



HAL
open science

Fast and General Automatic Differentiation for Finite-State Methods

Lucas Ondel Yang, Tina Raissi, Martin Kocour, Pablo Riera, Caio Corro

► **To cite this version:**

Lucas Ondel Yang, Tina Raissi, Martin Kocour, Pablo Riera, Caio Corro. Fast and General Automatic Differentiation for Finite-State Methods. 2026. <hal-05500843>

HAL Id: hal-05500843

<https://hal.science/hal-05500843v1>

Preprint submitted on 11 Feb 2026

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 CC BY 4.0 - Attribution - International License

Fast and General Automatic Differentiation for Finite-State Methods

Lucas Ondel Yang¹ Tina Raissi² Martin Kocour³ Pablo Riera⁴ Caio Corro⁵

Abstract

We propose a new method, that we coined the “morphism-trick”, to integrate custom implementations of vector-Jacobian products in automatic differentiation softwares, applicable to a wide range of semiring-based computations. Our approach leads to efficient and semiring-agnostic implementations of the backward pass of dynamic programming algorithms. For the particular case of finite-state methods, we introduce an algorithm that computes and differentiates the \oplus -sum of all paths’ weight of a finite-state automaton. Results show that, with minimal effort from the user, our novel library allows computing the gradient of a function w.r.t. to the weights of a finite state automaton orders of magnitude faster than state-of-the-art automatic differentiation systems. Implementations are made available via an open-source library distributed under a permissive license.

1. Introduction

Most machine learning algorithms rely on the optimization of a cost function. When the data exhibits some kind of organisation, it is often desirable to design functions capturing the underlying structure of the data or enforcing a priori adequate constraints. The problem or structured prediction (or structured learning) is the problem of designing and optimizing such functions.

A broad family of structured prediction methods (Goodman, 1999; Mohri, 2002; Rush, 2020; Ondel et al., 2022) is concerned with optimizing functions that are decomposable into a finite number of terms via two abstract operators, a “multiplication” (\otimes), and an “addition” (\oplus) forming a semiring. Such functions are conveniently represented as a directed

acyclic graph, facilitating the exploration of the structure, and they are efficiently evaluated via Dynamic Programming (DP) (Bellman, 1952). From a formal perspective, this family of finitely semiring-decomposable functions can be expressed in a unified way via the framework weighted finite state automata (Mohri et al., 2002). This formalism has enabled the development of general, efficient, and scalable softwares (Mohri et al., 2000; Riley et al., 2009) for DP algorithms on graphs for structured prediction problems.

However, DP algorithms, and semiring-abstracted computations in general, raises major challenges from the perspective of automatic differentiation. Because the semiring is user-defined and unknown beforehand, room for optimization is strongly limited. As a result, while it is technically possible to automatically differentiate semiring-based computations with tools such as (Innes, 2018), the computational cost is prohibitive. This is a major obstacle in the integration of large scale structured prediction methods in end-to-end pipelines which rely on gradient descent for optimisation method.

This work proposes the following contributions.

1. We introduce the “morphism-trick”: if the semiring’s \oplus -addition is isomorphic to the addition of real numbers, the computation graph to backpropagate the derivatives can be “flattened” and a vector-Jacobian product algorithm can be implemented without knowing the semiring beforehand.
2. We show that the morphism-trick can be applied to a broad family of semirings, including multi-valued semiring and the tropical semiring.
3. Using the morphism-trick and the WFSM formalism, we propose a general and efficient DP algorithm and its corresponding vector-Jacobian product. An implementation of this algorithm is provided in an open-source software.

2. Related Works

To alleviate the computational load of direct AD of semiring-based linear algebra operations, they implement optimized functions and the corresponding vector-Jacobian product for a group of carefully chosen semirings. This necessary

¹Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire du Numérique, Orsay, France ²RWTH Aachen University, Aachen, Germany ³Brno University of Technology, Faculty of Information Technology, Brno, Czech Republic ⁴Instituto de Investigación en Ciencias de la Computación (ICC), CONICET-UBA, Argentina ⁵INSA Rennes, IRISA, CNRS, Université de Rennes, Rennes, France. Correspondence to: Lucas Ondel Yang <lucas.ondel@cnsr.fr>.

optimization limits the flexibility of their approach: a new structured prediction algorithm requiring a different semiring involves rewriting and optimizing key operations and their corresponding vector-Jacobian product. Finally, since they use exclusively dense operations their approach is not suited for finite state automata which are inherently sparse objects.

Differentiation of finite state automata has been explored by Hannun et al. (2020) and Hannun (2021) where they focus exclusively on the log-semiring and, to a lesser extent, on the tropical semiring. In a similar vein, the k2-fsa software (Povey & Kuan, 2020) proposes a practical implementation of vector-Jacobian products for finite state automata operations limited to the log-semiring. Balakrishnan & Deshmukh (2024) proposes to directly use AD tools to differentiate through finite state automata computation over arbitrary differentiable semirings. However, they illustrate their approach on small problems (a 5-states and a 3-states automaton) far from the scale seen in, say, natural language processing.

Another line of work focuses on latent structures in neural networks (Niculae et al., 2025), i.e. structures that appear as intermediate computation inside a neural network (Kim et al., 2017). To this end, several methods have been proposed, including differentiating through regularized dynamic programming algorithms (Mensch & Blondel, 2018; Corro & Titov, 2019), efficient explicit marginalization via sparse distributions (Niculae et al., 2018; Correia et al., 2020), general frameworks based on perturbed optimization (Berthet et al., 2020; Stewart et al., 2023) and implicit differentiation (Blondel et al., 2022), among others.

Finally, a concept of derivatives (Brzozowski, 1964) and partial derivatives (Antimirov, 1996) for regular expressions and finite state automata has already been defined and extended to weighted finite state automata (Lombardy & Sakarovitch, 2005). Despite a resemblance with the present work, this is merely homonymy: we aim at computing the rate of change of a function with respect to some parameters whereas their concept of derivatives aims at providing a canonical representation of (weighted) regular expressions.

3. Preliminaries

In this section, we introduce relevant terms, concepts, and notations used throughout this work.

A **monoid** (Perrin, 1997) (M, \odot, ϵ) is an algebraic structure over a non-empty set M where $\odot : M \times M \rightarrow M$ is an associative binary operation, that is for all $x, y, z \in M$:

$$(x \odot y) \odot z = x \odot (y \odot z),$$

and $\epsilon \in M$ is a neutral element, that is $\forall x \in M$:

$$x \odot \epsilon = \epsilon \odot x = x.$$

A monoid M is said to be *commutative* if \odot is a commutative operator, i.e., $\forall x, y \in M : x \odot y = y \odot x$.

A **monoid morphism** (or simply morphism) is a mapping μ from a monoid (M, \odot, ϵ) to another monoid (M', \bullet, ϵ') such that for all $x, y \in M$:

$$\begin{aligned} \mu(x \odot y) &= \mu(x) \bullet \mu(y) \\ \mu(\epsilon) &= \epsilon'. \end{aligned}$$

A monoid morphism μ is an *isomorphism* if μ is bijective. The monoids (M, \odot, ϵ) and (M', \bullet, ϵ') are said to be *isomorphic* if they are linked by an isomorphism.

