
Learning in Integer Latent Variable Models with Nested Automatic Differentiation

Daniel Sheldon^{1,2} Kevin Winner¹ Debora Sujono¹

Abstract

We develop nested automatic differentiation (AD) algorithms for exact inference and learning in integer latent variable models. Recently, Winner, Sujono, and Sheldon showed how to reduce marginalization in a class of integer latent variable models to evaluating a probability generating function which contains many levels of nested high-order derivatives. We contribute faster and more stable AD algorithms for this challenging problem and a novel algorithm to compute exact gradients for learning. These contributions lead to significantly faster and more accurate learning algorithms, and are the first AD algorithms whose running time is polynomial in the number of levels of nesting.

1. Introduction

In a recent line of work, Winner & Sheldon (2016) and Winner et al. (2017) developed the first exact inference algorithms for a class of hidden Markov models (HMMs) with integer latent variables. Such models are used to model populations that change over time in ecology or epidemiology (Dail & Madsen, 2011; Heathcote, 1965). Standard inference techniques do not apply to these models because marginalization would require summing over the infinite support of the latent variables. Instead, they showed how to reformulate the forward algorithm for HMMs to use *probability generating functions* (PGFs) as its internal representation of probability distributions and messages. These PGFs can be represented compactly and encode all the information about the infinite-support distributions. Inference tasks such as computing the likelihood are reduced to evaluating PGFs and their derivatives.

¹College of Information and Computer Sciences, University of Massachusetts Amherst ²Department of Computer Science, Mount Holyoke College. Correspondence to: Daniel Sheldon <sheldon@cs.umass.edu>.




Figure 1. Function with nested high-order derivatives.

However, the PGFs are complex functions that are defined recursively in terms of high-order derivatives of other PGFs. Figure 1 illustrates such a function. The function f_k , on input x , first computes a value u , then computes the p th derivative of f_{k-1} with respect to u , and then uses the result to compute its own output value y . The function f_{k-1} itself involves nested derivatives of f_{k-2} , and so on, with k total levels of nesting. The goal is to compute the function $f_k(x)$, and possibly derivatives thereof. It is natural to consider automatic differentiation (AD) techniques for this task. However, this is a very difficult setting for AD. There are many levels of nesting and the total order of differentiation may number in the hundreds or thousands.

In this paper we substantially improve AD algorithms for inference in integer latent variable models, and develop new AD algorithms to support learning. For inference, we conceptually simplify the AD techniques of (Winner et al., 2017) and root them more firmly in the AD literature. We show that nested derivatives of univariate functions can be handled by an extension to the basic AD computation model that allows *nested derivative nodes*, which are just derivatives of another function defined by a computation graph (e.g., the “node” $\frac{\partial^p}{\partial u^p} f_{k-1}(u)$ in Figure 1). The AD algorithm needs only a thin adapter to handle the change of scope between the inner and outer functions. For the inference application, this places nearly all of the complexity in the general-purpose AD toolkit and simplifies the application.

We also make substantial stability and speed improvements to the AD algorithms. We show that implementing core AD operations using the *logarithmic number system* (LNS) to accurately represent signed real numbers with high dy-

dynamic range allows the algorithms to scale to very high order derivatives—something which is not possible with a standard floating point representation. We also adopt fast power series composition for core AD operations (Brent & Kung, 1978), which are asymptotically faster than those used in (Winner et al., 2017).

Finally, we contribute new AD algorithms to compute exact *gradients* of the log-likelihood in integer HMMs by extending (higher-order) forward-over-reverse AD to handle nested derivatives. This allows us to compute gradients $\frac{\partial}{\partial \theta} f(x, \theta)$ of functions like the one in Figure 1, where θ is a vector of parameters for the entire nested computation. We show experimentally that our new LNS-based AD algorithms are the only inference algorithms for integer latent variable models that are simultaneously accurate, fast, and stable, and that our novel algorithms for computing exact gradients lead to significantly faster and more accurate parameter estimation.

2. Model and Problem Statement

We consider the *integer hidden Markov model* (HMM) from (Winner et al., 2017), an HMM with integer latent variables n_1, \dots, n_K representing a population that changes over time through the processes of immigration, reproduction, and mortality, and which is partially observed at each time step. The model is:

$$n_k = \sum_{i=1}^{n_{k-1}} z_{k,i} + m_k \quad (1)$$

$$y_k \sim \text{Binomial}(n_k, \rho_k) \quad (2)$$

with the initial condition $n_0 = 0$. The variable n_k is the population size at the k th time step, and y_k is the observed number, assuming that each individual is observed with probability ρ_k . The population size n_k depends on an offspring process and an immigration process. First, each individual present at time $k - 1$ contributes $z_{k,i}$ individuals to the present time step, where $\{z_{k,i}\}$ are iid random variables from the *offspring distribution*. The “offspring” of one individual can include (or not include) itself, immediate offspring, or descendants of more than one generation, depending on the choices of the modeler. In particular, this distribution is used to model both survival and reproduction. Additionally, m_k individuals enter the population, where m_k is drawn from the *immigration distribution*. The offspring and immigration distributions can be arbitrary count-valued distributions, and will be specified through their PGFs as described below.

