polynomials can be computed in linear time [62].

While tropical polynomials are essentially combinatorial objects, this cannot be said for tps: since infs are not in general mins, a tps is a “limit” of tropical polynomials of higher and higher degree, and its behavior is in general way more difficult to study than that of tropical polynomials [61]. E.g., the tps $\varphi(x) := \inf_{n \in \mathbb{N}} \{nx + 2^{-n}\}$ (see Fig. 1) is the “limit” of the polynomials φ_n .

► **Remark 2.** There is a well-known relation between tropical polynomial/power series and usual polynomials/power series. For a power series $f(x) = \sum_{i \in I} a_i x^i$ (polynomials being the case for I finite) defined on $[0, 1]$ and coefficients a_i in $[0, 1]$, one can define the *tropicalisation* $\mathfrak{t}f : \mathbb{L} \rightarrow \mathbb{L}$ of f as the tropical polynomial/power series function $\mathfrak{t}f(\alpha) := \inf_{i \in I} \{-\log(a_i) + i\alpha\}$.

For instance, the tropicalisation sends the infinite power series $\sum_{n \geq 2} x^n$ to the tps $\varphi(x) = \inf_n (n + 2)x$. Yet the latter always coincides with the tropical monomial $\varphi(x) = 2x$, since $nx \geq 0$ for all $n \in \mathbb{N}$, which is in turn the tropicalisation of the polynomial x^2 . This shows that tropicalisation is *not* in general an injective operation and in fact, as we show in Theorem 3 below, tps (in finitely many variables) have a tendency to collapse, if not globally at least locally, onto tropical polynomials.

For instance, by looking at Fig. 1 it appears that, *far from 0*, φ behaves like some of the polynomials φ_n . In particular, φ coincides on $[\epsilon, \infty]$ with φ_n , for $\epsilon \geq 2^{-(n+1)}$ (the smallest tropical root of φ_n). However, at $x = 0$ we have that $\varphi(x = 0) = \inf_{n \in \mathbb{N}} 2^{-n} = 0$, and this is the only point where the inf is *not* a min. Also, while the derivative of f is bounded on all $(0, \infty)$, for $x \rightarrow 0^+$ it tends to ∞ . In fact, this is a general phenomenon, as showed below:

► **Theorem 3.** *For all tps $f(x) = \inf_{n \in \mathbb{N}^k} \{nx + \hat{f}(n)\}$, for all $0 < \epsilon < \infty$, there is a finite $\mathcal{F}_\epsilon \subseteq \mathbb{N}^k$ such that f coincides on all $[\epsilon, \infty]^k$ with $P_\epsilon(x) := \min_{n \in \mathcal{F}_\epsilon} \{nx + \hat{f}(n)\}$.*

As we'll see, the potential of collapsing infinitary objects (i.e. tps) into combinatorial ones (i.e. tropical polynomials), is one of the most intriguing features of tropical semantics.

For the interested reader, we provide the proof of Theorem 3 below. Let us first set the following:

► **Definition 4.** *Let \leq be the product order on \mathbb{N}^K (i.e. for all $m, n \in \mathbb{N}^K$, $m \leq n$ iff $m_i \leq n_i$ for all $1 \leq i \leq K$). Of course $m < n$ holds exactly when $m \leq n$ and $m_i < n_i$ for at least one $1 \leq i \leq K$. Finally, we set $m <_1 n$ iff $m < n$ and $\sum_{i=1}^K n_i - m_i = 1$ (i.e. they differ on exactly one coordinate).*

► **Remark 5.** If $U \subseteq \mathbb{N}^K$ is infinite, then U contains an infinite ascending chain $m_0 < m_1 < m_2 < \dots$. This is a consequence of König Lemma (KL): consider the directed acyclic graph $(U, <_1)$, indeed a K -branching tree; if there is no infinite ascending chain $m_0 < m_1 < m_2 < \dots$, then in particular there is no infinite ascending chain $m_0 <_1 m_1 <_1 m_2 <_1 \dots$ so the tree U has no infinite ascending chain; then by KL it is finite, contradicting the assumption.

Proof of Theorem 3. We will actually show the existence of $\mathcal{F}_\epsilon \subseteq_{\text{fin}} \mathbb{N}^k$ such that:

1. if $\mathcal{F}_\epsilon = \emptyset$ then $f(x) = +\infty$ for all $x \in \mathbb{L}^k$;
2. if $f(x_0) = +\infty$ for some $x_0 \in [\epsilon, \infty)^K$ then $\mathcal{F}_\epsilon = \emptyset$;
3. the restriction of f on $[\epsilon, \infty]^k$ coincides with $P_\epsilon(x) := \min_{n \in \mathcal{F}_\epsilon} \{nx + \hat{f}(n)\}$.

Let \mathcal{F}_ϵ be the complementary in \mathbb{N} of the set:

$$\{n \in \mathbb{N}^K \mid \text{either } \hat{f}(n) = +\infty \text{ or there is } m < n \text{ s.t. } \hat{f}(m) \leq \hat{f}(n) + \epsilon\}.$$

In other words, $n \in \mathcal{F}_\epsilon$ iff $\hat{f}(n) < +\infty$ and for all $m < n$, one has $\hat{f}(m) > \hat{f}(n) + \epsilon$. Suppose that \mathcal{F}_ϵ is infinite; then, using Remark 5, it contains an infinite ascending chain $\{m_0 < m_1 < \dots\}$. By definition of \mathcal{F}_ϵ we have then $+\infty > \hat{f}(m_0) > \hat{f}(m_1) + \epsilon > \hat{f}(m_2) + 2\epsilon > \dots$, so that $+\infty > \hat{f}(m_0) > \hat{f}(m_i) + i\epsilon \geq i\epsilon$ for all $i \in \mathbb{N}$. This contradicts the Archimedean property of \mathbb{R} . Hence \mathcal{F}_ϵ is finite.

1. We show that if $\mathcal{F}_\epsilon = \emptyset$, then $\hat{f}(n) = +\infty$ for all $n \in \mathbb{N}^K$. This immediately entails the desired result. We go by induction on the well-founded order $<$ over $n \in \mathbb{N}^K$:
 - if $n = 0^K \notin \mathcal{F}_\epsilon$, then $\hat{f}(n) = +\infty$, because there is no $m < n$.
 - if $n \notin \mathcal{F}_\epsilon$, with $n \neq 0^K$ then either $\hat{f}(n) = +\infty$ and we are done, or there is $m < n$ s.t. $\hat{f}(m) \leq \hat{f}(n) + \epsilon$. By induction $\hat{f}(m) = +\infty$ and, since $\epsilon < +\infty$, this entails $\hat{f}(n) = +\infty$.
2. If $f(x_0) = +\infty$ with $x_0 \in [\epsilon, \infty)^K$, then necessarily $\hat{f}(n) = +\infty$ for all $n \in \mathbb{N}^K$. Therefore, no $n \in \mathbb{N}^K$ belongs to \mathcal{F}_ϵ .
3. We have to show that $f(x) = P_\epsilon(x)$ for all $x \in [\epsilon, +\infty]^K$. By 1), it suffices to show that we can compute $f(x)$ by taking the inf, that is therefore a min, only in \mathcal{F}_ϵ (instead of all \mathbb{N}^K). If $\mathcal{F}_\epsilon = \emptyset$ then by 2) we are done (remember that $\min \emptyset := +\infty$). If $\mathcal{F}_\epsilon \neq \emptyset$, we show that for all $n \in \mathbb{N}^K$, if $n \notin \mathcal{F}_\epsilon$, then there is $m \in \mathcal{F}_\epsilon$ s.t. $\hat{f}(m) + mx \leq \hat{f}(n) + nx$. We do it again by induction on $<_1$:
 - if $n = 0^K$, then from $n \notin \mathcal{F}_\epsilon$, by definition of \mathcal{F}_ϵ , we have $\hat{f}(n) = +\infty$ (because there is no $n' < n$). So any element of $\mathcal{F}_\epsilon \neq \emptyset$ works.
 - if $n \neq 0^K$, then we have two cases: either $\hat{f}(n) = +\infty$, in which case we are done as before by taking any element of $\mathcal{F}_\epsilon \neq \emptyset$. Or $\hat{f}(n) < +\infty$, in which case (again by definition of \mathcal{F}_ϵ) there is $n' < n$ such that $\hat{f}(n') \leq \hat{f}(n) + \epsilon$ (\star). Therefore we have (remark that the following inequalities hold also for the case $x = +\infty$):

$$\begin{aligned}
 \hat{f}(n') + n'x &\leq \hat{f}(n) + \epsilon + n'x && \text{by } (\star) \\
 &\leq \hat{f}(n) + (n - n')x + n'x && \text{because } \epsilon \leq \min x \text{ and } \min x \leq (n - n')x \\
 &= \hat{f}(n) + nx.
 \end{aligned}$$

Now, if $n' \in \mathcal{F}_\epsilon$ we are done. Otherwise $n' \notin \mathcal{F}_\epsilon$ and we can apply the induction hypothesis on it, obtaining an $m \in \mathcal{F}_\epsilon$ s.t. $\hat{f}(m) + mx \leq \hat{f}(n') + n'x$. Therefore this m works. \blacktriangleleft

4 Tropical Semantics and First Order Effectful Programs

Before discussing how full-scale higher-order programming languages can be interpreted in terms of tropical power series, we highlight how such functions may naturally arise in the study of effectful programming languages. We will see that, when considering probabilistic and non-deterministic programs, tropical tools can be used to describe the behavior of programs in *the best/worst case*, and may lead to collapse the description of infinitely many possible behaviors into a combinatorial account of the optimal ones.

Maximum Likelihood Estimators for Probabilistic Languages. Let us start with a very basic probabilistic language: the terms are $M ::= \text{True} \mid \text{False} \mid M \oplus_p M$, for $p \in [0, 1]$, and the operational semantics is $M \oplus_p N \rightarrow pM$ and $M \oplus_p N \rightarrow (1 - p)N$, so that $M \oplus_p N$ plays the role of a probabilistic coin toss of bias p . Consider the program $M := (\text{True} \oplus_p \text{False}) \oplus_p ((\text{True} \oplus_p \text{False}) \oplus_p (\text{False} \oplus_p \text{True}))$. Calling $q = 1 - p$, to each occurrence of **True** or **False** in M , univocally determined by an address $\omega \in \{l, r\}^*$, is associated a monomial $P_\omega(p, q)$ which determines the probability of the event “ $M \rightarrow_\omega \text{True/False}$ ”, that is, that M reduces to **True/False** according to the choices in ω . Thinking of p, q as parameters, $P_\omega(p, q)$ can thus be read as the *likelihood function* of the event “ $M \rightarrow_\omega \text{True/False}$ ”. For instance, we have $P_{rll}(p, q) := qp^2$, $P_{rrr}(p, q) := q^3$, and $P_{rrl}(p, q) = P_{rlr}(p, q) := q^2p$. The polynomial function $Q_{\text{True}}(p, q) := P_{ll}(p, q) + P_{rll}(p, q) + P_{rrr}(p, q) = p^2 + p^2q + q^3$ gives instead the probability of the event “ $M \Rightarrow \text{True}$ ”, and analogously for $Q_{\text{False}}(p, q) := P_{lr}(p, q) + P_{rrl}(p, q) + P_{rtr}(p, q) = pq + 2pq^2$.