A **semiring** (Golan, 1999) $(S, \oplus, \otimes, \bar{0}, \bar{1})$ is an algebraic structure over a set S where $(S, \oplus, \bar{0})$ is a commutative monoid and $(S, \otimes, \bar{1})$ is a monoid such that \otimes distributes over \oplus , that is for all $x, y, z \in S$:

$$\begin{aligned} z \otimes (x \oplus y) &= (z \otimes x) \oplus (z \otimes y) \\ \text{and } (x \oplus y) \otimes z &= (x \otimes z) \oplus (y \otimes z). \end{aligned}$$

Moreover, $\bar{0}$ is an absorbing element with respect to \otimes , i.e., $\forall x \in S : x \otimes \bar{0} = \bar{0} \otimes x = \bar{0}$.

The **matrix semiring** is the semiring formed by the set of square matrices $S^{k \times k}$, with addition as the elementwise \oplus -addition of S , and multiplication as the standard matrix multiplication defined with \oplus -addition and \otimes -multiplication of S . When the context is clear, we use the standard notation, i.e., for $\mathbf{A}, \mathbf{B} \in S^{k \times k}$, the semiring-matrix addition is noted $\mathbf{A} + \mathbf{B}$, and the semiring-matrix multiplication is noted $\mathbf{A}\mathbf{B}$. The additive neutral element is the matrix noted $\mathbf{0}$ with all entries set to $\bar{0}$, and the multiplicative neutral element is the matrix noted \mathbf{I} with diagonal elements set to $\bar{1}$ and the others set to $\bar{0}$. We note \mathbf{A}^n the n -iterated multiplication of \mathbf{A} and we define $\mathbf{A}^0 \triangleq \mathbf{I}$.

A **weighted finite state automaton** (WFSA) is a 6-tuple $(S, Q, \Sigma, E, \lambda, \rho)$ consisting of a semiring $(S, \oplus, \otimes, \bar{0}, \bar{1})$, a finite set of states Q , a finite set of label Σ , a finite set of transition $E \subseteq Q \times Q \times \Sigma \times S$, an initial mapping function $\lambda : Q \mapsto S$, and a final mapping function $\rho : Q \mapsto S$.

For a transition $e \in E$, we write $o(e)$ its origin state, $d(e)$ its destination state, $\sigma(e)$ its label and $w(e)$ its weight.

We define the set of initial (resp. final) states I (resp. F) as the set of states for which the mapping λ (resp. ρ) is not $\bar{0}$, i.e. $I = \{q : \lambda(q) \neq \bar{0}, q \in Q\}$ (resp. $F = \{q | \rho(q) \neq \bar{0}, q \in Q\}$).

A path $p = (e_1, \dots, e_n) \in E^n$ is a sequence of transition such that, for all $1 \leq i \leq n$, we have $d(e_i) = o(e_{i+1})$. The

set of paths from state $q \in Q$ to state $r \in Q$ is noted $P(q, r)$. The weight of a path is defined as:

$$w(p) = w(e_1) \otimes \dots \otimes w(e_n).$$

The **shortest-distance** between $q \in Q$ and $r \in Q$ is defined as:

$$D[q, r] = \bigoplus_{p \in P(q, r)} w(p).$$

The shortest-distance between two states can be evaluated efficiently via DP using the following recursion:

$$D[q, r] = \bigoplus_{x \in P(s, r)} D[q, x] \otimes D[x, r]. \quad (1)$$

The **weight of an automaton** \mathcal{A} is defined as

$$\nu(\mathcal{A}) = \bigoplus_{(q, r) \in I \times F} \lambda(q) \otimes D[q, r] \otimes \rho(r). \quad (2)$$

The concept of weight of an automaton is central in structured prediction as it formalises the notion of cost function in structured prediction problems representable as a directed acyclic graph (assuming an acyclic automaton \mathcal{A}). Under various disguises, it is the likelihood function of many probabilistic models (e.g. Hidden Markov Models (Rabiner, 1989), Conditional Random Fields (Lafferty et al., 2001), ...), it is a special instance of the sum-product algorithm in factor graphs (Kschischang et al., 2001), and it is a key computation in most, if not all, structured loss functions for speech processing (Graves et al., 2006; Graves, 2012; Heigold et al., 2012; Hoffmeister et al., 2012; Hadian et al., 2018; Pratap et al., 2022; Variani et al., 2022; Laptev et al., 2023).

4. Automatic Differentiation of the Semiring Dot Product

Differentiating through semiring-based operations is very challenging when the semiring is unknown to the AD system. The problem does not come from evaluating the derivatives themselves but rather from scaling the computation. Even if the function to differentiate is efficiently implemented, the process of differentiation quickly becomes inefficient if not helped with specialized implementation of relevant vector-Jacobian products. We dissect here the steps of differentiating through semiring-based computations to illustrate this issue. Our discussion assumes the use of *reverse-mode* differentiation (Rall, 2006; Baydin et al., 2018) where the function to differentiate is first evaluated for a given input (*forward* step) and then the derivatives are backpropagated through the computation graph (*backward* step) via iterated vector-Jacobian products.

Let $(S, \oplus, \otimes, \bar{0}, \bar{1})$ be a semiring with \oplus and \otimes differentiable with respect to both arguments. We assume the following derivatives

$$\frac{\partial x \otimes y}{\partial x}, \quad \frac{\partial x \otimes y}{\partial y}, \quad \frac{\partial x \oplus y}{\partial x}, \quad \text{and} \quad \frac{\partial x \oplus y}{\partial y}$$

to be computable by the AD system with negligible cost.

We consider a dot product operation defined for two vectors $\mathbf{x}, \mathbf{y} \in S^K$ as:

$$z = (x_1 \otimes y_1) \oplus (x_2 \otimes y_2) \oplus \dots \oplus (x_K \otimes y_K). \quad (3)$$

Not only this operation is typical in semiring-based computation, it is also at the heart of DP algorithms as illustrated by (1). A concrete (sequential) implementation of (3) requires computing the \otimes -products:

$$u_i = x_i \otimes y_i \quad (4)$$

and then to perform the \oplus -accumulation

$$z_{K-1} = u_{K-1} \oplus u_K \quad (5)$$

$$z_i = u_i \oplus z_{i+1}. \quad (6)$$

The final result is obtained with $z = z_1$. In this case, assuming constant cost for the \oplus and \otimes operations, the forward computation has a linear complexity $O(K)$.

To compute the partial derivatives $\frac{\partial z}{\partial x_i}$ one has to unroll the computation by application of the chain rule:

$$\frac{\partial z}{\partial x_i} = \frac{\partial z}{\partial z_1} \underbrace{\dots \frac{\partial z_{i-1}}{\partial z_i} \frac{\partial z_i}{\partial u_i}}_{\frac{\partial z}{\partial u_i}} \frac{\partial u_i}{\partial x_i}. \quad (7)$$

An illustration of the computational graph for the forward and backward steps is shown in Fig. 6 (Appendix A).

Since we assume reverse-mode AD, (7) is calculated from left to right and corresponds to a top-down traversal of the computational graph. Note that the computation graphs of the forward and backward steps have the same topology, and therefore, both steps have the same algorithmic complexity¹ and, presumably, a similar run-time. In practice, however, the run-time between the forward and backward steps of semiring-based computations tends to differ significantly. We illustrate the issue in Fig. 1 where we report the run-time of the forward and backward steps of the dot product under the log-semiring (Droste & Kuich, 2009), where $x \oplus y = \log(e^x + e^y)$ and $x \otimes y = x + y$, using two popular AD softwares of the Julia programming language ecosystem: Zygote (Innes, 2018) and Enzyme (Moses & Churavy, 2020). Zygote benefits from highly optimized implementations of vector-Jacobian products for linear algebra

¹We assume that the \oplus and \otimes operations and their derivatives have the same complexity.