Problem Statement. The application goal is usually to estimate a parameter vector θ controlling the offspring and immigration distributions given some number of observations of this process (Dail & Madsen, 2011). We will focus on computing the log-likelihood and its gradient, which to-

gether will enable optimization routines to find maximum likelihood estimates. Let K be the total number of time steps and let $y_{1:K} = (y_1, \dots, y_K)$ (similar notation will be used throughout the paper). Our goal is to compute $\log p(y_{1:K}; \theta)$ and $\frac{\partial}{\partial \theta} \log p(y_{1:K}; \theta)$.

Inference via PGFs. Standard HMM inference algorithms such as the forward algorithm (Rabiner, 1989) do not apply here because the latent variables are unbounded. Hence, the messages have infinite length, and marginalization of any variable involves an infinite sum. Winner & Sheldon (2016) and Winner et al. (2017) showed how the forward algorithm can be reformulated to use PGFs to represent messages. Standard inference tasks, such as computing the likelihood, are then converted to the problem of evaluating recursively-defined PGFs.

Definition 1. *The PGF of a (not necessarily normalized) probability distribution $q(n)$ is the power series $F(s) = \sum_{n=0}^{\infty} q(n)s^n$ with probability values as coefficients.*

To formulate the forward algorithm using PGFs, define $\alpha_k(n_k) := p(n_k, y_{1:k})$ and $\gamma_k(n_k) := p(n_k, y_{1:k-1})$. These are the “messages” that are recursively computed within the standard forward algorithm. The PGFs of α_k and γ_k are defined (using the corresponding capital letters) as $A_k(s_k) = \sum_{n_k=0}^{\infty} \alpha_k(n_k) s_k^{n_k}$ and $\Gamma_k(u_k) = \sum_{n_k=0}^{\infty} \gamma_k(n_k) u_k^{n_k}$. The utility of switching to a PGF representation is summarized in the following proposition.

Proposition 1 (Winner & Sheldon 2016; Winner et al. 2017). *Let $F(u)$ and $G(u)$ be the PGFs for the offspring and immigration distributions, respectively. The PGFs Γ_k and A_k satisfy the following recurrence:*

$$\Gamma_k(u_k) = A_{k-1}(F(u_k)) \cdot G(u_k) \quad (3)$$

$$A_k(s_k) = \frac{(s_k \rho_k)^{y_k}}{y_k!} \cdot \Gamma_k^{(y_k)}(s_k(1 - \rho_k)) \quad (4)$$

with the base case $A_0(s_0) = 1$. The likelihood can be recovered from the final PGFs as $p(y_{1:K}) = A_K(1)$.

Equation (3) follows from the model definition and standard manipulations of PGFs, and is well known in the literature on branching processes (Heathcote, 1965). Equation (4) may appear surprising. It includes the y_k th derivative of the function Γ_k from Equation (3). The derivatives are related to the selection of particular terms in the joint PGF of n_k and y_k corresponding to the observed value of y_k .

Proposition 1 gives a recipe to compute the *exact* log-likelihood and its gradient. We need to compute $\log A_K(1; \theta)$ and $\frac{\partial}{\partial \theta} \log A_K(1; \theta)$, where A_K is defined in terms of the prior PGFs through Equations (3) and (4), and we have introduced the parameter vector θ , which includes the detection probabilities ρ_k and any parameters of F and G . Despite the somewhat complex appearance of

the recurrence, it implies a well-defined feed-forward computation to calculate $A_K(1)$ from the constituent functions $F(\cdot)$, $G(\cdot)$ and $A_0(\cdot)$. The key complication is the fact that this function contains deeply nested high-order derivatives: A_K calls $\Gamma_k^{(y_K)}$, which calls A_{K-1} , which calls $\Gamma_{K-1}^{(y_{K-1})}$, and so on. Winner et al. (2017) developed a method based on automatic differentiation to efficiently compute $A_K(1)$, which we will extend in this paper to be more robust and efficient, and to compute gradients.

3. Autodiff Approach

Our problem is to compute $f(x, \theta)$ and $\frac{\partial}{\partial \theta} f(x, \theta)$ for the function $f(x, \theta) = \log A(x, \theta)$, which includes nested high-order derivatives. To set up the proper recursion we will generalize this to the following:

Problem 1: compute $\frac{\partial^p}{\partial x^p} f(x, \theta)$ and $\frac{\partial}{\partial \theta} \frac{\partial^p}{\partial x^p} f(x, \theta)$

Now we will abstract away from the particular f from the previous section and consider any $f : \mathbb{R}^{m+1} \rightarrow \mathbb{R}$ defined according to a particular computation model, where $x \in \mathbb{R}$ will be called the *input* and the vector $\theta \in \mathbb{R}^m$ will be called the *parameters*. We will first describe how to solve both parts of Problem 1 in a basic computation model, and then extend the model to handle nested derivatives.