This way, the probabilistic evaluation of M is presented as a *hidden Markov model* [14], a fundamental statistical model, and notably one to which tropical methods are generally applied [64]. Then, a natural question in this case, for a fixed ω_0 , is the following: knowing that M reached a normal form, say **True**, what is the *maximum likelihood estimator* for the event “ $M \rightarrow_{\omega_0} \mathbf{True}$ ”? In other words, what is the choice of p, q that maximizes the conditional probability $\mathbb{P}(M \rightarrow_{\omega_0} \mathbf{True} \mid M \rightarrow \mathbf{True})$, i.e. that makes ω_0 the most likely path among those leading to **True**?

Answering this question amounts at finding a solution to the following constrained maximization problem in the unknown $p \in [0, 1]$:

$$P_{\omega_0}(p, 1 - p) = \max_{\omega} P_{\omega}(p, 1 - p)$$

which is related to the polynomial, say, $Q_{\mathbf{True}}(p, q)$. Since $-\log P_{\omega}(p, q) = (\mathbf{t}P_{\omega})(-\log p, -\log q)$, this is equivalent to finding a solution of the following constrained *minimization* problem in the unknowns $(p, x, y) \in [0, 1] \times [0, \infty] \times [0, \infty]$:

$$\mathbf{t}P_{\omega_0}(x, y) = \mathbf{t}Q_{\mathbf{True}}(x, y), \quad x = -\log p, \quad y = -\log(1 - p). \quad (\star)$$

Since the first equation can be solved easily (i.e. in linear time) by computing the tropical roots of $\mathbf{t}Q_{\mathbf{True}}$, we can obtain an explicit relation between x and y that can be used to solve the whole system, finally finding our p , as the next example shows.

► **Example 6.** For our running example M , let us suppose that we observed the event “ $M \rightarrow \mathbf{True}$ ”, so that our probabilities are conditioned under this observation. We have

$$\mathbf{t}Q_{\mathbf{True}}(x, y) = \min\{\mathbf{t}P_{ll}(x, y), \mathbf{t}P_{rll}(x, y), \mathbf{t}P_{rrr}(x, y)\} = \min\{2x, y + 2x, 3y\}.$$

The tropical roots of $\mathbf{t}Q_{\mathbf{True}}(x, y)$ are all the points of the form $(x, \frac{2}{3}x)$. Recall that these are the points where (\star) is satisfied for *at least two* distinct values of ω_0 (indeed for $\omega_0 \in \{ll, rrr\}$). From this it follows that $\mathbf{t}Q_{\mathbf{True}}(x, y) = \mathbf{t}P_{rrr}(x, y) = 3y$ holds iff $y \leq \frac{2}{3}x$, and $\mathbf{t}Q_{\mathbf{True}}(x, y) = 2x = \mathbf{t}P_{ll}(x, y)$ otherwise. In this way can find the maximum likelihood estimator for $\omega_0 = rrr$: via the substitution $x := -\log p$, $y := -\log(1 - p)$, the condition $y \leq \frac{2}{3}x$ is equivalent to $-\log(1 - p) \leq -\frac{2}{3}\log p$, i.e. $1 - p \geq p^{\frac{2}{3}}$. This means that, if $p \in [0, 1]$ satisfies $1 - p \geq p^{\frac{2}{3}}$ (for example, $p = \frac{1}{4}$), then $P_{rrr}(p, 1 - p) = \max_{\omega=ll, rll, rrr} P_{\omega}(p, 1 - p)$. In other words, knowing that M sampled **True** in its normal form, the most likely sampled occurrence of **True** is the one at the address rrr iff $1 - p \geq p^{\frac{2}{3}}$.

As we’ll see in Section 5, this analysis extends to PCF-style programs. For example, the program $M = \mathbf{Y}(\lambda x. \mathbf{True} \oplus_p x)$ yields the power series $Q_{\mathbf{True}}(p, q) = \sum_{n=0}^{\infty} pq^n = \frac{p}{1-q}$ that sums all *infinitely many* ways in which M may reduce to **True**. Notice that the tropicalised series $\mathbf{t}Q_{\mathbf{True}}(-\log p, -\log q) = \inf_{n \in \mathbb{N}} \{-\log p - n \log q\} = -\log p$ collapses onto a single monomial describing the *unique* most likely reduction path of M leading to **True**, namely the one that passes through a coin toss only once.

Best Case Analysis for Non-Deterministic Languages. This example is inspired from [50]. We consider now a basic non-deterministic language with terms $M ::= \mathbf{True} \mid \mathbf{Gen} \mid M + M$, with an operation semantics comprising a non-deterministic reduction rule $M_1 + M_2 \xrightarrow{\alpha} M_i$ and a generation rule $\mathbf{Gen} \xrightarrow{\beta} \mathbf{True} + \mathbf{Gen}$, where in each case the value $\alpha, \beta \in \mathbb{L}$ indicates a *cost* associated with the reduction (e.g. the estimated clock value for the simulation of each reduction on a given machine model). Then, any reduction $\omega : M \rightarrow N$ of a term to (one of its) normal form is associated with a tropical monomial $P_{\omega}(\alpha, \beta)$ consisting of

the sum of the costs of all reductions in ω . For a given normal form N , the reductions $\omega_i : M \rightarrow N$ give rise to a tps $\inf_{i \in I} P_{\omega_i}(\alpha, \beta)$. For example, consider the non-deterministic term $M := \mathbf{Gen} + ((\mathbf{True} + \mathbf{True}) + \mathbf{Gen})$. The (infinitely many) reduction paths leading to \mathbf{True} can be grouped as follows:

- left, then reduce \mathbf{Gen} $n + 1$ -times, then left;
- right, then left and then either left or right;
- right twice, then reduce \mathbf{Gen} $n + 1$ -times and then left.

This leads to the tps $\varphi_{M \rightarrow \mathbf{True}}(\alpha) = \inf_{n \in \mathbb{N}} \{(n + 2)\alpha + (n + 1)\beta, 3\alpha, (n + 3)\alpha + (n + 1)\beta\} = \min\{2\alpha + \beta, 3\alpha\}$, which describes all possible behaviors of M . Notice that, since α and β are always positive, the power series $\varphi_{M \rightarrow \mathbf{True}}(\alpha)$ is indeed equivalent to the tropical polynomial $\min\{2\alpha + \beta, 3\alpha\}$. In other words, of the infinitely many behaviors of M , only finitely many have chances to be *optimal*: either left + \mathbf{Gen} + left, or right + left + (left,right). Also in this case, reducing to best-case analysis leads to collapse the infinitary description of *all* behaviors to a purely combinatorial description of the finitely many optimal ones.

Once reduced $\varphi_{M \rightarrow \mathbf{True}}$ to a polynomial, the best behavior among these will depend on the values of α and β , and by studying the tropical polynomial $\varphi_{M \rightarrow \mathbf{True}}$ one can thus answer questions analogous to 2) above, that is, what are the best choices of costs α, β making a *chosen* reduction of M to \mathbf{True} the cheapest one?

5 Tropical Semantics of Higher-Order Programs

In this section we first recall a general and well-known construction that yields, for any *continuous* semiring Q , a model $Q\mathbf{Rel}_!$ of effectful extensions of the simply typed λ -calculus and PCF, and we show how, when $Q = \mathbb{L}$, it captures optimal program behavior; moreover, we discuss how this model adapts to *graded* and *differential* variants of STLC.

Linear/Non-Linear Algebra on Q -Modules. For any *continuous* semiring Q (i.e. a cpo equipped with an order-compatible semiring structure), one can define a category $Q\mathbf{Rel}$ ([50] calls it $Q^{\mathbb{I}}$) of “ Q -valued matrices” as follows: $Q\mathbf{Rel}$ has sets as objects and set-indexed matrices with coefficients in Q as morphisms, i.e. $Q\mathbf{Rel}(X, Y) := Q^{X \times Y}$. The composition $st \in Q^{X \times Z}$ of $t \in Q^{X \times Y}$ and $s \in Q^{Y \times Z}$ is given by $(st)_{a,c} := \sum_{b \in Y} s_{b,c} t_{a,b}$ (observe that this series always converges because Q is continuous). For any set X , Q^X is a Q -semimodule and we can identify $Q\mathbf{Rel}(X, Y)$ with the set of linear maps from Q^X to Q^Y , which have shape $f(x)_b := \sum_{a \in X} \hat{f}_{a,b} x_a$, for some matrix $\hat{f} \in Q^{X \times Y}$. Notice that usual linear algebra conventions correspond to work in $Q\mathbf{Rel}^{\text{op}}$, e.g. the usual matrix-vector product defines a map $Q^Y \rightarrow Q^X$. Following [29, 47, 50], we are instead working with transpose matrices.

$Q\mathbf{Rel}$ admits a comonad $!$ which acts on objects by taking the finite multisets. Remember that the coKleisli category $\mathcal{C}_!$ of a category \mathcal{C} w.r.t. a comonad $!$ is the category whose objects are the same of \mathcal{C} , and $\mathcal{C}_!(X, Y) := \mathcal{C}(!X, Y)$, with composition $\circ_!$ defined via the co-multiplication of $!$. Now, although a matrix $t \in Q\mathbf{Rel}_!(X, Y)$ yields a linear map $\mathbb{L}^{!X} \rightarrow \mathbb{L}^Y$, by exploiting the coKleisli structure we can also “express it in the base X ”, which leads to the *non-linear* map $t^! : Q^X \rightarrow Q^Y$ defined by the power series $t^!(x) = t \circ_! x : b \mapsto \sum_{\mu \in !X} t_{\mu,b} \cdot x^\mu$, where $x^\mu = \prod_{a \in x} x_a^{\mu(a)}$.

When we instantiate $Q = \mathbb{L}$, we obtain the category $\mathbb{L}\mathbf{Rel}$ of \mathbb{L} -valued matrices. As one might expect, this category is tightly related to Lawvere’s theory of (generalized) metric spaces. For the moment, let us just observe that a (possibly ∞) metric on a set X is nothing but a “ \mathbb{L} -valued square matrix” $d : X \times X \rightarrow \mathbb{L}$ satisfying axioms like e.g. the triangular law. We will come back to this viewpoint in Section 8.