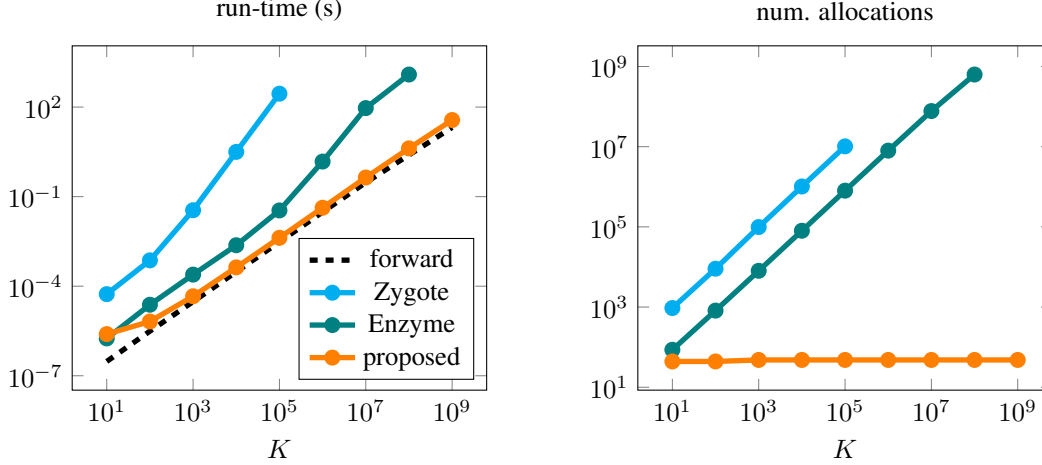


Figure 1. **(left)** Run-time of the product under the log-semiring (forward) and the AD of the computation with Zygote (Innes, 2018), Enzyme (Moses & Churavy, 2020) and Zygote augmented with a generic vector-jacobian product (our proposed method). Our approach scales well as it bypasses the needs to allocate and maintain the computational graph in memory. **(right)** Number of heap allocations realized by the different AD methods. Note that the computation of the semiring dot-product itself does not require dynamic memory allocations.

operations (over the field of reals), Enzyme has a minimal set of optimized implementations of vector-Jacobian products, but achieves high-performance by differentiating the optimize low-level LLVM’s intermediate representation of the program. In both cases the AD process is several orders of magnitude slower than the forward computation despite having the same theoretical complexity (Griewank & Walther, 2008). This performance gap is explained by the need to store in memory the computation graph necessary for the backward step. As one can see in Fig. 1, the total number of allocations made by the AD system grows with the depth of the computation tree strongly impacting the run-time.

5. Differentiation via the “Morphism-Trick”

We have seen that the performance bottleneck in AD of semiring-based computations is related to the depth of the computation graph, which incurs many unavoidable heap allocations. In this section we describe a method to “flatten” the computation graph for the backward step by using a unique implementation of vector-Jacobian products for a broad family of semirings.

Proposition 5.1. (Morphism-trick) Let $(S, \oplus, \otimes, \bar{0}, \bar{1})$ be a semiring such that the monoid $(S, \oplus, \bar{0})$ is isomorphic to $(\mathbb{R}, +, 0)$, and let $\mu : S \rightarrow \mathbb{R}$ be the associated morphism and μ^{-1} its inverse. For all $\mathbf{x}, \mathbf{y} \in S^K$, if

$$z = (x_1 \otimes y_1) \oplus \dots \oplus (x_n \otimes y_n),$$

then:

$$\frac{\partial z}{\partial x_i} = \frac{\partial z}{\partial \mu(z)} \frac{\partial \mu(x_i \otimes y_i)}{\partial x_i \otimes y_i} \frac{\partial x_i \otimes y_i}{\partial x_i}.$$

Proof. Using the morphism, we have

$$\begin{aligned} z &= \mu^{-1}[\mu(z)] \\ &= \mu^{-1}[\mu(x_1 \otimes y_1) + \dots + \mu(x_n \otimes y_n)] \end{aligned}$$

and applying the chain rule gives:

$$\frac{\partial z}{\partial x_i} = \frac{\partial z}{\partial \mu(z)} \frac{\partial \mu(z)}{\partial x_i \otimes y_i} \frac{\partial x_i \otimes y_i}{\partial x_i}.$$

Observing that $\frac{\mu(z)}{\partial x_i \otimes y_i}$ simplifies to $\frac{\partial \mu(x_i \otimes y_i)}{\partial x_i \otimes y_i}$ leads to the desired result. \square

Proposition 5.1 shows that, under certain conditions, the computation of $\frac{\partial z}{\partial w_i}$ does not require any more unrolling the recursion in (6). From a graph perspective, this implies that the computation graph for the backward step has a depth of $O(1)$ (see Figure 2) and no longer requires storing the z_i values saving therefore many heap allocations.

Using this approach, one can incorporate an optimized implementation of vector-Jacobian products in the AD system valid for any semiring equipped with a morphism μ . It is sufficient for the end-user to provide to the AD system routines to evaluate the following computations:

$$\frac{\partial a}{\partial \mu(a)} \frac{\partial \mu(b)}{\partial b}, \frac{\partial a \otimes b}{\partial a}, \text{ and } \frac{\partial a \otimes b}{\partial b}. \quad (8)$$




Figure 2. Computation graph of the backward step of z via the morphism-trick.

An example of a pseudo-algorithm of the vector-Jacobian product of semiring dot product is provided in Alg. 2 (Appendix B).

As an example, consider the log-semiring and observe that the morphism $\mu(x) = e^x$ satisfies the constraints in Proposition 5.1 and is invertible. It follows that

$$\begin{aligned} \frac{\partial a}{\partial \mu(a)} \frac{\partial \mu(b)}{\partial b} &= \frac{e^b}{e^a} \\ \frac{\partial a \otimes_{\log} b}{\partial a} &= \frac{\partial a \otimes_{\log} b}{\partial b} = 1. \end{aligned}$$

Using these definitions in Alg. 2, allows to backpropagate the derivatives without storing the computation graph of the forward step.

In Fig. 1, we measure the run-time of the AD of the semiring dot product using a general implementation of the vector-Jacobian product with Zygote. In this case, there is no significant difference between the forward step and the backward step: the forward graph is not stored and the overhead of the AD process is negligible.

We end this section on a practical remark: the morphism μ allows to shortcut dependencies in the chain of derivatives. As a result, the implementation of backpropagation can be decoupled from the semiring, relying on an abstract definition of the functions in (8). This mirrors the design of finite state automata libraries, as described in Mohri et al. (2000), where concepts of object oriented programming help to separate graph-related algorithms from the semiring abstraction of the transitions' weight.

6. Applicable Semirings

We now present examples of semirings satisfying this constraint, but also possible workarounds to differentiate some idempotent semirings, that is extending the proposed approach to semirings beyond ones strictly satisfying the conditions of Proposition 5.1.

6.1. Semirings Isomorphic to the Real Line

The proposed approach is applicable when the semiring's monoid $(S, \oplus, \bar{0})$ is isomorphic to $(\mathbb{R}, +, 0)$, where “+” is

the natural addition, under the monoid morphism $\mu : S \mapsto \mathbb{R}$. The existence of the isomorphism μ for a semiring S implies the \oplus -operation to be defined as:

$$x \oplus y = \mu^{-1}[\mu(x) + \mu(y)]. \quad (9)$$

This evidently applies with $\mu(x) = x$ for the semiring of real numbers and with the semiring of complex numbers frequently used in quantum finite-state methods (Say & Yakaryılmaz, 2014).