3.1. Computation Model and Dual Numbers

The basic computation model for f is shown in Algorithm 1, following (Griewank & Walther, 2008). In Line 3, the function φ_j is a *primitive operation* that operates on the variables v_{A_j} , where $A_j \subseteq \{0, \dots, j-1\}$ is the set of *predecessors* of j , and $v_{A_j} = (v_i)_{i \in A_j}$ is the subvector of $v_{0:n}$ corresponding to index set A_j . The predecessor relationship defines the *computation graph* G , a directed acyclic graph (DAG) with edges from i to j for all $i \in A_j$.

Partial computations and dual numbers. We wish to express derivatives $\frac{dv_\ell}{dv_i}$ of a variable v_ℓ with respect to a preceding variable v_i . To make this precise, for $i \leq \ell$, we denote the partial computation from i to ℓ as $f_{i \rightarrow \ell}(v_{0:i})$. This is defined as the function that maps from fixed values of $v_{0:i}$ to the value of v_ℓ obtained by executing the procedure above starting with the assignment to v_{i+1} and ending with the assignment to v_ℓ . A formal recurrence for $f_{i \rightarrow \ell}$ is given in the supplement. We can now define precisely the derivative of v_ℓ with respect to v_i :

$$\frac{dv_\ell}{dv_i} = \frac{\partial}{\partial v_i} f_{i \rightarrow \ell}(v_{0:i}).$$

Definition 2. For $i \leq \ell$, a generalized dual number $\langle v_\ell, dv_i \rangle_p$ is the sequence of derivatives of v_ℓ with respect to v_i up to order p :

$$\langle v_\ell, dv_i \rangle_p = \left(\frac{\partial^q}{\partial v_i^q} f_{i \rightarrow \ell}(v_{0:i}) \right)_{q=0}^p$$

Algorithm 1 Basic computation model

Input: $x \in \mathbb{R}$, $\theta \in \mathbb{R}^m$, **Output:** $f(x, \theta)$

- 1: Set $v_0 = x$
 - 2: Set $v_j = \theta_j$ for $j = 1, \dots, m$
 - 3: Set $v_j = \varphi_j(v_{A_j})$ for $j = m+1, \dots, n$
 - 4: Output v_n
-

We say that $\langle v_\ell, dv_i \rangle_p$ is a dual number of order p with respect to v_i . We will commonly write dual numbers as:

$$\langle s, du \rangle_p = \left(s, \frac{ds}{du}, \dots, \frac{d^p s}{du^p} \right)$$

in which case it is understood that $s = v_\ell$ and $u = v_i$ for some $0 \leq i \leq \ell$, and $f_{i \rightarrow \ell}(\cdot)$ will be clear from context.

This definition generalizes standard (higher-order) dual numbers by explicitly tracking the variable v_i with respect to which we differentiate. In standard forward-mode autodiff, this would always be x . The generalization is important for nested differentiation, where we will instantiate intermediate dual numbers with respect to different variables.

Remark 1 (Dual numbers as Taylor series coefficients). A dual number can be viewed as holding Taylor series coefficients. In particular, the dual number $\langle v_\ell, dx \rangle_p$ encodes the first p coefficients of the Taylor series of $f_{0 \rightarrow \ell}$ about x :

$$f_{0 \rightarrow \ell}(x + \epsilon) = \sum_{q=0}^{\infty} \frac{f_{0 \rightarrow \ell}^{(q)}(x)}{q!} \epsilon^q$$

Most forward AD methods store and propagate dual numbers as truncated Taylor series (Griewank & Walther, 2008; Pearlmutter & Siskind, 2007).

3.2. Higher-Order Forward Mode

We now discuss how to solve the first part of Problem 1—computing $\frac{\partial^p}{\partial x^p} f(x, \theta)$ in the basic computation model—using forward AD. Higher-order forward mode works by propagating dual numbers instead of real numbers through the computation. For this to work, each local computation must be “lifted” to accept dual numbers as inputs and to output a dual number.

Definition 3 (Lifted Function). Let $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}$ be a function of variables u_1, \dots, u_k . The lifted function $\mathcal{L}\varphi$ is the function that accepts as input dual numbers $\langle u_1, dx \rangle_p, \dots, \langle u_k, dx \rangle_p$ and returns the dual number $\langle \varphi(u_1, \dots, u_k), dx \rangle_p$.