Composition in $\mathbb{L}\text{Rel}$ reads as $(st)_{a,c} := \inf_{b \in Y} \{s_{b,c} + t_{a,b}\}$, and the non-linear maps $t^! : \mathbb{L}^X \rightarrow \mathbb{L}^Y$ have shape $t^!(x)_b = \inf_{\mu \in !X} \{\mu x + t_{\mu,b}\}$, where $\mu x := \sum_{a \in X} \mu(a)x_a$. These correspond to the generalisation of tps with *possibly infinitely* many variables (in fact, as many as the elements of X). By identifying $!*\} \simeq \mathbb{N}$ and $\mathbb{L}^{*\} \simeq \mathbb{L}$, the tps generated by the morphisms in $\mathbb{L}\text{Rel}_!(\{*\}, \{*\})$ are exactly the usual tps's of one variable. For example, the φ of Figure 1 is indeed of shape $\varphi = t^!$, for $t \in \mathbb{L}^{!\{*\} \times \{*\}}$, $t_{\mu,*} := 2^{-\#\mu}$.

► **Remark 7.** The operation $f \mapsto f^!$ turning a matrix into a function is reminiscent of the well-known operation of taking the *convex conjugate* f^* of a function f defined over a vector space (itself a generalization of the Legendre transformation). Indeed, let X, Y be sets and let $\langle _, _ \rangle : X \times Y \rightarrow \mathbb{R}$. For $f : X \rightarrow \mathbb{R}$, let $f^* : Y \rightarrow \mathbb{R}$ be defined by $f^*(y) := \sup_{x \in X} \{\langle x, y \rangle - f(x)\}$. Then for $X = !A, Y = \mathbb{L}^A$, where A is a set, and $\langle \mu, y \rangle := \mu y$, we have $f^!(y) = (-f)^*(-y)$ for all $f \in \mathbb{L}^A$.

The \mathbb{L} -Weighted Relational Model. The categories $Q\text{Rel}$ are well-known to yield a model of both probabilistic and non-deterministic versions of PCF (see e.g. [34, 50]), which are called *weighted relational models*. The interpretation of the simply typed λ -calculus STLC in $\mathbb{L}\text{Rel}$ relies on the fact that all categories $Q\text{Rel}_!$ are cartesian closed [50], with cartesian product and exponential objects acting on objects X, Y as, respectively, $X + Y$ and $!X \times Y$. Hence, any typable term $\Gamma \vdash M : A$ gives rise to a morphism $\llbracket \Gamma \vdash M : A \rrbracket \in \mathbb{L}\text{Rel}_!(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket)$, and thus to a generalized tps $\llbracket \Gamma \vdash M : A \rrbracket^! : \mathbb{L}^{\llbracket \Gamma \rrbracket} \rightarrow \mathbb{L}^{\llbracket A \rrbracket}$.

► **Example 8.** The evaluation morphism $\text{ev} \in \mathbb{L}\text{Rel}_!((!X \times Y) + X, Y)$ yields the tps $\text{ev}^! : \mathbb{L}^{!X \times Y} \times \mathbb{L}^X \rightarrow \mathbb{L}^Y$ given by $b \in Y \mapsto \text{ev}^!(F, x)_b = \inf_{\mu} \{F_{\mu,b} + \mu x\}$. So, for instance, supposing the ground type o of STLC is interpreted as the singleton set $\{*\}$, and recalling the identification $!*\} \simeq \mathbb{N}$, the interpretation of the term $x : o, z : (o \rightarrow o \rightarrow o) \vdash zxx : o$, involving two consecutive evaluations, yields the tps $\varphi : \mathbb{L} \times \mathbb{L}^{\mathcal{M}_{\text{fin}}(\mathbb{N} \times \mathbb{N})} \rightarrow \mathbb{L}$ given by $\varphi(x, z) = \inf_{n, n'} \{z_{[(n, n')]} + (n + n')x\}$.

This interpretation extends to PCF by interpreting the fixpoint combinator \mathbf{Y} via the matrices $\text{fix}^X = \inf_n \{\text{fix}_n^X\} \in \mathbb{L}^{!(X \times X) \times X}$, where $\text{fix}_0^X = 0$ and $\text{fix}_{n+1}^X = \text{ev} \circ_! \langle \text{fix}_n^X, \text{id} \rangle$.

One can easily check, by induction on a typing derivation, that for any program of STLC or PCF, the associated matrix is *discrete*, that is, its values are included in $\{0, \infty\}$. Indeed, as suggested in Section 3, the actual interest of tropical semantics lies in the interpretation of effectful programs. As the homsets $\mathbb{L}\text{Rel}_!(X, Y)$ are \mathbb{L} -modules, it is possible to interpret in it extensions of STLC and PCF comprising \mathbb{L} -module operations $\alpha \cdot M$ and $M + N$ [50], by letting $\llbracket \Gamma \vdash \alpha \cdot M : A \rrbracket = \llbracket \Gamma \vdash M : A \rrbracket + \alpha$ and $\llbracket \Gamma \vdash M + N : A \rrbracket = \min\{\llbracket \Gamma \vdash M : A \rrbracket, \llbracket \Gamma \vdash N : A \rrbracket\}$. More precisely, [50] considers a language PCF^Q corresponding to PCF extended with Q -module operations, with an operational semantics given by rules $M \xrightarrow{1} M'$ for each rule $M \rightarrow M'$ of PCF (here 1 is the monoidal unit of Q) as well as $M_1 + M_2 \xrightarrow{1} M_i$ and $\alpha \cdot M \xrightarrow{\alpha} M$. Hence, any reduction $\omega = \rho_1 \dots \rho_k : M \rightarrow N$ is naturally associated with a weight $w(\omega) = \sum_{i=1}^k w(\rho_i) \in Q$. In particular, from [Theorem V.6] [50] we deduce the following adequation result:

► **Proposition 9.** $\llbracket \vdash_{\text{PCF}^L} M : \text{Nat} \rrbracket_n = \inf \{w(\omega) \mid \omega : M \rightarrow \underline{n}\}$ for all $n \in \mathbb{N}$.

The previous result allows to relate the tropical semantics of a program with its best-case operational behavior. Observe that the two examples shown in Section 3 can easily be rephrased in the language $\text{PCF}^{\mathbb{L}}$. For instance, for the probabilistic example, one can use the translation $(M \oplus_p N)^\circ = \min\{M^\circ + p, N^\circ + (1-p)\}$, so that the reductions $M \oplus_p N \xrightarrow{p} M$ and $M \oplus_p N \xrightarrow{1-p} N$ translate into a sequence of two reductions $(M \oplus_p N)^\circ \xrightarrow{0} M^\circ \xrightarrow{p} M$ and

$x :_1 A \vdash x : A$	
$\frac{\Gamma \vdash M : A}{\Gamma, x :_0 B \vdash M : A}$	$\frac{\Gamma, x :_n B, y :_m B \vdash M : A}{\Gamma, x :_{n+m} B \vdash M\{x/y\} : A}$
$\frac{\Gamma, x :_n A \vdash M : B}{\Gamma \vdash \lambda x.M : !_n A \multimap B}$	$\frac{\Gamma \vdash M : !_n A \multimap B \quad \Delta \vdash N : A}{\Gamma + n\Delta \vdash MN : B}$

■ **Figure 2** Typing rules for bSTLC.

$(M \oplus_p N)^\circ \xrightarrow{0} N^\circ \xrightarrow{1-p} N^\circ$. Let PPCF (for *probabilistic* PCF [34]) be standard PCF extended with the constructor $M \oplus_p N$ ($p \in [0, 1]$) and its associated reduction rules. From the above discussion we deduce the following:

► **Corollary 10.** *Let $\vdash_{\text{PPCF}} M : \text{Nat}$ and $n \in \mathbb{N}$. Considering its interpretation as a function of $p, 1 - p$, we have that $\llbracket \vdash_{\text{PPCF}} M^\circ : \text{Nat} \rrbracket_n(-\log(p), -\log(1 - p))$ is the minimum negative log-probability of any reduction from M to \underline{n} , i.e. the negative log-probability of (any of) the (equiprobable) most likely reduction path from M to n .*

Remark that this implies that all solution $p \in [0, 1]$ to the equation $-\log w(\omega) = \llbracket \vdash_{\text{PPCF}^\mathbb{L}} M^\circ : \text{Nat} \rrbracket_n(-\log(p), -\log(1 - p))$ are the values of the probabilistic parameter which make the reduction ω the most likely.

► **Remark 11.** The function $\llbracket \vdash_{\text{PPCF}} M^\circ : \text{Nat} \rrbracket_n(\alpha, \beta)$ is a tps, and Theorem 3 ensures that this function coincides *locally* with a tropical polynomial. This means that, for any choice of $p, 1 - p$, the most likely reduction path of M can be searched for within a *finite* space.

Finally, [50] obtained a similar result for a non-deterministic version of PCF, by translating each term into $\text{PCF}^\mathbb{L}$ via $(\lambda x.M)^\circ = \lambda x.M^\circ + 1$ and $(\mathbf{Y}M)^\circ = \mathbf{Y}(M^\circ + 1)$ ([50] considers the discrete tropical semiring $\mathbb{N} \cup \{\infty\}$, but the result obviously transports to \mathbb{L}), and in that case [50, Corollary VI.10] gives that $\llbracket \vdash M^\circ : \text{Nat} \rrbracket_n$ computes the minimum number of β - and fix- redexes reduced in a reduction sequence from M to \underline{n} .

N-Graded Types. We now show how to interpret in $\mathbb{L}\text{Rel}$ a graded version of STLC, that we call bSTLC, indeed a simplified version of the well-studied language Fuzz [68]. This language is based on a graded exponential $!_n A$, corresponding to the possibility of using an element of type A *at most* n times. In particular, if a function $\lambda x.M$ of type $!_n A \multimap B$, then, for any N of type A , x is duplicated *at most* n times in any reduction of $(\lambda x.M)N$ to the normal form.

Graded simple types are defined by $A ::= o \mid !_n A \multimap A$; the contexts of the typing judgements are sets of declarations of the form $x :_n A$, with $n \in \mathbb{N}$; given two contexts Γ, Δ , we define $\Gamma + \Delta$ recursively as follows: if Γ and Δ have no variable in common, then $\Gamma + \Delta = \Gamma \cup \Delta$; otherwise, we let $(\Gamma, x :_m A) + (\Delta, x :_n A) = (\Gamma + \Delta), x :_{m+n} A$. Moreover, for any context Γ and $m \in \mathbb{N}$, we let $m\Gamma$ be made all $x :_{mn} A$ for $(x :_n A) \in \Gamma$. The typing rules of bSTLC are illustrated in Fig. 2,

On the side of $\mathbb{L}\text{Rel}$, one can check that the comonad $!$ of $\mathbb{L}\text{Rel}$ can be “decomposed” into a family of “graded exponentials functors” $!_n : \mathbb{L}\text{Rel} \rightarrow \mathbb{L}\text{Rel}$ ($n \in \mathbb{N}$), where $!_n X$ is the set of multisets on X of cardinality *at most* n . The sequence $(!_n)_{n \in \mathbb{N}}$ gives rise to a so-called *N-graded linear exponential comonad* on (the SMC) $\mathbb{L}\text{Rel}$ [49]. As such, $(\mathbb{L}\text{Rel}, (!_n)_{n \in \mathbb{N}})$ yields then a model of bSTLC. More details can be found in the extended version of this article.