Example 6.1 (log-semiring). The log-semiring on the set \mathbb{R} is defined as follows:

$$\begin{aligned} x \oplus y &\triangleq \tau^{-1} \log(\exp(\tau x) + \exp(\tau y)), \\ x \otimes y &\triangleq x + y, \\ \bar{0} &= -\infty \text{ and } \bar{1} = 0, \end{aligned}$$

where $\tau > 0$ is a configuration temperature. The isomorphism is given by:

$$\mu(x) \triangleq \exp(\tau x)$$

The log-semiring is an important example as it appears in the computation of the log-partition function for conditional random fields and hidden Markov models.

More generally, κ -calculus (Kaniadakis, 2002) defines deformed exponential and logarithmic functions that are basis for power-law distributions:

$$\exp_{\kappa}(x) = \frac{x^{\kappa} - x^{-\kappa}}{2\kappa}, \quad (10)$$

$$\log_{\kappa}(x) = \left(\sqrt{1 + \kappa^2 x^2} + \kappa x\right)^{\frac{1}{\kappa}}. \quad (11)$$

The generalisation of the log-semiring with these functions is compatible with our framework.

Example 6.2 (\log_{κ} -semirings). The \log_{κ} -semiring on the set \mathbb{R} is defined as follows:

$$\begin{aligned} x \oplus_{\kappa} y &= \log_{\kappa}[\exp_{\kappa}(x) + \exp_{\kappa}(y)], \\ x \otimes_{\kappa} y &= x\sqrt{1 + \kappa^2 y^2} + y\sqrt{1 + \kappa^2 x^2}, \\ \bar{0} &= -\infty \text{ and } \bar{1} = 0. \end{aligned}$$

The isomorphism is given by:

$$\mu(x) \triangleq \exp_{\kappa}(x)$$

The \log_{κ} -semiring reduces to the standard log-semiring when κ approaches 0.

6.2. Multi-valued Semiring

Our approach is also applicable to *multi-valued semirings*, i.e. semirings for which the set S is a *ary tuple* of values. A

typical example is the expectation semiring (Eisner, 2002) and its higher order generalization (Li & Eisner, 2009). Consider a couple-valued set $S = S_1 \times S_2$ and the couple of binary operations defined as $\oplus = (\oplus_1, \oplus_2)$. Then, for $(x, a) \in S$ and $(y, b) \in S$, we write:

$$(x, a) \oplus (y, b) = (x \oplus_1 y, a \oplus_2 b). \quad (12)$$

Notice that when dealing with multi-valued semirings, the partial derivatives in (8) are not scalar but Jacobian matrices. This is not an issue in practice as we do not need to instantiate these matrices explicitly to compute the vector-Jacobian products.

Example 6.3 (log-expectation semiring). The log-expectation semiring is defined on the multi-valued set $\mathbb{R} \times \mathbb{R}$

$$\begin{aligned} (x, a) \oplus (y, b) &= (\log(e^x + e^y), a + b) \\ (x, a) \otimes (y, b) &= (x + y, e^x b + e^y a) \\ \bar{0} &= (-\infty, 0) \text{ and } \bar{1} = (0, 0). \end{aligned}$$

The isomorphism is given by:

$$\mu[(x, a)] \triangleq (\exp(x), a)$$

6.3. Idempotent Semirings

The morphism-trick is not applicable generally on idempotent semirings, i.e., semirings for which $x \oplus x = x$ for any $x \in S$, as they do not satisfy (9). Nevertheless, some idempotent semirings can be integrated into our framework by considering weaker constraints and augmenting the semiring with other values necessary for the backward step. Consider the sum

$$z = z_1 \oplus \dots \oplus z_N \quad (13)$$

with $z_i \in S$, and let assume that:

1. There is a mapping $f_z(\cdot)$ such that

$$f_z(z) = f_z(z_1) + \dots + f_z(z_N). \quad (14)$$

2. $f_z(\cdot)$ is invertible at z .

With these two conditions, it is easy to see that the partial derivative of z with respect to any term of the sum undergoes a simplification similar to the morphism case. We draw the reader's attention that f_z is not necessarily a morphism and that another \oplus -summation, say z' , may require different mapping $f_{z'}$, and $f_z(x)$ is not equal to $f_{z'}(x)$ in general for $x \in S$. These conditions do not suppose the existence of the isomorphism μ but require implementing a general vector-Jacobian product for the whole family of mappings

$f_z(\cdot)$. As an example, we illustrate how the (idempotent) tropical semiring, defined with $x \oplus y = \min(x, y)$, can be differentiated with our framework using these relaxed constraints. Let $1_m(x)$ be the indicator function which returns 1 if $x = m$ and 0 otherwise. We define the function f parameterized by m and C as:

$$f_{m,C}(x) = \frac{x 1_m(x)}{C}. \quad (15)$$

For each $z = \min(z_1, \dots, z_n)$ we define $C_z = \sum_{i=1}^n 1_m(z_i)$, i.e., C_z is the number of argument z_i equals to z . Then, we have the following relationship:

$$\begin{aligned} z &= z_1 \oplus \dots \oplus z_N \\ &= \min(z_1, \dots, z_n) \\ &= f_{z,C_z}(z_1) + \dots + f_{z,C_z}(z_n). \end{aligned} \quad (16)$$

Therefore, a tropical \oplus -summation z can be decomposed as a regular sum using $f_{z,C_z}(\cdot)$ and the latter is invertible (only) at z . Consequently, the backpropagation of tropical computations can be theoretically differentiated with our approach as any other semiring satisfying (9). A difference subsists as the partial derivatives

$$\frac{\partial z}{\partial z_i} = \frac{\partial z}{\partial f_{z,C_z}(z)} \frac{\partial f_{z,C_z}(z)}{\partial z_i} = \frac{1_z(z_i)}{C_z} \quad (17)$$

depend on a parameter C_z , which must be stored for the backward step. This is easily done by considering the following composite semiring

$$(\mathbb{R} \times \mathbb{R}^+, \oplus, \otimes, (+\infty, 0), (0, 1)), \quad (18)$$

where

$$(x, C_x) \oplus (y, C_y) = (\min(x, y), C_x + C_y) \quad (19)$$

$$(x, C_x) \otimes (y, C_y) = (\min(x, y), C_x C_y). \quad (20)$$

The advantage of integrating the evaluation of the C_z parameters in the semiring computation is that it avoids the need to implement a specific vector-Jacobian product for the tropical semiring. It is sufficient to use the extended-tropical semiring defined above with the appropriate implementation of (8) and (17). A similar treatment applies to any idempotent semiring with \oplus -addition defined with \min or \max . This includes, among others, the tropical semiring, the arctic semiring, and the Łukasiewicz semiring (Droste & Kuich, 2009).