We defer the details of *how* a function is lifted for the moment. If we are able to lift each primitive operation in the original procedure, we can execute it using dual numbers (with respect to x) in place of real numbers. The procedure is shown in Algorithm 2. The output \tilde{v}_n is a dual number

Algorithm 2 Higher-Order Forward AD

Input: x, θ **Output:** $\frac{\partial^p}{\partial x^p} f(x, \theta)$

- 1: Set $\tilde{v}_0 = \langle x, dx \rangle_p = (x, 1, 0, \dots, 0)$
 - 2: Set $\tilde{v}_j = \theta_j$ for $j = 1, \dots, m$
 - 3: Set $\tilde{v}_j = \mathcal{L}\varphi_j(\tilde{v}_{A_j})$ for $j = m + 1, \dots, n$
 - 4: Extract and return the p th derivative from \tilde{v}_n
-

Algorithm 3 Forward-Over-Reverse AD

Input: x, θ **Output:** $\frac{\partial}{\partial \theta} \frac{\partial^p}{\partial x^p} f(x, \theta)$

- 1: Run forward AD (Algorithm 2) to compute $\tilde{v}_0, \dots, \tilde{v}_n$
- 2: Initialize $\bar{v}_n = \langle 1, dx \rangle_p = (1, 0, \dots, 0)$
- 3: Initialize $\bar{v}_i = \langle 0, dx \rangle_p = (0, 0, \dots, 0)$ for $i < n$
- 4: For $j = n$ down to $m + 1$ and for all $i \in A_j$:

$$\bar{v}_i \leftarrow \bar{v}_i + \bar{v}_j \cdot \mathcal{L} \frac{\partial}{\partial v_i} \varphi_j(\tilde{v}_{A_j})$$

- 5: Extract and return the p th derivatives from $\bar{v}_1, \dots, \bar{v}_m$.
-

with respect to x , from which we can extract the derivatives $\frac{d^q v_n}{dx^q} = \frac{\partial^q}{\partial x^q} f(x, \theta)$ for $q = 0, \dots, p$. Note that in Line 2, we did not write the parameter values as dual numbers; we will allow in our notation a real number θ to be used as an input to a lifted operation $\mathcal{L}\varphi$, with the understanding that: (1) θ does not depend on x , and (2) θ will be promoted to the dual number $\langle \theta, dx \rangle_p = (\theta, 0, \dots, 0)$.

It is well known how to lift basic mathematical operations:

Proposition 2 (Griewank & Walther 2008). *The arithmetic operations $x + cy$, $x * y$, x/y , x^r , viewed as functions of x and y , and the mathematical functions $\ln(x)$, $\exp(x)$, $\sin(x)$, $\cos(x)$ can be lifted to operate on dual numbers of order p . In each case, the lifted function runs in $O(p^2)$ time.*

As a result, forward-mode autodiff can compute the first p derivatives of any function f that uses only these primitive operations in time $O(p^2)$ times the running time of f .

3.3. Forward-Over-Reverse AD

For the purpose of learning, we need to solve the second part of Problem 1 and compute the gradient $\frac{\partial}{\partial \theta} \frac{\partial^p}{\partial x^p} f(x, \theta)$. This can be accomplished using “forward-over-reverse” AD, which is based on the relatively simple observation that we can switch the order of differentiation to see that

$$\frac{\partial}{\partial \theta} \frac{\partial^p}{\partial x^p} f(x, \theta) = \frac{\partial^p}{\partial x^p} \frac{\partial}{\partial \theta} f(x, \theta),$$

which reveals that what we want is also the higher-order derivative of $g(x, \theta) := \frac{\partial}{\partial \theta} f(x, \theta)$ with respect to x . The well-known *reverse* mode of autodiff (or backpropagation)

provides a procedure to compute $g(x; \theta)$. Then, by executing reverse mode using dual numbers instead of real numbers, we can obtain $\frac{\partial^p}{\partial x^p} g(x; \theta)$. Forward-over-reverse AD is shown in Algorithm 3. It incrementally computes the *adjoint*:

$$\bar{v}_i := \frac{\partial^p}{\partial x^p} \frac{\partial}{\partial v_i} f_{i \rightarrow n}(v_{0:i}),$$

for all i in a reverse sweep through the computation graph. The final adjoint values are dual numbers, from which the derivatives with respect to x of the parameter gradients can be extracted. As in the forward mode, we need to lift the primitive operations of the procedure. In this case the primitive operations are those that compute the partial derivatives of φ_j with respect to its own inputs. When φ_j is a simple mathematical operation ($+$, \times , $/$, \exp , \log , \sin , etc.), these partial derivatives are also simple and can be lifted by standard techniques (cf. Proposition 2).

3.4. Forward Mode with Nested Derivatives

We now extend the forward AD algorithms to handle nested derivatives within the computation procedure. For the most part, we use ideas that were present in (Winner et al., 2017). However, we simplify the conceptual framework considerably, so that all we need to do is lift a function φ_j that takes the derivative of another function, which leads to a conceptual improvement of the AD techniques and a dramatic simplification of the inference application.

Extended computation model: nested derivative nodes.