$\frac{}{\Gamma \vdash 0 : A}$	$\frac{}{\Gamma, x : A \vdash x : A}$
$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x.M : A \rightarrow B}$	$\frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash \mathbb{T} : A}{\Gamma \vdash M\mathbb{T} : B}$
$\frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash D[M, N] : A \rightarrow B}$	$\frac{\Gamma \vdash M_1 : A \cdots \Gamma \vdash M_n : A}{\Gamma \vdash M_1 + \cdots + M_n : A} \quad (n \geq 2)$

■ **Figure 3** Typing rules of ST ∂ LC.

Notice that arrow types are interpreted via $\llbracket !_n A \multimap B \rrbracket := !_n \llbracket A \rrbracket \times \llbracket B \rrbracket$ and that, whenever $\llbracket * \rrbracket$ is finite, the set $\llbracket A \rrbracket$ is finite for any type A of bSTLC.

The Differential λ -Calculus. We recall the interpretation in $\mathbb{LRel}_!$ of the simply typed *differential* λ -calculus ST ∂ LC, an extension of STLC ensuring exact control of duplications. The syntax of ST ∂ LC (see [20, Section 3]) is made of *terms* M and *sums* \mathbb{T} , mutually generated by: $M ::= x \mid \lambda x.M \mid M\mathbb{T} \mid D[M, M]$ and $\mathbb{T} ::= 0 \mid M \mid M + \mathbb{T}$, quotiented by equations that make $+, 0$ form a commutative monoid on the set of sums, by linearity of $\lambda x.(_)$, $D[_, _]$ and $(_) \mathbb{T}$ (but *not* of $M(_)$) and by irrelevance of the order of consecutive $D[_, _]$. We follow the tradition of quotienting also for the idempotency of $+$. The typing rules are illustrated in Figure 3, where a context Γ is a list of typed variable declarations. The main feature of this language is that $D^n[\lambda x.M, N^n]0$ has a non-zero normal form iff x is duplicated exactly n times during reduction.

The categorical models of ST ∂ LC are called *cartesian closed differential λ -categories* (CC ∂ λ C) [16, 17, 20]. These are CCCs enriched over commutative monoids (i.e. morphisms are summable and there is a 0 morphism), with the cartesian closed structure compatible with the additive structure, and equipped with a certain *differential operator* D , turning a morphism $f : A \rightarrow B$ into a morphism $Df : A \times A \rightarrow B$, and generalising the usual notion of differential, see e.g. [15]. An example is the CC ∂ λ C of convenient vector spaces with smooth maps, where D is the “real” differential of smooth maps.

Applying [52, Theorem 6.1] one can check that $\mathbb{LRel}_!$ is a CC ∂ λ C when equipped with $D : \mathbb{LRel}_!(X, Y) \rightarrow \mathbb{LRel}_!(X \& X, Y)$ defined as $(Dt)_{\mu \oplus \rho, b} = t_{\rho + \mu, b}$ if $\#\mu = 1$ and as ∞ otherwise (using the isomorphism $(\mu, \rho) \in !Z \times !Z' \mapsto \mu \oplus \rho \in !(Z + Z')$). The differential operator D of $\mathbb{LRel}_!$ translates into a differential operator $D_!$ turning a tps $f : \mathbb{L}^X \rightarrow \mathbb{L}^Y$ into a tps $D_!f : \mathbb{L}^X \times \mathbb{L}^X \rightarrow \mathbb{L}^Y$, linear in its first variable, and given by $D_!f(x, y)_b = \inf_{a \in X, \mu \in !X} \left\{ \hat{f}_{\mu+a, b} + x_a + \mu y \right\}$. One can check that, when f is a tropical polynomial, $D_!f$ coincides with the standard tropical derivative (see e.g. [45]).

6 On Tropical Power Series

As seen in Section 3, tropical polynomials are piecewise-linear functions, hence concave and Lipschitz-continuous. Moreover, tps in finitely many variables are locally equivalent to tropical polynomials (except at some singular points), and are thus also concave and locally Lipschitz-continuous. In this section we show that much of these properties extend also to tps with infinitely many variables, as those arising from the tropical relational model. The literature on tropical power series is often recent (e.g. [61]), and several results we prove in this section are, to our knowledge, new.

Notice that, as a set, $\mathbb{L}^X = [0, \infty]^X$, and with the usual $+$ and \cdot it is a $\mathbb{R}_{\geq 0}$ -semimodule, let us call it $\overline{\mathbb{R}}_{\geq 0}^X$. Together with the usual sup-norm $\|x\|_{\infty} := \sup_{a \in X} x_a$, it can be showed to be a Scott-complete *normed cone* (see [71] or the extended version). Suitable categories of cones have been recently investigated as models of probabilistic computation ([23, 31, 35]). The cone structure of $\overline{\mathbb{R}}_{\geq 0}^X$ also induces a partial order on it, its *cone partial-order*: $x \leq y$ iff $y = x + z$ for some (unique) $z \in \overline{\mathbb{R}}_{\geq 0}^X$. It actually coincides with the pointwise order on $\overline{\mathbb{R}}_{\geq 0}^X$ (and makes it a Scott-continuous dcpo). In this section we consider tps w.r.t. this structure.

Continuity of tps. Looking at Fig. 1, we see that φ , just like the polynomials φ_n , is non-decreasing and concave. This is indeed always the case:

► **Proposition 12.** *All tps are non-decreasing and concave, w.r.t. the pointwise order on $\overline{\mathbb{R}}_{\geq 0}^X$.*

The tps φ is continuous on $\mathbb{R}_{\geq 0}$ (w.r.t. the usual norm of real numbers). We can generalise this property, dropping the case of x having some 0 coordinate. But we have to be careful, because while in the finite dimensional \mathbb{R}^n , every real convex function is continuous because it is necessarily locally bounded from above (the sup-norm and the euclidean one are equivalent) [22, Proposition 4.7], in infinite dimensions the former condition is no longer true [22, Example 4.8]. However, [22, Proposition 4.4.(3)] shows that it is the only requirement to ask: if a real-valued convex function with domain a convex open subset of a locally convex topological \mathbb{R} -vector space (LCTVS) is, locally around any point, bounded from above by a finite non-zero constant, then it is continuous on all its domain.

► **Theorem 13.** *All tps $f : \overline{\mathbb{R}}_{\geq 0}^X \rightarrow \overline{\mathbb{R}}_{\geq 0}$ are continuous on $(0, \infty)^X$, w.r.t. to the norm $\|\cdot\|_{\infty}$.*

Proof. By Proposition 12, $-f$ is convex. Since $f \geq 0$ on all $\overline{\mathbb{R}}_{\geq 0}^X$, we have e.g. $-f \leq 1$ on $\overline{\mathbb{R}}_{\geq 0}^X$. Now $(0, \infty)^X \subseteq \mathbb{R}^X$ is open and convex, so [22, Proposition 4.4.(3)] entails the continuity of $-f$ on $(0, \infty)^X$, hence that of f on it. ◀

In analogy with [27, Proposition 17], we also have:

► **Theorem 14.** *All monotone (w.r.t. pointwise order) and $\|\cdot\|_{\infty}$ -continuous functions $f : (0, \infty)^X \rightarrow (0, \infty)$ are Scott-continuous. In particular, all tps $f : \overline{\mathbb{R}}_{\geq 0}^X \rightarrow \overline{\mathbb{R}}_{\geq 0}$ are Scott-continuous on $(0, \infty)^X$ w.r.t. the pointwise orders.*


Lipschitz-continuity of tps. Let us first look at what happens with those tps which are either *linear* or obtained via bounded exponentials. The result below is in analogy with what happens in the usual metric semantics of Fuzz, where linear functions are non-expansive and n -bounded functions are n -Lipschitz [68].

► **Proposition 15.**

1. *If a tps $f : \overline{\mathbb{R}}_{\geq 0}^X \rightarrow \overline{\mathbb{R}}_{\geq 0}^Y$ arises from a matrix $\hat{f} : X \times Y \rightarrow \overline{\mathbb{R}}_{\geq 0}$ (i.e. it is linear), then f is non-expansive (i.e. 1-Lipschitz).*
2. *If $f : \overline{\mathbb{R}}_{\geq 0}^X \rightarrow \overline{\mathbb{R}}_{\geq 0}^Y$ arises from a matrix $\hat{f} : !_n X \times Y \rightarrow \overline{\mathbb{R}}_{\geq 0}$, then f is n -Lipschitz-continuous.*

Proof sketch.

- 1). Using the fact that $f(x)_b = \inf_{a \in X} \{\hat{f}_{a,b} + x_a\}$, the problem reduces to: $|(\hat{f}_{a,b} - x_a) - (\hat{f}_{a,b} - y_a)| = |x_a - y_a| \leq \|x - y\|_{\infty}$.
- 2). Follows from 1. and the remark that, for all $x \in \mathbb{L}^X$, $\|!_n x - !_n y\|_{\infty} \leq n \cdot \|x - y\|_{\infty}$, where $!_n x$ is the restriction of $!x$ to $\mathcal{M}_{\leq n}(X)$. ◀



■ **Figure 4** Drawing of the proof of Theorem 18.

Observe that on the hom-sets $\mathbb{L}\text{Rel}_!(X, Y)$ there are two natural notions of distance: the metric $\|f - g\|_\infty$ arising from the norm and the one arising from the usual sup-metric $d_\infty(f, g) := \sup_{x \in \mathbb{L}^X} \|f^!(x) - g^!(x)\|_\infty$. In general one has $\|f - g\|_\infty \geq d_\infty(f, g)$, the equality holding when $f^!, g^!$ are linear (i.e. when they arise from morphisms of $\mathbb{L}\text{Rel}(X, Y)$).