7. Differentiating DP Algorithms with Finite State Methods

So far, we have exposed our approach on the generalized dot product, the core operation of DP algorithms for structure prediction problems. In this section, we derive a general

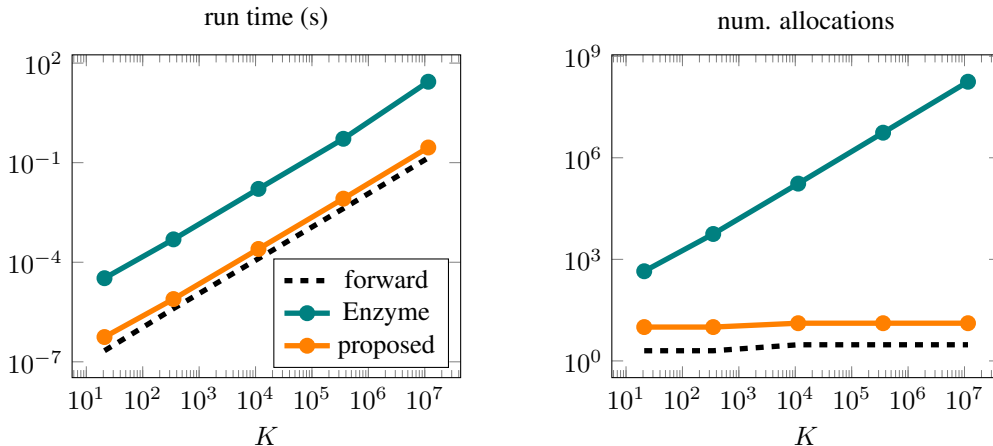


Figure 3. (left) Run-time of computing $\nu(\mathcal{A})$ under the log-semiring (forward) and the AD of the computation, Enzyme (Moses & Churavy, 2020) and Zygote augmented with a general vector-jacobian product (proposed). (right) Number of heap allocations realized by the different AD methods. For both plots, K represents the number of states of plus the number of transitions of \mathcal{A} . The automata used for this benchmark were artificially created by iterated concatenations of the automaton illustrated in Fig. 7a (Appendix C).

differentiable DP algorithm to compute the weight of an automata $\nu(\mathcal{A})$. Recall from Section 3 that $\nu(\mathcal{A})$ formalises all possible cost functions for structure prediction problems on directed acyclic graphs. The DP recursion to evaluate $\nu(\mathcal{A})$, as defined in (1), requires exploring the graph of the structural constraints represented by the automaton \mathcal{A} . This is usually done via a queue-like mechanism, where states to explore are added progressively (Mohri, 2002). From a software perspective, queue-like mechanisms present major challenges as, (i) they are sequential in nature and cannot easily leverage parallel optimization, and (ii) they are difficult to differentiate through as it is necessary to keep track of the internal state of the queue over time to ensure gradient correctness. Because of these difficulties, it is common to have task-specific implementation of DP algorithms (and their derivatives), rather than general ones. A glaring example is the widely used CTC loss (Graves et al., 2006): common implementations, as in PyTorch (Paszke et al., 2019), rely upon a fixed structure to derive a parallel and optimized implementation. Therefore, exploring a slightly different structure, as in (Pratap et al., 2022), necessitates the development of new complex softwares (Hannun et al., 2020).

Inspired by modern numerical methods for graphs (Kepner & Gilbert, 2011), we present, in this section, a general DP algorithm expressed entirely in terms of semiring linear algebra operations, bypassing the difficulty of the queue-based graph exploration. Then, we apply the morphism-trick on this new algorithm to derive an efficient and general of the corresponding vector-Jacobian product.

7.1. Computing $\nu(\mathcal{A})$

Let $\mathcal{A} = (S, Q, \Sigma, E, \lambda, \rho)$ be an acyclic and topologically sorted finite state automaton. We define the matrix $\mathbf{T} \in S^{|Q| \times |Q|}$ such that the element at the i th row and j th column, denoted T_{ij} , is the \oplus -sum of all the transitions' weight going from state i to state j , i.e.,

$$T_{ij} = \bigoplus_{\sigma \in \Sigma} (\delta_{ij}, \sigma). \quad (21)$$

We further define the vector α (resp. ω), taking value in S^Q , such that $\alpha_i = \lambda(i)$ (resp. $\omega_i = \rho(i)$). Since \mathcal{A} is topologically sorted, it follows that \mathbf{T} is strictly upper triangular, i.e. $T_{ij} = \bar{0}$ for when $i \geq j$.

Using \mathbf{T} , α , and ω , the weight of \mathcal{A} can be expressed as:

$$\nu(\mathcal{A}) = \alpha^\top \mathbf{T}^* \omega, \quad (22)$$

where \mathbf{T}^* denotes the star operation in the matrix semiring, i.e., $\mathbf{T}^0 + \mathbf{T}^1 + \mathbf{T}^2 + \dots$. Intuitively, (22) is the sum of the weights of all the paths of length 0, 1, 2, and so on. The intermediate computation

$$\mathbf{d}^\top = \alpha^\top \mathbf{T}^* \quad (23)$$

is known as the single-source shortest-distance algorithm (Mohri, 2002) and is the bulk of the algorithm's complexity.

To derive an efficient recursive algorithm for \mathbf{d} , we make use of the following proposition (adapted from Theorem 2.17 in Ésik & Kuich (2009)):

Proposition 7.1. Let $(S, \oplus, \otimes, \bar{0}, \bar{1})$ be a semiring, k a positive integer, $\alpha \in S^k$ a vector, and $\mathbf{T} \in S^{k \times k}$ a matrix. If $\mathbf{d}^\top = \alpha^\top \mathbf{T}^*$, then, the following equality

holds:

$$\mathbf{d}^\top = \boldsymbol{\alpha}^\top + \mathbf{d}^\top \mathbf{T}.$$

Proof. Plugging $\mathbf{d}^\top = \boldsymbol{\alpha}^\top \mathbf{T}^*$ on the right-hand side of the equality gives:

$$\begin{aligned} \mathbf{d}^\top &= \boldsymbol{\alpha}^\top + (\boldsymbol{\alpha}^\top \mathbf{T}^*) \mathbf{T} \\ &= \boldsymbol{\alpha}^\top (\mathbf{I} + \mathbf{T}^1 + \mathbf{T}^2 + \dots) \end{aligned}$$

and noting that, by definition, $\mathbf{T}^0 = \mathbf{I}$, completes the proof. \square

From Proposition 7.1 and the fact that \mathbf{T} is strictly upper triangular, it follows that the i th dimension of \mathbf{d} is given by:

$$d_i = \alpha_i \oplus \left(\bigoplus_{j=1}^{i-1} d_j \otimes T_{ji} \right). \quad (24)$$

Since d_i depends only on other values d_j when $j < i$, it can be evaluated recursively starting from $d_1 = \alpha_1$.

Since each element T_{ij} will be multiplied and added exactly one times, and taking into account the addition of the vector $\boldsymbol{\alpha}$, computing \mathbf{d} requires $O(|Q| + |E|)$ semiring operations. This matches the complexity of other DP algorithms on topologically sorted directed acyclic graphs (Mohri, 2002). On an ideal parallel system, the cost of adding $\boldsymbol{\alpha}$ becomes constant and the sums in (21) and in (24) can be distributed logarithmically which leads to an algorithmic depth of $O(|Q| \log \sigma)$, where σ is the maximum number of transitions leaving a state in \mathcal{A} .