Although the functions φ_j are usually conceptualized as simple “primitive” operations, they may be arbitrarily complex as long as we know how to lift them—i.e., modify them to propagate dual numbers. So, assume now that one or more of the φ_j functions is a *nested derivative node*, which takes the derivative of some other function g with respect to one of its inputs:

$$\varphi_j(v_{A_j}) := \frac{\partial^q}{\partial v_k^q} g(v_k, \pi). \quad (5)$$

Here we consider $\pi = v_{A_j \setminus k}$ to be “parameters” of the nested computation. We need to reason about how to lift φ_j to propagate dual numbers with respect to x . We will make the following restriction on φ_j , with which we can reason about the sensitivity of φ_j to x through v_k alone:

Assumption 1. *There is no path in G from x to π .*

Lifting a univariate primitive via composition of Taylor series.

We wish to lift a nested derivative node $\varphi_j(v_k, \pi)$. To do so, we will discuss the general procedure to lift a univariate function $\varphi_j(v_k)$. We may temporarily suppress the parameters π from notation because they are constant with respect to x . Here is the general setting. We have completed the partial computation to compute the first p derivatives of $v_k = f_{0 \rightarrow k}(x)$ with respect to x . We wish to compute

Algorithm 4 COMPOSE($\langle u, dv \rangle_p, \langle v, dx \rangle_p$)

- 1: Unpack dual numbers to Taylor coefficients $q_i = \frac{1}{i!} \frac{d^i u}{dv^i}$ and $r_i = \frac{1}{i!} \frac{d^i v}{dx^i}$ for $i = 1$ to p , and scalar value u
 - 2: Let $Q(\tau) = \sum_{i=1}^p q_i \tau^i + O(\tau^{p+1})$
 - 3: Let $R(\epsilon) = \sum_{i=1}^p r_i \epsilon^i + O(\epsilon^{p+1})$
 - 4: Compute the first p coefficients s_1, \dots, s_p of ϵ in the power series $S(\epsilon) = Q(R(\epsilon))$ using a power series composition algorithm.
 - 5: Return $\langle u, dx \rangle_p = (u, 1!s_1, 2!s_2, \dots, p!s_p)$
-

Algorithm 5 DIFF $_q(\langle y, dv \rangle_{q+p})$

- 1: Return $\left(\frac{d^q y}{dv^q}, \dots, \frac{d^{q+p} y}{dv^{q+p}} \right)$ // shift left q positions
-

the first p derivatives of $v_j = \varphi_j(v_k) = \varphi_j(f_{0 \rightarrow k}(x)) = f_{0 \rightarrow j}(x)$. This can be done by composing the Taylor series of φ_j and $f_{0 \rightarrow k}$.

Proposition 3. *Suppose φ_j and $f_{0 \rightarrow k}$ are analytic. Let $\varphi_j(v_k + \tau) = v_j + \sum_{i=1}^{\infty} q_i \tau^i := v_j + Q(\tau)$ be the Taylor series expansion of φ_j about v_k , and let $f_{0 \rightarrow k}(x + \epsilon) = v_k + \sum_{i=1}^{\infty} r_i \epsilon^i := v_k + R(\epsilon)$ be the Taylor series expansion of $f_{0 \rightarrow k}$ about x . Then the Taylor series expansion of $f_{0 \rightarrow j}$ about x is:*

$$f_{0 \rightarrow j}(x + \epsilon) = v_j + Q(R(\epsilon)).$$

The first p coefficients of ϵ in $Q(R(\epsilon))$ encode the first p derivatives of $f_{0 \rightarrow j}$, and can be computed by power series composition algorithms from the first p coefficients of the power series Q and R .

Proposition 3 (proved in the supplement) is the foundation of higher-order forward mode. The input to $\mathcal{L}\varphi_j$ is the dual number $\langle v_k, dx \rangle_p$, which gives us the first p coefficients of $R(\epsilon)$, i.e., $r_i = \frac{1}{i!} \cdot \frac{d^i v_k}{dx^i}$. The function φ_j is typically a simple mathematical primitive (e.g., log, exp, sin), which will have a very special Taylor series $Q(\tau)$. The goal is to compute the first p coefficients of $Q(R(\epsilon))$. In most cases of “true” primitive operations, specialized routines exist to do this in $O(p^2)$ or $O(p \log p)$ time (Griewank & Walther, 2008; Brent & Kung, 1978). For general φ_j , Brent & Kung (1978) describe fast algorithms to compute the first p coefficients of $Q(R(\epsilon))$ from the first p coefficients of Q and R . The compose operation we implement will run in $O(p^{2.5})$ time—see Section 4.

Algorithm 4 shows the COMPOSE operation for two dual numbers $\langle u, dv \rangle_p$ and $\langle v, dx \rangle_p$, which extracts Taylor coefficients from the dual numbers and then performs power series composition to compute $\langle u, dx \rangle_p$.