When a tps is expressed as an inf of finitely many monomials we simply call it a tropical polynomial. For any such tropical polynomial $\varphi : \overline{\mathbb{R}}_{\geq 0}^X \rightarrow \overline{\mathbb{R}}_{\geq 0}^Y$, the associated matrix has shape $!_{\text{deg}(\varphi)}(X) \times Y \rightarrow \overline{\mathbb{R}}_{\geq 0}$ (as a monomial $\mu_i x + c_i$ yields a matrix entry on $!_{\# \mu_i} X \times Y$). Hence, using Proposition 15 2., we have:

► **Corollary 16.** *Any tropical polynomial $\varphi : \overline{\mathbb{R}}_{\geq 0}^X \rightarrow \overline{\mathbb{R}}_{\geq 0}$ is $\text{deg}(\varphi)$ -Lipschitz continuous.*

We now show that, if we consider the *full* exponential $!$, i.e. arbitrary tps, we can still prove that a local Lipschitz condition holds. In [22, Theorem 6.4] a locally Lipschitz property is obtained for locally convex topological vector spaces, under the hypothesis of continuity. [22, Proposition 6.5] shows that continuity is used in order to have a locally bounded condition, the crucial ingredient of the proof. Instead of showing how our case fits into such theorems, we prefer to state the following theorem, basically a particular case of [22, Theorem 6.9, Lemma 6.10]:

► **Theorem 17.** *All tps $f : \overline{\mathbb{R}}_{\geq 0}^X \rightarrow \overline{\mathbb{R}}_{\geq 0}$ are locally Lipschitz on $(0, \infty)^X$. Moreover, the Lipschitz constant of f on $\overline{B}_\delta(x)$ can be chosen to be $\frac{1}{\delta} \max_{\overline{B}_{3\delta}(x)} f$.*

The fundamental ingredient of the proof of this result is the following Theorem 18, from which Theorem 17 immediately follows, since $(0, \infty)^X$ is open and convex in $(\mathbb{R}^X, \|\cdot\|)$ and all tps are non-negative.

► **Theorem 18.** *Let $f : V \subseteq (\mathbb{R}^X, \|\cdot\|) \rightarrow (\mathbb{R}, |\cdot|)$, with V open and convex and $\|\cdot\|$ any norm. If f is concave and locally bounded, then f is locally Lipschitz. Moreover, the Lipschitz constant of f on $\overline{B}_\delta(x)$ can be chosen to be $\frac{1}{\delta} \max_{\overline{B}_{3\delta}(x)} |f|$.*

Proof. Call $\overline{B}_\delta(x) := B_1$, $\overline{B}_{3\delta}(x) := B_3$. It suffices to show that for all $x \in V$, there is $\delta > 0$ s.t. $B_3 \subseteq \text{interior}(V)$, $K := \max_{B_3} |f|$ exists and f is $\frac{K}{\delta}$ -Lipschitz on B_1 . A δ satisfying the first two conditions exists since V is open and because f is locally bounded and B_3 is compact. We show that the third condition holds for all such δ . For that, fix $y, z \in B_1$ and call $r := \frac{d(y, z)}{2\delta} \in [0, 1]$. We want to show that $|f(y) - f(z)| \leq \frac{K}{\delta} d(y, z) = 2Kr$. Wlog $y \neq z$, otherwise there is nothing to prove.

So $r \neq 0$ and we can consider $u := \frac{1+r}{r}z - \frac{1}{r}y$, $v := \frac{1}{r}y - \frac{r-1}{r}z$. We have $u, v \in \overline{B}_{2\delta}(z) =: B_2$. Indeed, $d(u, z) = \|u - z\| = \|\frac{z}{r} + z - \frac{y}{r} - z\| = \|\frac{z-y}{r}\| = 2\delta$ and similarly $d(v, z) = 2\delta$. Geometrically, those are actually the intersections between B_2 and the line generated by y and

z , see Figure 4. Now we have the convex combinations $z = \frac{1}{1+r}y + \frac{r}{1+r}u$ and $y = (1-r)z + rv$, so the concavity of f entails on one hand: $f(z) \geq \frac{1}{1+r}f(y) + \frac{r}{1+r}f(u) \geq \frac{f(y)}{1+r} - \frac{rK}{1+r}$, i.e. $f(y) - f(z) \leq r(K + f(z)) \leq 2rK$, and on the other hand: $f(y) \geq (1-r)f(z) + rf(v) \geq f(z) - r(f(z) + K)$, i.e. $f(z) - f(y) \leq r(f(z) + K) \leq 2rK$. In the previous inequalities we have used that $f(u), f(v) \geq -K$. This follows from the fact that $u, v \in B_2 \subseteq B_3$, as it can be immediately checked, thus $|f(u)|, |f(v)| \leq K$. Putting the final inequalities together, we have $|f(y) - f(z)| \leq 2rK$, i.e. the thesis. \blacktriangleleft

7 Lipschitz Meets Taylor

In this section we finally relate the metric and differential analysis of higher-order programs in the tropical relational model.

The key ingredient is the notion of Taylor expansion $\mathcal{T}(M)$ of a λ -term M . This is a set of terms of the differential λ -calculus defined inductively as: $\mathcal{T}(x) = \{x\}$, $\mathcal{T}(\lambda x.M) = \{\lambda x.t \mid t \in \mathcal{T}(M)\}$ and $\mathcal{T}(MN) = \{t \cdot \langle u_1, \dots, u_k \rangle \mid k \in \mathbb{N}, t \in \mathcal{T}(M), u_i \in \mathcal{T}(N)\}$, where $t \cdot \langle u_1, \dots, u_k \rangle$ is an abbreviation for $D^k[t, u_1, \dots, u_k]0$. Observe that in the terms appearing in $\mathcal{T}(M)$ all applications are *bounded*: they may use an exact number of copies of their input. Such terms are usually called *resource λ -terms* [56, 65]. One can easily check that for all terms $\Gamma \vdash_{\text{STLC}} M : A$ and $t \in \mathcal{T}(M)$, also $\Gamma \vdash_{\text{ST}\partial\text{LC}} t : A$ holds.

► **Example 19.** Considering the term $M = zxx$ from Example 8, all terms $t_{n,m} = z \langle x^n \rangle \langle x^m \rangle$, for $n, m \in \mathbb{N}$, are in $\mathcal{T}(M)$. Notice that the interpretation of $t_{n,m}$ yields a tropical polynomial $\llbracket t_{n,m} \rrbracket^!(x)(z) = y_{[n,m]} + (n+m)x$, rather than a tps. However, this is not a general fact: consider $y : (o \rightarrow o) \rightarrow (o \rightarrow o), x : (o \rightarrow o) \vdash t : (o \rightarrow o)$ with $t = y \cdot \langle y \cdot \langle x \rangle \rangle \in \mathcal{T}(y(yx))$. Then $\llbracket t \rrbracket^! : \mathbb{L}^{\mathbb{N} \times \mathbb{N}} \times \mathbb{L}^{\mathbb{N}} \rightarrow \mathbb{L}^{\mathbb{N}}$ is given by $\llbracket t \rrbracket^!(y, x)_i = \inf_{m,n \in \mathbb{N}} \{y_{[m],i} + y_{[n],m} + x_n\}$, which is not a polynomial. Yet, $\llbracket t \rrbracket^!$ is Lipschitz, more precisely, 1-Lipschitz in x and 2-Lipschitz in y . This is a general fact, as shown by Theorem 22 below.

We have already shown that the tropical differential makes $\mathbb{L}\text{Rel}_!$ a model of the differential λ -calculus. We now show that it also models the Taylor expansion (this needs not be true for *any* $\text{CC}\partial\lambda\text{C}$). First, it can be patiently checked that (see [56, Definition 4.22]):

► **Theorem 20.** *Morphisms in $(\mathbb{L}\text{Rel}_!, D)$ can be Taylor-expanded: for all $t \in \mathbb{L}\text{Rel}_!(Z, !X \multimap Y)$, $s \in \mathbb{L}\text{Rel}_!(Z, X)$ we have $\text{ev} \circ_! \langle t, s \rangle = \inf_{n \in \mathbb{N}} \{((\dots((\Lambda^- t) \star s) \star \dots) \star s) \circ_! \langle \text{id}, \infty \rangle\}$.*

The equation above is a tropical reformulation of the Taylor formula from the Introduction: $u \star s = (Du) \circ_! \langle \langle \infty, s \circ_! \pi_1 \rangle, \text{id} \rangle$ corresponds to the application of the derivative of u on s , and Λ^- is the uncurry operator. Hence the right-hand term corresponds to the inf of the n -th derivative of $\Lambda^- t$ applied to “ n copies” of s .

Second, since $\mathbb{L}\text{Rel}_!$ has countable sums (all countable infs converge), an immediate adaptation of the proof of [56, Theorem 4.23] shows:

► **Corollary 21.** $\llbracket \Gamma \vdash_{\text{STLC}} M : A \rrbracket = \inf_{t \in \mathcal{T}(M)} \llbracket \Gamma \vdash_{\text{ST}\partial\text{LC}} t : A \rrbracket$.

Using the results of the previous section, as well as the results above, we now deduce the following properties:

► **Theorem 22.** *Let \mathcal{S} be one of $\text{PCF}^{\mathbb{L}}, \text{STLC}, \text{bSTLC}, \text{ST}\partial\text{LC}$. Let $\Gamma \vdash_{\mathcal{S}} M : A$ and $a \in \llbracket A \rrbracket$.*

1. *For $\mathcal{S} = \text{bSTLC}$, $\llbracket \Gamma \vdash_{\mathcal{S}} M : A \rrbracket_a^!$ is a tropical polynomial, and thus Lipschitz;*
2. *For $\mathcal{S} = \text{ST}\partial\text{LC}$, if $t \in \mathcal{T}(M)$, then $\llbracket \Gamma \vdash_{\mathcal{S}} t : A \rrbracket_a^!$ is Lipschitz;*
3. *For $\mathcal{S} = \text{STLC}, \text{PCF}^{\mathbb{L}}$, then $\llbracket \Gamma \vdash_{\mathcal{S}} M : A \rrbracket_a^!$ is locally Lipschitz;*
4. *For $\mathcal{S} = \text{STLC}$, $\mathcal{T}(M)$ decomposes $\llbracket \Gamma \vdash_{\text{STLC}} M : A \rrbracket_a^!$ as an $\inf_{t \in \mathcal{T}(M)} \llbracket \Gamma \vdash_{\text{ST}\partial\text{LC}} t : A \rrbracket_a^!$ of Lipschitz functions.*

Proof.

- 1). From Proposition 15 2. and the remark that for any type of bSTLC, $\llbracket A \rrbracket$ is finite.
- 2). From Proposition 15 2. observing that a resource term $t(x)$ may use a variable x a fixed number n times, so that its matrix lies in $\mathbb{L}^{!n.X \times Y}$.
- 3). From Theorem 17.
- 4). It follows from Corollary 21 plus the fact that, for $(f_n)_{n \in \mathbb{N}} \subseteq \mathbb{L}^{!X \times Y}$, we have $(\inf_{n \in \mathbb{N}} f_n)^! = \inf_{n \in \mathbb{N}} f_n^!$. \blacktriangleleft

We conclude our discussion with an application of the Taylor expansion in $\mathbb{L}\text{Rel}$: as proved in the previous section, all tps are locally Lipschitz; now, Theorem 22 can be used to compute approximations of the Lipschitz constants of an actual higher-order program.