7.2. Differentiation with the morphism-trick

We now turn to the problem of propagating back the derivatives through the three steps of the computation, namely (21), (23), and the dot product $\mathbf{d}^\top \boldsymbol{\omega}$. The main challenge consists in computing the derivatives of $\nu(\mathcal{A})$ with respect to \mathbf{d} from which the derivatives of $\boldsymbol{\alpha}$, $\boldsymbol{\omega}$ and \mathbf{T} naturally follows. First, observe that irrespective of the semiring we have the recursion:

$$\frac{\partial \nu(\mathcal{A})}{\partial d_i} = \frac{\partial \mathbf{d}^\top \boldsymbol{\omega}}{\partial d_i} + \sum_{j=i+1}^Q \frac{\partial d_j}{\partial d_i} \frac{\partial \nu(\mathcal{A})}{\partial d_j}, \quad (25)$$

which requires computing the derivatives of a semiring dot product $\frac{\partial \mathbf{d}^\top \boldsymbol{\omega}}{\partial d_i}$, which we have already discussed, and the term $\frac{\partial d_j}{\partial d_i}$. Since (24) has the same functional form as (3), we can use the morphism trick, i.e. differentiating through $\mu^{-1}[\mu(d_i)]$, which yields:

$$\frac{\partial d_i}{\partial d_j} = \frac{\partial d_i}{\partial \mu(d_i)} \frac{\partial \mu(d_j \otimes T_{ij})}{\partial d_j \otimes T_{ij}} \frac{\partial d_j \otimes T_{ij}}{\partial d_j}. \quad (26)$$

As in the case of the dot-product, the gradient depends on the semiring only through terms with the same form as (8) which can be provided externally by the practitioner allowing therefore a general implementation of the vector-Jacobian product of $\nu(\mathcal{A})$. The pseudo-code for such implementation is provided in Alg. 1 (Appendix B) as an example. The comparison of run-time of the AD of $\nu(\mathcal{A})$ with Enzyme on the one hand and Zygote augmented with our Alg. 1² on the other hand is reported in Fig. 3. As for the case of the semiring dot product, the speed up is drastic thanks to a minimal number of dynamic memory allocations.

The DP algorithm defined by the recursion in (24) is made available in the open-source WFS library `TensorAutomata.jl`³. Our implementation makes uses of the sparse row compressed storage format (Tinney & Walker, 1967) for the matrix \mathbf{T} , enabling large scale application. The library also implements the corresponding general vector-Jacobian product via the `ChainRules.jl` library⁴. Therefore, our algorithm is differentiable *as is* for all AD engines compatible with `ChainRules.jl`. In practice, we have used our software to differentiate through the DP recursion on graphs with millions of states and hundreds of millions of transitions. Computation on GPU is currently not supported but planned for future releases. Illustration of gradients computed with `TensorAutomata.jl` are shown in Appendix C.

8. Conclusion

This work proposes a principled method to address the implementation of general vector-Jacobian products for DP algorithms and semiring-based computations in general. To avoid reimplementing these routines for each possible semiring, it is assumed that the \oplus -summations in a chosen semiring can be morphed into a regular summation, allowing simplification in the chain of derivatives. We coined our method the “morphism-trick”. As a consequence, the back-propagation graph can be “flattened” saving many dynamic memory allocations, a major performance bottleneck in AD systems. To integrate a new semiring into the framework it is sufficient for the end-user to provide implementation of three elementary derivatives listed in (8). Experimental results show that augmenting AD systems with semiring-agnostic implementation of vector-Jacobian products brings several order of magnitude speed up compared to state-of-the-art AD systems. Our approach is applicable to semirings representing a smooth deformation of the semiring of

²The results for Zygote alone are not shown as they were significantly worse and impractical to measure for non-trivial size automata.

³<https://gitlab.lisn.upsaclay.fr/PTAL/Automata/TensorAutomata.jl>

⁴<https://github.com/JuliaDiff/ChainRules.jl>

real/complex such as the log-semiring or the \log_{κ} semiring. It is also applicable to some idempotent semirings (e.g. the tropical semiring) but requires in this case to augment the semiring with an extra field to evaluate quantities necessary for the backpropagation.

9. Acknowledgement

This project was partially funded by the French Government Defense Innovation Agency (convention études et de recherche no. 2022 65 0079). It has also received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101007666. Caio Corro was supported by the SEMIAMOR (CE23-2023-0005) and InExtenso (ANR-23-IAS1-0004) project grants given by the French National Research Agency (ANR). We also acknowledge support from the JSALT 2023 workshop, hosted at Le Mans University, France, and sponsored by Johns Hopkins University.

Finally, we would like to warmly thank Matthew Wiesner for his careful proofreading and his insightful comments on the manuscript.

References

- Antimirov, V. Partial derivatives of regular expressions and finite automaton constructions. *Theoretical Computer Science*, 155(2):291–319, 1996. ISSN 0304-3975. doi: [https://doi.org/10.1016/0304-3975\(95\)00182-4](https://doi.org/10.1016/0304-3975(95)00182-4). URL <https://www.sciencedirect.com/science/article/pii/0304397595001824>.
- Balakrishnan, A. and Deshmukh, J. V. Differentiable weighted automata. In *ICML 2024 Workshop on Differentiable Almost Everything: Differentiable Relaxations, Algorithms, Operators, and Simulators*, 2024. URL <https://openreview.net/forum?id=k2hIQYqHTh>.
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18(153):1–43, 2018. URL <http://jmlr.org/papers/v18/17-468.html>.
- Bellman, R. On the theory of dynamic programming. *Proceedings of the National Academy of Sciences*, 38(8):716–719, 1952. doi: 10.1073/pnas.38.8.716. URL <https://www.pnas.org/doi/abs/10.1073/pnas.38.8.716>.
- Berthet, Q., Blondel, M., Teboul, O., Cuturi, M., Vert, J.-P., and Bach, F. Learning with differentiable perturbed optimizers. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 9508–9519. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/6bb56208f672af0dd65451f869fedfd9-Paper.pdf.
- Blondel, M., Berthet, Q., Cuturi, M., Frostig, R., Hoyer, S., Linares-Lopez, F., Pedregosa, F., and Vert, J.-P. Efficient and modular implicit differentiation. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 5230–5242. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/228b9279ecf9bbafe582406850c57115-Paper-Conference.pdf.
- Brzozowski, J. A. Derivatives of regular expressions. *J. ACM*, 11(4):481–494, October 1964. ISSN 0004-5411. doi: 10.1145/321239.321249. URL <https://doi.org/10.1145/321239.321249>.
- Correia, G., Niculae, V., Aziz, W., and Martins, A. Efficient marginalization of discrete and structured latent variables via sparsity. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 11789–11802. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/887caadc3642e304ede659b734f79b00-Paper.pdf.
- Corro, C. and Titov, I. Learning latent trees with stochastic perturbations and differentiable dynamic programming. In Korhonen, A., Traum, D., and Márquez, L. (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5508–5521, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1551. URL <https://aclanthology.org/P19-1551/>.
- Droste, M. and Kuich, W. *Semirings and Formal Power Series*, pp. 3–28. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. ISBN 978-3-642-01492-5. doi: 10.1007/978-3-642-01492-5_1. URL https://doi.org/10.1007/978-3-642-01492-5_1.
- Eisner, J. Parameter estimation for probabilistic finite-state transducers. In Isabelle, P., Charniak, E., and Lin, D. (eds.), *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, Philadelphia, Pennsylvania, USA, jul 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073085. URL <https://aclanthology.org/P02-1001/>.