Lifting a nested derivative node. We now return to the specifics of lifting a nested derivative node $\varphi(v, \pi) =$

Algorithm 6 Lifted Nested Derivative $\mathcal{L}\varphi = \mathcal{L} \frac{\partial^q}{\partial v^q} g(\cdot, \pi)$

Input: Function $g(v, \pi)$, dual number $\langle v, dx \rangle_p$, and π , which does not depend on x

- 1: Unpack scalar v from $\langle v, dx \rangle_p$ and initialize a new dual number $\langle v, dv \rangle_{q+p}$ for the inner scope of φ_j
 - 2: Let $\langle y, dv \rangle_{q+p} = \mathcal{L}g(\langle v, dv \rangle_{q+p}, \pi)$ be the result of the lifted version of g on the newly initialized dual number (recursively apply forward-mode autodiff to g)
 - 3: Let $\langle u, dv \rangle_p = \text{DIFF}_q(\langle y, dv \rangle_{q+p})$
 - 4: Return COMPOSE($\langle u, dv \rangle_p, \langle v, dx \rangle_p$)
-

$\frac{\partial^q}{\partial v^q} g(v, \pi)$ to propagate dual numbers with respect to x (we temporarily drop subscripts without ambiguity). Note that the semantics of the partial differentiation operation $\frac{\partial^q}{\partial v^q}$ would be simple if we were propagating dual numbers with respect to v : we would just shift all derivatives in the sequence lower by q positions (see the DIFF operator in Algorithm 5). So, our solution to nested derivatives will be to temporarily instantiate a new dual number *with respect to v* within the scope of $\mathcal{L}\varphi$, propagate this through $\mathcal{L}g$ to compute the derivatives of g with respect to v , apply the DIFF operator to the result, and then apply the COMPOSE operation to convert back to a dual number with respect to x . The whole procedure is detailed in Algorithm 6.

By using this procedure to lift nested derivative nodes, we can now run forward AD in the extended computation model. Furthermore, because we lift the nested function by recursively applying forward AD to it, we can handle recursively nested derivatives (of univariate functions, i.e., subject to Assumption 1). One can prove inductively that the running time increases by a modest factor, regardless of the number of levels of nested derivatives. Specifically, let d be the maximum order of differentiation, which, for a nested derivative node, is equal to its own order *plus* that of the calling procedure. In the lifted procedure, each operation takes $O(T(d))$ times its original running time, where $T(d)$ is the maximum time of an operation on dual numbers of order d , which in our models is the COMPOSE operation.

Proposition 4. *Higher-order forward AD can be extended to handle nested univariate derivatives by recursively calling forward AD using Algorithm 6. The running time is bounded by $O(T(d))$ times the number of primitive operations, where d is the maximum order of differentiation, and $T(d)$ is the time to compose two power series of order d .*

3.5. Nested Forward-Over-Reverse

We now wish to extend forward-over-reverse AD to the computation model in which φ_j can be a nested derivative node. By examining Algorithm 3, we see in Line 4 that we will need to lift a function of the form $\frac{\partial}{\partial v_i} \varphi_j(v_{A_j}) =$

$\frac{\partial}{\partial v_i} \frac{\partial^q}{\partial v_k^q} g(v_k, \pi)$. By again switching the order of differentiation, the function to be lifted is:

$$\frac{\partial^q}{\partial v_k^q} \left(\frac{\partial}{\partial v_i} g(v_k, \pi) \right)$$

where $\frac{\partial}{\partial v_i} g(v_k, \pi)$ is the scalar-valued function that computes the partial derivative of $g(v_k, \pi)$ with respect to v_i . Observe that this is just another nested derivative node—for the function that computes the partial derivative of g —which we can lift using Algorithm 6. In practice, we don’t need to apply Algorithm 6 separately for each partial derivative. We can compute $\frac{\partial^q}{\partial v_k^q} \frac{\partial}{\partial v_i} g(v_k, \pi)$ for all i simultaneously by (recursively) applying forward-over-reverse AD to g , within a procedure similar to Algorithm 6 (details omitted). This will be more efficient because it only needs one forward and reverse sweep through the computation graph for g . In summary:

Proposition 5. *Forward-over-reverse AD can also be extended to handle nested derivative nodes by recursively calling forward-over-reverse AD using Algorithm 6 (or a more efficient variant). The running time is $O(T(d))$ times that of the original computation.*

4. Application and Implementation

We now return to inference and learning in integer HMMs. The log-likelihood $f(x, \theta) = \log A_K(x; \theta)$ (Equation 3) clearly fits within the extended computation model: it involves only basic mathematical primitives and a nested derivative of Γ_K . Therefore, we can apply the nested forward-over-reverse AD algorithms from Section 3 to compute the log-likelihood and its gradient. Observe that the total number of levels of nesting is K and the maximum order of differentiation is $d = \sum_k y_k$.