► **Corollary 23.** *Suppose $x : A \vdash_{\text{STLC}} M : B$ and $\vdash_{\text{STLC}} N : A$. Then for all $t \in \mathcal{T}(M)$ such that $\llbracket t \rrbracket^!(\llbracket N \rrbracket) \neq \infty$, and $\delta > 0$, the tps $\llbracket x : A \vdash_{\text{STLC}} M : B \rrbracket^!$ is $\frac{\llbracket t \rrbracket^!(\llbracket N \rrbracket) + 3\delta}{\delta}$ -Lipschitz over the open ball $B_\delta(\llbracket N \rrbracket)$.*

Proof. Thm. 17 yields the estimate $\max_{B_{3\delta}(\llbracket N \rrbracket)} \llbracket M \rrbracket^!$. As from Thm. 22 4. it follows that $\llbracket t \rrbracket^! \geq \llbracket M \rrbracket^!$, we deduce that $K = \max_{B_{3\delta}(\llbracket N \rrbracket)} \llbracket t \rrbracket^! \geq \max_{B_{3\delta}(\llbracket N \rrbracket)} \llbracket M \rrbracket^!$ is also a local Lipschitz constant for $\llbracket M \rrbracket^!$. Moreover, since $\llbracket t \rrbracket^!$ is concave and non-decreasing, the max of $\llbracket t \rrbracket^!$ is attained at the maximum point of $B_{3\delta}(\llbracket N \rrbracket)$, that is, $K = \llbracket t \rrbracket^!(x + 3\delta)$. Finally, from $\llbracket t \rrbracket^!(\llbracket N \rrbracket) < \infty$ and the continuity of $\llbracket t \rrbracket^!$ we deduce $K < \infty$. \blacktriangleleft

► **Example 24.** Consider again the term $M = zxx$ from Example 8. The (generalized) tps $\llbracket M \rrbracket^!(x)(y) = \inf_{n, n' \in \mathbb{N}} \{y_{[(n, n')]} + (n + n')x\}$ is not (globally) Lipschitz: for any $L > 0$, choose a natural number $N > L$, let $Y \in \mathbb{L}^{\mathcal{M}_{\text{fin}}(\mathbb{N} \times \mathbb{N})}$ be such that $Y_\mu < \infty$ only if $\mu = [(n, n')]$ with $n + n' \geq N$; then $|\llbracket M \rrbracket^!(x)(Y) - \llbracket M \rrbracket^!(x + \epsilon)(Y)| \geq N\epsilon > L\epsilon$. Now take the approximant $t = z\langle x^{N-1} \rangle x \in \mathcal{T}(M)$ (chosen so that $\llbracket t \rrbracket^!(x)(Y) < \infty$). Its interpretation is the monomial $\llbracket t \rrbracket^!(x)(Y) = Y_{[(N-1, 1)]} + Nx$. We can then compute a Lipschitz-constant for $\llbracket M \rrbracket^!$ around $\langle x, Y \rangle$ as $\frac{1}{\delta} \llbracket t \rrbracket^!(\langle x, Y \rangle + \delta) = 3N + 3 + \frac{Y_{[(N-1, 1)]} + Nx}{\delta}$.

8 Generalized Metric Spaces and \mathbb{L} -Modules

As we have seen, the morphisms of $\mathbb{L}\text{Rel}$ can be seen as continuous functions between the \mathbb{L} -modules \mathbb{L}^X , when the latter are taken with the metric induced by the ∞ -norm. This viewpoint gives a metric flavor to $\mathbb{L}\text{Rel}$, and allowed us to relate differential and metric structure. Yet, how far can this correspondence be pushed? In particular, is this correspondence restricted to \mathbb{L} -modules of the form \mathbb{L}^X (i.e. with a fixed base), or does it hold in some sense for arbitrary \mathbb{L} -modules? Is this correspondence restricted to the ∞ -norm metric, or does it hold for other metrics too?

8.1 \mathbb{L} -Modules and Cocomplete \mathbb{L} -Categories

An answer to the questions above comes from an elegant categorical correspondence between tropical linear algebra and the theory of *generalized metric spaces*, initiated by Lawvere's pioneering work [51], and at the heart of the emergent field of *monoidal topology* [47, 74].

On the one hand we have \mathbb{L} -modules: these are triples (M, \leq, \star) where (M, \leq) is a sup-lattice, and $\star : \mathbb{L} \times M \rightarrow M$ is a continuous (left-)action of \mathbb{L} on it, where continuous means that \star commutes with both joins in \mathbb{L} and in M . A \mathbb{L} -module homomorphism is a map $f : M \rightarrow N$ commuting with both joins and the \mathbb{L} -action. We let $\mathbb{L}\text{Mod}$ indicate the category of \mathbb{L} -modules and their homomorphisms.

On the other hand we have Lawvere's *generalized metric spaces* [47, 51, 74]: Lawvere was the first to observe that a metric space can be described as a \mathbb{L} -enriched category. Indeed, spelling out the definition, a \mathbb{L} -enriched category (in short, a \mathbb{L} -category) is given by a set X together with a "hom-set" $X(-, -) : X \times X \rightarrow \mathbb{L}$ satisfying $0 \geq X(x, x)$ and $X(y, z) + X(x, y) \geq X(x, z)$. This structure clearly generalizes the usual definition of metric spaces, which are indeed precisely the \mathbb{L} -categories which are *skeletal* (i.e. $X(x, y) = 0$ implies $x = y$) and *symmetric* (i.e. $X(x, y) = X^{\text{op}}(x, y)$, where $X^{\text{op}}(x, y) = X(y, x)$). A basic example of \mathbb{L} -category is \mathbb{L} itself, with the distance $\mathbb{L}(x, y) = y \dot{-} x$.

Moreover, a \mathbb{L} -enriched *functor* between \mathbb{L} -categories is nothing but a non-expansive map $f : X \rightarrow Y$, since functoriality reads as $Y(f(x), f(y)) \leq X(x, y)$. Functors of shape $\Phi : Y^{\text{op}} \times X \rightarrow \mathbb{L}$ are called *distributors* and noted $\Phi : X \leftrightarrow Y$. Two distributors $\Phi : Y \leftrightarrow X$ and $\Psi : Z \leftrightarrow Y$ can be composed just like ordinary matrices in $\mathbb{L}\text{Rel}$: $\Psi \diamond \Phi : Z \leftrightarrow X$ is given by $(\Psi \diamond \Phi)_{x,z} = \inf_{y \in Y} \{\Psi(y, z) + \Phi(x, y)\}$.

Lawvere also observed that the usual notion of Cauchy-completeness can be formulated, in this framework, as the existence of suitable colimits [51]. Let us recall the notion of weighted colimit in this context:

► **Definition 25** (weighted colimits). *Let X, Y, Z be \mathbb{L} -categories, $\Phi : Z \leftrightarrow Y$ be a distributor and $f : Y \rightarrow X$ be a functor. A functor $g : Z \rightarrow X$ is the Φ -weighted colimit of f over X , noted $\text{colim}(\Phi, f)$, if for all $z \in Z$ and $x \in X$*

$$X(g(z), x) = \sup_{y \in Y} \{X(f(y), x) \dot{-} \Phi(y, z)\}$$

A functor $f : X \rightarrow Y$ is said *cocontinuous* if it commutes with all existing weighted colimits in X , i.e. $f(\text{colim}(\Phi, g)) = \text{colim}(\Phi, f \circ g)$. A \mathbb{L} -enriched category. A \mathbb{L} -category X is said *cocomplete* if all weighted colimits over X exist. We let $\mathbb{L}\text{CCat}$ indicate the category of cocomplete \mathbb{L} -categories and cocontinuous \mathbb{L} -enriched functors as morphisms.

Observe that cocompleteness for a symmetric \mathbb{L} -category X implies the usual Cauchy completeness. Indeed, a Cauchy sequence $(x_n)_{n \in \mathbb{N}}$ in X yields two adjoint distributors $\phi^* : 1 \leftrightarrow X$ and $\phi_* : X \leftrightarrow 1$, where $\phi^*(x') = \lim_{n \rightarrow \infty} X(x', x_n)$ and $\phi_*(x') = \lim_{n \rightarrow \infty} X(x_n, x')$. Hence, $\text{colim}(\phi^*, 1_X) : 1 \rightarrow X$ must be a point x satisfying $0 = X(x, x) = \sup_{y \in X} \lim_{n \rightarrow \infty} (X(y, x) \dot{-} X(y, x_n))$, which implies $\lim_{n \rightarrow \infty} X(x_n, x) = 0$.

It turns out that the notions of \mathbb{L} -module and cocomplete \mathbb{L} -category are indeed equivalent. More precisely, the categories $\mathbb{L}\text{Mod}$ and $\mathbb{L}\text{CCat}$ are isomorphic [73]. First, any \mathbb{L} -module (M, \leq, \star) can be endowed with the structure of a \mathbb{L} -category by letting $M(x, y) = \inf \{\epsilon \mid \epsilon \star x \geq y\}$. Moreover, a homomorphism of \mathbb{L} -modules induces a cocontinuous functor of the associated \mathbb{L} -categories. Conversely, in cocomplete \mathbb{L} -categories it is possible to introduce a continuous \mathbb{L} -action via suitable weighted colimits called *tensors*:

► **Definition 26** (tensors, cf. [74]). *Let X be a \mathbb{L} -category, $x \in X$ and $\epsilon \in \mathbb{L}$. The tensor of x and ϵ , if it exists, is the colimit $\epsilon \otimes x := \text{colim}([\epsilon], \Delta x)$, where $[\epsilon] : 1 \rightarrow 1$ is the constantly ϵ distributor and $\Delta x : 1 \rightarrow X$ is the constant functor.*

A cocomplete \mathbb{L} -category can thus be endowed with a \mathbb{L} -module structure with order given by $x \leq_X y$ iff $X(y, x) = 0$, and action given by tensors $\epsilon \otimes x$. Moreover, a cocontinuous functor between cocomplete \mathbb{L} -categories is the same as a homomorphism of the associated \mathbb{L} -modules.

8.2 Exponential and Differential Structure of $\mathbb{L}\text{Mod} \simeq \mathbb{L}\text{CCat}$

We now show that the correspondence between \mathbb{L} -modules and cocomplete \mathbb{L} -categories lifts to a model of the differential λ -calculus, generalizing the co-Kleisli category $\mathbb{L}\text{Rel}$.