- Ésik, Z. and Kuich, W. *Finite Automata*, pp. 69–104. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. ISBN 978-3-642-01492-5. doi: 10.1007/978-3-642-01492-5_3. URL https://doi.org/10.1007/978-3-642-01492-5_3.
- Golan, J. S. *Hemirings and Semirings: Definitions and Examples*, pp. 1–18. Springer Netherlands, Dordrecht, 1999. ISBN 978-94-015-9333-5. doi: 10.1007/978-94-015-9333-5_1. URL https://doi.org/10.1007/978-94-015-9333-5_1.
- Goodman, J. Semiring parsing. *Comput. Linguist.*, 25(4): 573–605, December 1999. ISSN 0891-2017.
- Graves, A. Sequence transduction with recurrent neural networks. *CoRR*, abs/1211.3711, 2012. URL <http://arxiv.org/abs/1211.3711>.
- Graves, A., Fernández, S., Gomez, F., and Schmidhuber, J. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, pp. 369–376, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595933832. doi: 10.1145/1143844.1143891. URL <https://doi.org/10.1145/1143844.1143891>.
- Griewank, A. and Walther, A. *Evaluating Derivatives*. Society for Industrial and Applied Mathematics, second edition, 2008. doi: 10.1137/1.9780898717761. URL <https://epubs.siam.org/doi/abs/10.1137/1.9780898717761>.
- Hadian, H., Sameti, H., Povey, D., and Khudanpur, S. End-to-end speech recognition using lattice-free mmi. In *Interspeech 2018*, pp. 12–16, 2018. doi: 10.21437/Interspeech.2018-1423.
- Hannun, A. An introduction to weighted automata in machine learning. 2021. URL https://awnihannun.com/writing/automata_ml/automata_in_machine_learning.pdf.
- Hannun, A. Y., Pratap, V., Kahn, J., and Hsu, W. Differentiable weighted finite-state transducers. *CoRR*, abs/2010.01003, 2020. URL <https://arxiv.org/abs/2010.01003>.
- Heigold, G., Ney, H., Schluter, R., and Wiesler, S. Discriminative training for automatic speech recognition: Modeling, criteria, optimization, implementation, and performance. *IEEE Signal Processing Magazine*, 29(6): 58–69, 2012. doi: 10.1109/MSP.2012.2197232.
- Hoffmeister, B., Heigold, G., Rybach, D., Schluter, R., and Ney, H. Wfst enabled solutions to asr problems: Beyond hmm decoding. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(2):551–564, 2012. doi: 10.1109/TASL.2011.2162402.
- Innes, M. Don’t unroll adjoint: Differentiating ssa-form programs. *CoRR*, abs/1810.07951, 2018. URL <http://arxiv.org/abs/1810.07951>.
- Kaniadakis, G. Statistical mechanics in the context of special relativity. *Phys. Rev. E*, 66:056125, Nov 2002. doi: 10.1103/PhysRevE.66.056125. URL <https://link.aps.org/doi/10.1103/PhysRevE.66.056125>.
- Kepner, J. and Gilbert, J. *Graph Algorithms in the Language of Linear Algebra*. Society for Industrial and Applied Mathematics, 2011. doi: 10.1137/1.9780898719918. URL <https://epubs.siam.org/doi/abs/10.1137/1.9780898719918>.
- Kim, Y., Denton, C., Hoang, L., and Rush, A. M. Structured attention networks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=HkE0Nvqlg>.
- Kschischang, F., Frey, B., and Loeliger, H.-A. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, 47(2):498–519, 2001. doi: 10.1109/18.910572.
- Lafferty, J. D., McCallum, A., and Pereira, F. C. N. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pp. 282–289, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1558607781.
- Laptev, A., Bataev, V., Gitman, I., and Ginsburg, B. Powerful and extensible wfst framework for rnn-transducer losses. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023. doi: 10.1109/ICASSP49357.2023.10096679.
- Li, Z. and Eisner, J. First- and second-order expectation semirings with applications to minimum-risk training on translation forests. In Koehn, P. and Michaela, R. (eds.), *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pp. 40–51, Singapore, aug 2009. Association for Computational Linguistics. URL <https://aclanthology.org/D09-1005/>.
- Lombardy, S. and Sakarovitch, J. Derivatives of rational expressions with multiplicity. *Theoretical*

- Computer Science*, 332(1):141–177, 2005. ISSN 0304-3975. doi: <https://doi.org/10.1016/j.tcs.2004.10.016>. URL <https://www.sciencedirect.com/science/article/pii/S0304397504007054>.
- Mensch, A. and Blondel, M. Differentiable dynamic programming for structured prediction and attention. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3462–3471. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/mensch18a.html>.
- Mohri, M. Semiring frameworks and algorithms for shortest-distance problems. *J. Autom. Lang. Comb.*, 7:321–350, 2002. URL <https://api.semanticscholar.org/CorpusID:6022263>.
- Mohri, M., Pereira, F., and Riley, M. The design principles of a weighted finite-state transducer library. *Theoretical Computer Science*, 231(1):17–32, 2000. ISSN 0304-3975. doi: [https://doi.org/10.1016/S0304-3975\(99\)00014-6](https://doi.org/10.1016/S0304-3975(99)00014-6). URL <https://www.sciencedirect.com/science/article/pii/S0304397599000146>.
- Mohri, M., Pereira, F., and Riley, M. Weighted finite-state transducers in speech recognition. *Computer Speech & Language*, 16(1):69–88, 2002. ISSN 0885-2308. doi: <https://doi.org/10.1006/csla.2001.0184>. URL <https://www.sciencedirect.com/science/article/pii/S0885230801901846>.
- Moses, W. and Churavy, V. Instead of rewriting foreign code for machine learning, automatically synthesize fast gradients. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 12472–12485. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/9332c513ef44b682e9347822c2e457ac-Paper.pdf>.
- Niculae, V., Martins, A., Blondel, M., and Cardie, C. SparseMAP: Differentiable sparse structured inference. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3799–3808. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/niculae18a.html>.
- Niculae, V., Corro, C., Nangia, N., Mihaylova, T., and Martins, A. F. T. Discrete latent structure in neural networks. *Foundations and Trends® in Signal Processing*, 19(2):99–211, 2025. ISSN 1932-8346. doi: 10.1561/20000000134. URL <http://dx.doi.org/10.1561/20000000134>.
- Ondel, L., Lam-Yee-Mui, L.-M., Kocour, M., Corro, C. F., and Burget, L. Gpu-accelerated forward-backward algorithm with application to lattice-free mmi. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8417–8421, 2022. doi: 10.1109/ICASSP43922.2022.9746824.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- Perrin, D. *Words*, pp. 1–17. Cambridge Mathematical Library. Cambridge University Press, 1997.
- Povey, D. and Kuan, F. k2, 2020. URL <https://github.com/k2-fsa/k2>.
- Pratap, V., Hannun, A., Synnaeve, G., and Collobert, R. Star temporal classification: Sequence modeling with partially labeled data. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 13392–13403. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/57587d8d6a7ede0e5302fc22d0878c53-Paper-Conference.pdf.
- Rabiner, L. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989. doi: 10.1109/5.18626.
- Rall, L. B. *Perspectives on Automatic Differentiation: Past, Present, and Future?*, pp. 1–14. Springer Berlin Heidelberg, 2006. ISBN 9783540284383. doi: 10.1007/3-540-28438-9_1. URL http://dx.doi.org/10.1007/3-540-28438-9_1.
- Riley, M., Allauzen, C., and Jansche, M. OpenFst: An open-source, weighted finite-state transducer library and its applications to speech and language. In Chelba, C., Kantor, P., and Roark, B. (eds.), *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Tutorial Abstracts*, pp. 9–10, Boulder, Colorado, May 2009. Association for Computational Linguistics. URL <https://aclanthology.org/N09-4005/>.

Rush, A. Torch-struct: Deep structured prediction library.

In Celikyilmaz, A. and Wen, T.-H. (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 335–342, Online, jul 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-demos.38. URL <https://aclanthology.org/2020.acl-demos.38/>.

Say, A. C. and Yakaryilmaz, A. *Quantum Finite Automata: A Modern Introduction*, pp. 208–222. Springer International Publishing, Cham, 2014. ISBN 978-3-319-13350-8. doi: 10.1007/978-3-319-13350-8_16. URL https://doi.org/10.1007/978-3-319-13350-8_16.