This is a unique application of AD. The total order d may number in the hundreds or thousands with ten or more levels of nesting. For comparison, typical values of d , even in high-order applications, are single digits (Griewank & Walther, 2008). Applications of nested AD are rare; reported examples involve only one level of nesting (Siskind & Pearlmutter, 2008; Foo et al., 2008; Maclaurin et al., 2015; Domke, 2012). Because of the very high order and nesting, we faced significant implementation challenges related to two interrelated issues: numerical stability and asymptotic efficiency of power series operations. We describe in this section the steps needed to overcome these.

4.1. Numerical Stability: Logarithmic Number System

The core computations in our approach are lifted primitive operations for dual numbers (Proposition 2), and the COMPOSE operation. Both operate on Taylor series coefficients of the form $r_i = \frac{1}{i!} \frac{d^i v}{dx^i}$ for i from 0 to d . Due to the $\frac{1}{i!}$

factor in each coefficient and the fact that the derivatives themselves may have high dynamic range, overflow and underflow are significant problems and unavoidable for d greater than a few hundred if coefficients are stored directly in floating point. Such problems are familiar in probabilistic inference, where a standard solution is to log transform all values and use stable operations such as logsumexp whenever this transformation needs to be undone.

Our solution is similar, but more difficult for several reasons. To see why, it is helpful to see the nature of algorithms to lift primitive operations, which use convolutions or related recurrences (Griewank & Walther, 2008). As an example, here is the recurrence for the division operation to compute the Taylor coefficients $\{s_i\}$ of the function $1/v(x)$ given the Taylor coefficients $\{r_i\}$ of $v(x)$:

$$s_i = \frac{1}{r_0} \left[r_i - \sum_{j=0}^{i-1} s_j r_{i-j} \right]$$

Observe that we must handle subtraction and negative values, unlike standard probabilistic inference. Also, both multiplication and addition are required in the innermost loop of the algorithm (the right hand side includes multiplication and addition of the most recently computed value s_{i-1}) so it is not possible to batch multiplications and additions, with transformation to and from log-space only between batches.

Our solution, which generalizes the “log-space trick” for probabilistic inference, is to store coefficients using a *logarithmic number system* (LNS, Swartzlander & Alexopoulos 1975), and implement core power series operations in LNS. In LNS, a real number is represented as $X = s \cdot b^x$ where $s = \text{sign}(X)$ is the sign bit and $x = \log_b(|X|)$ is the log of the magnitude of X , stored in fixed precision. The multiplication, division, and power operations are simple and efficient in LNS, while addition and subtraction require more effort. For example, if X and Y are both positive, then $\log_b(X + Y) = x + \log_b(1 + b^{y-x})$. We implemented LNS in C using a floating-point instead of fixed point internal representation to interface more easily with existing mathematical functions. Addition and subtraction use the C standard library’s `log1p` function, and are much more expensive than floating point operations, despite the fact they are, in principle, single “arithmetic operations”, and could be comparable to floating point operations with appropriate hardware (Coleman et al., 2000).

4.2. Fast and Accurate Power Series Operations

The lifting procedures for primitive operations that we describe above all take $O(d^2)$ time and have recurrences that resemble convolutions. Multiplication can be done in $O(d \log d)$ time using FFT-based convolution, and $O(d \log d)$ algorithms for most other primitives can be also be derived using the FFT (Brent & Kung, 1978). Although

our application certainly reaches the regime where FFT is faster than direct convolution, existing FFT implementations require that we transform coefficients out of LNS into floating point, which led to inaccurate results and numerical problems when we tried this. Similar observations have been made in previous applications of FFT for probabilistic inference (Wilson & Keich, 2016). A possible future remedy is to implement the FFT in LNS (Swartzlander et al., 1983), however, our current implementation uses the direct $O(d^2)$ algorithms.

The COMPOSE operation is a bottleneck and should be implemented as efficiently as possible. Winner et al. (2017) used a naive $O(d^3)$ composition algorithm. Brent & Kung (1978) present two fast algorithms—BK 2.1 and BK 2.2—for power series composition. BK 2.2 has the fastest known running time of $O((d \log d)^{1.5})$ with FFT convolution, but may be slower in practice (Johansson, 2015) and we have already ruled out FFT convolution. Instead, we use BK 2.1, which, in our setting, runs in $O(d^{2.5})$ time and achieves substantial speedups over the naive $O(d^3)$ algorithm.

5. Experiments

We conducted experiments on the accuracy, speed, and stability of AD inference algorithms for integer HMMs, and of parameter estimation using AD gradient algorithms.

Inference: Accuracy, Speed, Stability. We evaluated four methods for inference in integer HMMs. TRUNC and TRUNC-FFT are variants of a *truncated forward algorithm*, which is currently used in practice in ecological applications (Dail & Madsen, 2011). These place an *a priori* upper bound N on the values of the hidden variables and then apply the standard forward algorithm for discrete HMMs. They require a convolution of the offspring and immigration distributions to compute the transition probabilities (details are given in the supplement), and take $O(KN^3)$ time with direct convolution and $O(KN^2 \log N)$ time with FFT convolution. They can be implemented in log-space easily, but FFT convolutions must be done in linear space, which can lead to accuracy loss (Wilson & Keich, 2016). Selecting N in the truncated algorithms is a challenge. A too-small value will cut off some of the probability mass and lead to the wrong result, and a too-large value will lead to high running time. For these experiments, we iteratively doubled N until the log-likelihood converged (to a precision of 5 decimal places) or until N reached an absolute limit of 2500, and *only measured the running time of the final iteration*. Note that this is a conservative comparison that hides the cost of tuning this value in practice.