In order to define a Lafont exponential $!$ over $\mathbb{L}\text{Mod}$, we exploit a well-known recipe from [50, 60]. The first step is to define a symmetric algebra $\text{Sym}_n(M)$ as the equalizer of all permutative actions on n -tensors $M \otimes \cdots \otimes M$. Notice that each element of $!_n M$ can be described as a join of “multisets” $[x_1, \dots, x_n]$, where the latter is the equivalence class of the tensors $x_1 \otimes \cdots \otimes x_n \in M^{\otimes n}$ under the action of permutations $\sigma \in \mathfrak{S}_n$. The \mathbb{L} -module $!_n M$ is a cocomplete \mathbb{L} -category with distance function defined on basic “multisets” as follows:

$$(!_n M)(\alpha, \beta) = \sup_{\sigma \in \mathfrak{S}_n} \inf_{\tau \in \mathfrak{S}_n} \sum_{i=1}^n X(x_{\sigma(i)}, y_{\tau(i)}) \quad (1)$$

where $\alpha = [x_1, \dots, x_n]$ and $\beta = [y_1, \dots, y_n]$, and extended to arbitrary elements $\alpha = \bigvee_i \alpha_i$ and $\beta = \bigvee_j \beta_j$ by $(!_n M)(\alpha, \beta) = \sup_i \inf_j (!_n M)(\alpha_i, \beta_j)$.

Next, we define $!M$ as the infinite biproduct $\prod_n !_n M$, yielding the cofree commutative comonoid over M (cf. [60, Proposition 1]). Using the fact that biproducts commute with tensors in $\mathbb{L}\text{Rel}$, by standard results [60], we obtain that the coKleisli category $\mathbb{L}\text{Mod}_!$ is a CCC. Moreover, the constructions for $\mathbb{L}\text{Mod}$ generalize those of $\mathbb{L}\text{Rel}$, in the sense that $!(\mathbb{L}^X) \simeq \mathbb{L}^{\mathcal{M}^{\text{fin}}(X)}$ and that $\mathbb{L}\text{Mod}_!(\mathbb{L}^X, \mathbb{L}^Y) \simeq \mathbb{L}\text{Rel}_!(X, Y)$.

Finally, since the coKleisli category of a Lafont category with biproducts is always a CC ∂ C [53, Theorem 21], we can endow $\mathbb{L}\text{Mod}_!$ with a differential operator E , generalizing $D^!$, given by

$$Ef(\alpha) = \bigvee \left\{ f(\beta \cup [x]) \mid \iota_n(\beta) \otimes \iota_1(x) \leq S(\alpha) \right\} \quad (2)$$

where $\iota_k : M_k \rightarrow \prod_{i \in I} M_i$ is the injection morphism given by $\iota_k(x)(k) = x$ and $\iota_k(x)(i \neq k) = \infty$, and $S : !(M \times N) \rightarrow !M \otimes !N$ is the Seelye isomorphism [60], and conclude that:

► **Theorem 27.** $(\mathbb{L}\text{Mod}_!/\mathbb{L}\text{CCat}_!, E)$ is a CC ∂ C.

9 Related Work

The applications of tropical mathematics in computer science abound, e.g. in automata theory [48, 72], machine learning [57, 64, 77], optimization [4, 5], and convex analysis [54].

As we said, the relational semantics over the tropical semiring was quickly explored in [50], to provide a “best case” resource analysis of a PCF-like language with non-deterministic choice. The connections between differential λ -calculus (and differential linear logic), relational semantics, and non-idempotent intersection types are very well-studied (see [28], and more recently, [59] for a more abstract perspective, and [40, 63] for a 2-categorical, or proof-relevant, extension). *Probabilistic coherent spaces* [33], a variant of the relational semantics, provide an interpretation of higher-order probabilistic programs as analytic functions. In [32] it was observed that such functions satisfy a local Lipschitz condition somehow reminiscent of our examples in Section 3.

The study of linear or bounded type systems for sensitivity analysis was initiated in [44] and later developed [26, 68, 70]. Related approaches, although not based on metrics, are provided by *differential logical relations* [24] and *change action* models [7]. More generally, the literature on program metrics in denotational semantics is vast. Since at least [75] metric spaces, also in Lawvere’s generalized sense [51], have been exploited as an alternative framework to standard, domain-theoretic, denotational semantics; notably models of STLC and PCF based on *ultra*-metrics and *partial* metrics have been proposed [25, 37, 66].

Motivated by connections with computer science and fuzzy set-theory, the abstract study of generalized metric spaces in the framework of *quantale*- or even *quantaloid*-enriched categories has led to a significant literature in recent years (e.g. [47, 74]), and connections with tropical mathematics also have been explored e.g. in [38, 76]. Moreover, applications of quantale-modules to both logic and computer science have also been studied [1, 69].

Finally, connections between program metrics and the differential λ -calculus have been already suggested in [66]; moreover, *cartesian difference categories* [6] have been proposed as a way to relate derivatives in differential categories with those found in change action models.

10 Conclusion and Future Work

The main goals of this paper are two. Firstly, to demonstrate the existence of a conceptual bridge between two well-studied quantitative approaches to higher-order programs, and to highlight the possibility of transferring results and techniques from one approach to the other. Secondly, to suggest that tropical mathematics, a field which has been largely and successfully applied in computer science, could be used for the quantitative analysis of functional programming languages. While the first goal was here developed in detail, and at different levels of abstraction, for the second goal we only sketched a few interesting directions, and we leave their development to a second paper of this series.

While the main ideas of this article only use basic concepts from the toolbox of tropical mathematics, an exciting direction is that of looking at potential applications of more advanced tools from tropical algebraic and differential geometry (e.g. Newton polytopes, tropical varieties, tropical differential equations). Another interesting question is how much of our results on tps and their tropical Taylor expansion can be extended to the abstract setting of generalized metric spaces and continuous functors.

References

- 1 Samson Abramsky and Steven Vickers. Quantales, observational logic and process semantics. *Mathematical Structures in Computer Science*, 3(2):161–227, 1993. doi:10.1017/S0960129500000189.
- 2 Beniamino Accattoli, Ugo Dal Lago, and Gabriele Vanoni. The (in)efficiency of interaction. In *Proceedings POPL 2021*, volume 5, New York, NY, USA, 2021. Association for Computing Machinery.
- 3 Beniamino Accattoli, Ugo Dal Lago, and Gabriele Vanoni. Multi types and reasonable space. *Proceedings ICFP 2022*, 6, 2022.
- 4 Marianne Akian, Stéphane Gaubert, and Alexander Guterman. Tropical polyhedra are equivalent to mean payoff games. *International Journal of Algebra and Computation*, 22(01):1250001, 2023/01/16 2012. doi:10.1142/S0218196711006674.
- 5 Marianne Akian, Stéphane Gaubert, Viorel Nițică, and Ivan Singer. Best approximation in max-plus semimodules. *Linear Algebra and its Applications*, 435(12):3261–3296, 2011. doi:10.1016/j.laa.2011.06.009.
- 6 Mario Alvarez-Picallo and Jean-Simon Pacaud Lemay. Cartesian difference categories. In Jean Goubault-Larrecq and Barbara König, editors, *Proceedings FoSSaCS 2020*, pages 57–76, Cham, 2020. Springer International Publishing.
- 7 Mario Alvarez-Picallo and C.-H. Luke Ong. Change actions: Models of generalised differentiation. In Mikołaj Bojańczyk and Alex Simpson, editors, *Proceedings FoSSaCS 2019*, pages 45–61, Cham, 2019. Springer International Publishing.
- 8 Mário S. Alvim, Miguel E. Andrés, Konstantinos Chatzikokolakis, Pierpaolo Degano, and Catuscia Palamidessi. Differential privacy: On the trade-off between utility and information leakage. In *Proceedings FAST 2011*, FAST–11, pages 39–54, Berlin, Heidelberg, 2011. Springer-Verlag. doi:10.1007/978-3-642-29420-4_3.

- 9 Melissa Antonelli, Ugo Dal Lago, and Paolo Pistone. Curry and Howard Meet Borel. In *Proceedings LICS 2022*, pages 1–13. IEEE Computer Society, 2022.
- 10 Arthur Azevedo de Amorim, Marco Gaboardi, Justin Hsu, Shin-ya Katsumata, and Ikram Cherigui. A semantic account of metric preservation. In *Proceedings POPL 2017*, pages 545–556, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3009837.3009890.
- 11 Marco Azevedo de Amorim, Gaboardi, Arthur, Justin Hsu, and Shin-ya Katsumata. Probabilistic relational reasoning via metrics. In *Proceedings LICS 2019*. IEEE Computer Society, 2019.
- 12 Paolo Baldan, Filippo Bonchi, Henning Kerstan, and Barbara König. Coalgebraic behavioral metrics. *Logical Methods in Computer Science*, 14(3), 2018. doi:10.23638/LMCS-14(3:20)2018.
- 13 Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. Probabilistic relational reasoning for differential privacy. In *Proceedings POPL 2012*. ACM Press, 2012. doi:10.1145/2103656.2103670.
- 14 Leonard E. Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state markov chains. *The Annals of Mathematical Statistics*, 37(6):1554–1563, 2023/01/17/1966. URL: <http://www.jstor.org/stable/2238772>.
- 15 Richard Blute, Thomas Ehrhard, and Christine Tasson. A convenient differential category. *Cahiers de Topologie et Géométrie Différentielle Catégorique*, LIII:211–232, 2012. arXiv:1006.3140.
- 16 Richard F. Blute, Robin Cockett, J.S.P. Lemay, and R.A.G. Seely. Differential categories revisited. *Applied Categorical Structures*, 28:171–235, 2020.
- 17 Richard F. Blute, Robin Cockett, and R.A.G. Seely. Cartesian Differential Categories. *Theory and Applications of Categories*, 22(23):622–672, 2009.
- 18 Gérard Boudol. The lambda-calculus with multiplicities. In Eike Best, editor, *Proceedings CONCUR'93*, pages 1–6, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg.
- 19 Flavien Breuvar and Ugo Dal Lago. On intersection types and probabilistic lambda calculi. In *Proceedings PPDP 2018*, PPDP '18, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3236950.3236968.
- 20 Antonio Bucciarelli, Thomas Ehrhard, and Giulio Manzonetto. Categorical models for simply typed resource calculi. In *Proceedings MFPS 2010*, volume 265 of *Electronic Notes in Theoretical Computer Science*, pages 213–230, 2010. doi:10.1016/j.entcs.2010.08.013.
- 21 Maria Manuel Clementino and Dirk Hofmann. Exponentiation in V-categories. *Topology and its Applications*, 153(16):3113–3128, 2006. doi:10.1016/j.topol.2005.01.038.
- 22 Stefan Cobzaş. Lipschitz properties of convex functions. *Advances in Operator Theory*, 2(1):21–49, 2017. doi:10.22034/aot.1610.1022.
- 23 Raphaëlle Crubillé. Probabilistic stable functions on discrete cones are power series. In *Proceedings LICS 2018*, pages 275–284. ACM, 2018. doi:10.1145/3209108.3209198.
- 24 Ugo Dal Lago, Francesco Gavazzo, and Akira Yoshimizu. Differential logical relations, part I: the simply-typed case. In *Proceedings ICALP 2019*, pages 111:1–111:14, 2019. doi:10.4230/LIPIcs.ICALP.2019.111.
- 25 Ugo Dal Lago, Furio Honsell, Marina Lenisa, and Paolo Pistone. On quantitative algebraic higher-order theories. In *Proceedings FSCD 2022*, volume 228 of *LIPIcs*, pages 4:1–4:18, 2022.
- 26 Ugo Dal Lago and Ulrich Schöpp. Computation by interaction for space-bounded functional programming. *Information and Computation*, 248:150–194, June 2016. doi:10.1016/j.ic.2015.04.006.
- 27 Vincent Danos and Thomas Ehrhard. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Information and Computation*, 209(6):966–991, 2011. doi:10.1016/j.ic.2011.02.001.
- 28 Daniel de Carvalho. Execution time of λ -terms via denotational semantics and intersection types. *Mathematical Structures in Computer Science*, 28(7):1169–1203, 2018.