Stewart, L., Bach, F., Llinares-Lopez, F., and Berthet, Q. Differentiable clustering with perturbed spanning forests. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 31158–31176. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/637a456d89289769aclab29617ef7213-Paper-Conference.pdf.

Tinney, W. and Walker, J. Direct solutions of sparse network equations by optimally ordered triangular factorization. *Proceedings of the IEEE*, 55(11):1801–1809, 1967. doi: 10.1109/PROC.1967.6011.

Variani, E., Wu, K., Riley, M., Rybach, D., Shannon, M., and Allauzen, C. Global normalization for streaming speech recognition in a modular framework. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.

A. Reverse-Mode differentiation

An illustration of the reverse-mode differentiation for the semiring dot-product is shown in Fig. 6.

B. vector-Jacobian product pseudo-algorithms

- Pseudo-Algorithm 2 for the vector-Jacobian product of the semiring dot-product.
- Pseudo-Algorithm 1 for the vector-Jacobian product of $\nu(\mathcal{A})$.

C. Examples of gradient of $\nu(\mathcal{A})$

- Figure 7 shows examples of gradients of $\nu(\mathcal{A})$ for a finite-state automaton under the log-semiring and the \log_{κ} -semiring.

Algorithm 1 Pseudo-code of a vector-Jacobian product for the weight of an automata \mathcal{A} . \mathbf{T} , α , ω are the input and z is the evaluated output. Δz is the derivative backpropagated by the AD system.

```


 $\nabla \mathbf{d} \leftarrow \mathbf{0}$ 
for  $i$  in  $[Q, \dots, 1]$  do
     $u_i = d_i \otimes \omega_i$  { $u_i$  is reconstructed as it was not stored
    during the forward step}
     $\nabla d_i \leftarrow \Delta z \cdot \frac{\partial z}{\partial \mu(z)} \frac{\partial \mu(u_i)}{\partial u_i} \cdot \frac{\partial u_i}{\partial d_i}$ 
    for  $j$  in  $[i + 1, \dots, Q]$  do
         $\frac{\partial d_i}{\partial d_j} \leftarrow \frac{\partial d_i}{\partial \mu(d_i)} \frac{\partial \mu(d_j \otimes T_{ij})}{\partial d_j \otimes T_{ij}} \frac{\partial d_j \otimes T_{ij}}{\partial d_j}$ 
         $\nabla d_j \leftarrow \nabla d_j + \frac{\partial d_i}{\partial d_j} \cdot \nabla d_j$ 
    end for
end for

return  $\nabla \mathbf{d}$ 
    
```


Algorithm 2 Pseudo-code of the vector-Jacobian product for the semiring dot product operation. \mathbf{x} , \mathbf{y} are the input of the function and z the output. Δz is the derivative backpropagated by the AD system. The operations \otimes , $\frac{\partial z}{\partial \mu(z)} \frac{\mu(u_i)}{u_i}$, $\frac{\partial u_i}{x_i}$, and $\frac{\partial u_i}{y_i}$ are external to the AD system and are provided by the practitioner.

```

 $\Delta \mathbf{x} \leftarrow \mathbf{0}$ 
 $\Delta \mathbf{y} \leftarrow \mathbf{0}$ 
for  $i \in \{1, \dots, K\}$  do
     $u_i = x_i \otimes y_i$  { $u_i$  is reconstructed as it was not stored
    during the forward step}
     $\Delta x_i \leftarrow \Delta z \cdot \frac{\partial z}{\partial \mu(z)} \frac{\partial \mu(u_i)}{\partial x_i} \cdot \frac{\partial u_i}{\partial x_i}$ 
     $\Delta y_i \leftarrow \Delta z \cdot \frac{\partial z}{\partial \mu(z)} \frac{\partial \mu(u_i)}{\partial y_i} \cdot \frac{\partial u_i}{\partial y_i}$ 
end for
return  $\Delta \mathbf{x}$ ,  $\Delta \mathbf{y}$ 
    
```




(a) $\nabla\nu(\mathcal{A})$ (tropical)




(b) $\nabla\nu(\mathcal{A})$ (arctic)

Figure 4. Examples of gradient of $\nu(\mathcal{A})$ (see \mathcal{A} in Fig. 7) under the tropical and arctic semirings. Dashed gray lines indicates zero-value partial derivatives.



(a) \mathcal{B}



(b) $\nabla\nu(\mathcal{B})$ (log-expectation)

Figure 5. (a) finite state automaton \mathcal{B} over the log-expectation semiring: $(\mathbb{R} \times \mathbb{R}, \oplus, \otimes, (-\infty, 0), (0, 0))$ where $(x, a) \oplus (y, b) = (\log(e^x + e^y), a + b)$ and $(x, a) \otimes (y, b) = (x + y, e^x b + e^y a)$. (b) Gradient of z where $\nu(\mathcal{B}) = (z, c)$ with respect to the parameters of \mathcal{B} .

- Figure 4 shows examples of gradients of $\nu(\mathcal{A})$ for a finite-state automaton under the tropical-semiring and the arctic-semiring.
- Figure 5 shows an example of gradient for a multi-valued semiring.




Figure 6. Illustration of reverse-mode automatic differentiation for the semiring-based dot product $z = x_1 \otimes y_1 \oplus \dots \oplus x_k \otimes y_k$. The terms u_i are placeholders for the intermediate \otimes -products $x_i \otimes y_i$. First z is evaluated for a given \mathbf{x} and \mathbf{y} (Fig. 6a) and then, the partial derivatives are backpropagated via iterated vector-Jacobian products from the root to the leaves of the computation graph (Fig. 6b). The whole process requires building and maintaining the computation graph in memory between the forward and backward computation.

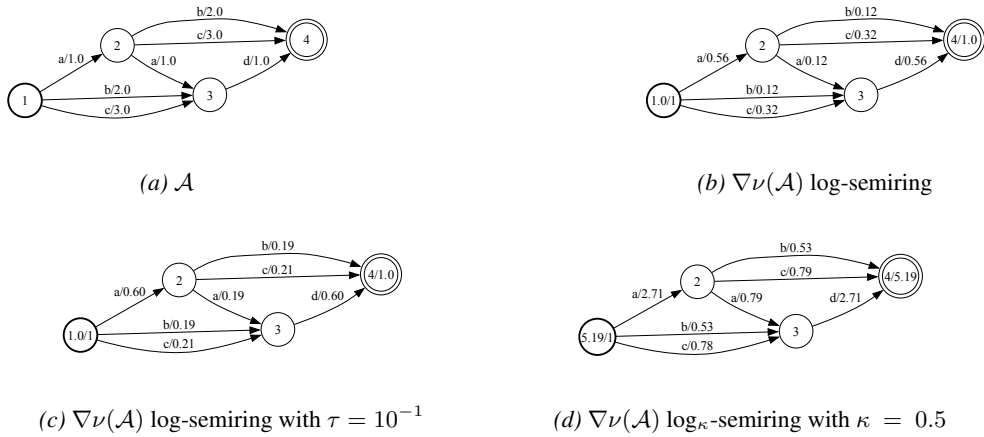


Figure 7. Examples of gradient of $\nu(\mathcal{A})$ evaluated under different semirings with respect to the parameters of \mathcal{A} . A number x displayed as “ x/q ” (resp. “ q/x ”) for a state circle with number q is the partial derivatives of $\nu(\mathcal{A})$ with respect to the initial (resp. final) weight of the state q .