- 29 Thomas Ehrhard. Finiteness spaces. *Mathematical Structures in Computer Science*, 15(4):615–646, 2005. URL: <https://www.cambridge.org/core/article/finiteness-spaces/E5E9CE1FA4050A56EF25CFB6F6A5754F>.
- 30 Thomas Ehrhard. An introduction to differential linear logic: proof-nets, models and antiderivatives. *Mathematical Structures in Computer Science*, pages 1–66, February 2017. doi:10.1017/S0960129516000372.
- 31 Thomas Ehrhard. Cones as a model of intuitionistic linear logic. In *Proceedings LICS 2020*, pages 370–383. IEEE Computer Society, 2020. doi:10.1145/3373718.3394758.
- 32 Thomas Ehrhard. Differentials and distances in probabilistic coherence spaces. *Logical Methods in Computer Science*, 18(3):2:1–2:33, 2022.
- 33 Thomas Ehrhard, Michele Pagani, and Christine Tasson. The computational meaning of probabilistic coherence spaces. In *Proceedings LICS 2011*, pages 87–96. IEEE Computer Society, 2011. doi:10.1109/LICS.2011.29.
- 34 Thomas Ehrhard, Michele Pagani, and Christine Tasson. Full Abstraction for Probabilistic PCF. *Journal of the ACM*, 65(4), 2018.
- 35 Thomas Ehrhard, Michele Pagani, and Christine Tasson. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. In *Proceedings POPL 2018*, volume 2, pages 59:1–59:28, 2018. doi:10.1145/3158147.
- 36 Thomas Ehrhard and Laurent Regnier. The differential lambda-calculus. *Theoretical Computer Science*, 309(1):1–41, December 2003. doi:10.1016/S0304-3975(03)00392-X.
- 37 Martín Hötzen Escardó. A metric model of PCF, 1999. Unpublished note presented at the Workshop on Realizability Semantics and Applications, June 1999. Available at the author’s webpage.
- 38 Soichiro Fuji. Enriched categories and tropical mathematics, 2019. arXiv:1909.07620.
- 39 Marco Gaboardi, Andreas Haeberlen, Justin Hsu, Arjun Narayan, and Benjamin C. Pierce. Linear dependent types for differential privacy. *SIGPLAN Not.*, 48(1):357–370, January 2013. doi:10.1145/2480359.2429113.
- 40 Zeinab Galal. A bicategorical model for finite nondeterminism. In *Proceedings FSCD 2021*, volume 195 of *LIPICs*, pages 10:1–10:17, 2021.
- 41 D. Gebler and S. Tini. SOS specifications for uniformly continuous operators. *Journal of Computer and System Science*, 92:113–151, 2018.
- 42 Guillaume Geoffroy and Paolo Pistone. A partial metric semantics of higher-order types and approximate program transformations. In *Proceedings CSL 2021*, volume 183 of *LIPICs*, pages 35:1–35:18, 2021.
- 43 Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–101, 1987. doi:10.1016/0304-3975(87)90045-4.
- 44 Jean-Yves Girard, Andre Scedrov, and Philip J. Scott. Bounded linear logic: A modular approach to polynomial time computability. *Theoretical Computer Science*, 97:1–66, 1992. Extended abstract in *Feasible Mathematics*, S. R. Buss and P. J. Scott editors, Proceedings of the MCI Workshop, Ithaca, NY, June 1989, Birkhauser, Boston, pp. 195–209.
- 45 Dima Grigoriev. Tropical differential equations. *Advances in Applied Mathematics*, 82:120–128, 2017. doi:10.1016/j.aam.2016.08.002.
- 46 Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. A convenient category for higher-order probability theory. In *Proceedings LICS 2017*. IEEE Computer Society, 2017.
- 47 Dirk Hofmann, Gavin J Seal, and W Tholen. *Monoïdal Topology: a Categorical Approach to Order, Metric and Topology*. Cambridge University Press, New York, 2014.
- 48 C. Kahlert and L.O. Chua. The complete canonical piecewise-linear representation. I. the geometry of the domain space. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 39(3):222–236, 1992. doi:10.1109/81.128016.
- 49 Shin-ya Katsumata. A double category-theoretic analysis of graded linear exponential comonads. In *Proceedings FoSSaCS 2018*, pages 110–127. Springer International Publishing, 2018.

- 50 Jim Laird, Giulio Manzonetto, Guy McCusker, and Michele Pagani. Weighted relational models of typed lambda-calculi. In *Proceedings LICS 2013*, pages 301–310. IEEE Computer Society, 2013. doi:10.1109/LICS.2013.36.
- 51 F. William Lawvere. Metric spaces, generalized logic, and closed categories. *Rendiconti del Seminario Matematico e Fisico di Milano*, 43(1):135–166, December 1973. doi:10.1007/BF02924844.
- 52 Jean-Simon Pacaud Lemay. Convenient antiderivatives for differential linear categories. *Mathematical Structures in Computer Science*, 30(5):545–569, 2020.
- 53 Jean-Simon Pacaud Lemay. Coderelections for Free Exponential Modalities. In *Proceedings CALCO 2021*, volume 211 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:21, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CALCO.2021.19.
- 54 Yves Lucet. What shape is your conjugate? a survey of computational convex analysis and its applications. *SIAM Journal on Optimization*, 20(1):216–250, 2009. doi:10.1137/080719613.
- 55 Diane Maclagan and Bernd Sturmfels. *Introduction to tropical geometry*, volume 161 of *Graduate Studies in Mathematics*. American Mathematical Society, 2015.
- 56 Giulio Manzonetto. What is a categorical model of the differential and the resource λ -calculi? *Mathematical Structures in Computer Science*, 22(3):451–520, 2012. doi:10.1017/S0960129511000594.
- 57 Petros Maragos, Vasileios Charisopoulos, and Emmanouil Theodosios. Tropical geometry and machine learning. *Proceedings of the IEEE*, 109(5):728–755, 2021. doi:10.1109/JPROC.2021.3065238.
- 58 Radu Mardare, Prakash Panangaden, and Gordon Plotkin. Quantitative algebraic reasoning. In *Proceedings LICS 2016*. IEEE Computer Society, 2016.
- 59 Damiano Mazza, Luc Pellissier, and Pierre Vial. Polyadic approximations, fibrations and intersection types. In *Proceedings POPL 2018*. ACM, 2018.
- 60 Paul-André Melliès, Nicolas Tabareau, and Christine Tasson. An explicit formula for the free exponential modality of linear logic. *Mathematical Structures in Computer Science*, 28(7):1253–1286, 2018. doi:10.1017/S0960129516000426.
- 61 Gian Maria Negri Porzio, Vanni Noferini, and Leonardo Robol. Tropical Laurent series, their tropical roots, and localization results for the eigenvalues of nonlinear matrix functions, 2021. arXiv:2107.07982.
- 62 Vanni Noferini, Meisam Sharify, and Françoise Tisseur. Tropical roots as approximations to eigenvalues of matrix polynomials. *SIAM J. Matrix Anal. Appl.*, 36(1):138–157, January 2015. doi:10.1137/14096637X.
- 63 Federico Olimpieri. Intersection type distributors. In *Proceedings LICS 2021*. IEEE Computer Society, 2021. doi:10.1109/LICS52264.2021.9470617.
- 64 Lior Pachter and Bernd Sturmfels. Tropical geometry of statistical models. *Proceedings of the National Academy of Sciences*, 101(46):16132–16137, 2023/01/16 2004. doi:10.1073/pnas.0406010101.
- 65 Michele Pagani and Paolo Tranquilli. Parallel reduction in resource λ -calculus. In *Proceedings APLAS 2009*, pages 226–242, 2009.
- 66 Paolo Pistone. On generalized metric spaces for the simply typed λ -calculus. In *Proceedings LICS 2021*, pages 1–14. IEEE Computer Society, 2021.
- 67 Gordon Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5(3):223–255, 1977.
- 68 Jason Reed and Benjamin C. Pierce. Distance makes the types grow stronger. *Proceedings ICFP 2010*, pages 157–168, 2010.
- 69 Ciro Russo. *Quantale Modules, with Applications to Logic and Image Processing*. PhD thesis, Università degli Studi di Salerno, available at arXiv:0909.4493, 2007.

- 70 Ulrich Schöpp. Stratified Bounded Affine Logic for Logarithmic Space. In *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 411–420, July 2007. doi:10.1109/LICS.2007.45.
- 71 Peter Selinger. Towards a semantics for higher-order quantum computation. In *Proceedings QPL 2004*, TUCS General Publication No 33, pages 127–143, 2004.
- 72 Imre Simon. On semigroups of matrices over the tropical semiring. *Informatique Théorique et Applications*, 28:277–294, 1994.
- 73 Isar Stubbe. Categorical structures enriched in a quantaloid: Tensored and cotensored categories. *Theory and Applications of Categories*, 16(14):283–306, 2006.
- 74 Isar Stubbe. An introduction to quantaloid-enriched categories. *Fuzzy Sets and Systems*, 256:95–116, 2014. Special Issue on Enriched Category Theory and Related Topics (Selected papers from the 33rd Linz Seminar on Fuzzy Set Theory, 2012). doi:10.1016/j.fss.2013.08.009.
- 75 Franck van Breugel. An introduction to metric semantics: operational and denotational models for programming and specification languages. *Theoretical Computer Science*, 258(1):1–98, 2001. doi:10.1016/S0304-3975(00)00403-5.
- 76 Simon Willerton. Tight spans, Isbell completions and semi-tropical modules. *Theory and Applications of Categories*, 28(22):696–732, 2013.
- 77 Liwen Zhang, Gregory Naitzat, and Lek-Heng Lim. Tropical geometry of deep neural networks. In *Proceedings ICML 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 5819–5827. PMLR, 2018. URL: <http://proceedings.mlr.press/v80/zhang18i.html>.