The AD and AD-LNS methods use nested AD to compute $\log A_K(1; \theta)$, with the latter storing coefficients in LNS. With our conceptual simplification of nested AD, these algo-

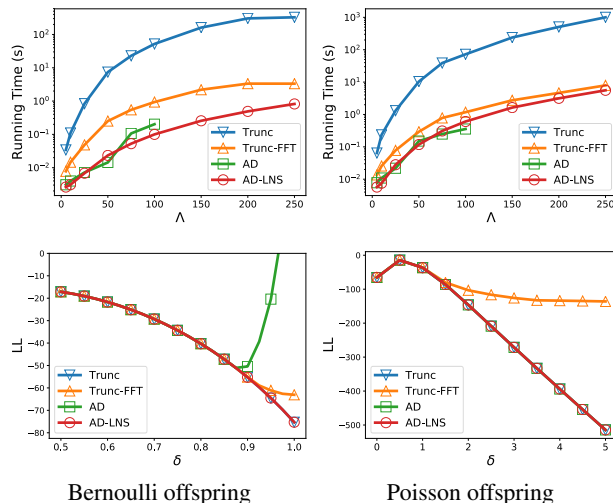


Figure 2. Accuracy, stability, and running time of inference. Top: running time vs. immigration parameter Λ (population size). Bottom: Log-likelihood vs offspring parameter δ . Left: Bernoulli offspring. Right: Poisson offspring.

gorithms are now nearly a direct translation of Equations (3) and (4), with all of work done by the AD routines. We show the full 15-line Python implementation in the supplementary material. AD and AD-LNS both have $O(KY^{2.5})$ running time, where $Y = \sum_k y_k$.

We generated data from integer HMMs, performed inference with each algorithm, and compared the resulting values and running times. In Figure 2, top, we scaled the population size by generating data with immigration distribution $m_k \sim \text{Poisson}(\Lambda)$ for increasing Λ , and fixed offspring distributions $z_{k,i} \sim \text{Bernoulli}(0.5)$ or $z_{k,i} \sim \text{Poisson}(0.5)$. Here, Λ controls the immigration rate, and the expectation of all variables scales in proportion to Λ , as do the parameters N and Y controlling the running time of the algorithms. We ran inference using the true parameters. The plots show running time on a log scale. AD and AD-LNS are the fastest algorithms, but AD is unstable (due to overflow) and does not compute a result for $\Lambda > 100$. TRUNC-FFT is somewhat slower than the AD algorithms. AD-LNS is especially fast for Bernoulli offspring because the Bernoulli PGF is linear, which leads to many zero power series coefficients, and LNS addition of zero is a very fast special case. The Poisson results should be considered the general case.

Figure 2, bottom, compares the accuracy of different methods. Here, we generated data from a model with $\rho_k = 0.5$, immigration $m_k \sim \text{Poisson}(\lambda_k)$ for $\lambda_{1:5} = (12.5, 55, 105, 75, 20)$, and offspring distributions $z_{k,i} \sim \text{Bernoulli}(\delta)$ and $z_{k,i} \sim \text{Poisson}(\delta)$ with $\delta = 0.5$. We performed inference for *different* values of δ . This simulates the situation encountered during estimation when an opti-

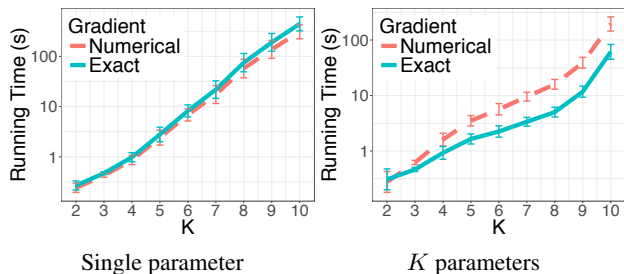


Figure 3. MLE running time vs number of time steps K . Left: with just one parameter, exact and numerical gradients are similar. Right: with K parameters, optimization with exact gradients is much faster. See text for experiment details.




Figure 4. Objective trace of optimizer with different gradients on same problem. Left: a case where iterates are identical. Right: a case where exact gradients lead to a better solution.

mizer queries the log-likelihood at an unlikely parameter setting; it is important that the methods are robust in this case. We see that as δ increases from the true value of 0.5, TRUNC-FFT (in both models) and AD (in Bernoulli) deviate from the correct answer. This is likely due to loss of numerical precision in floating point operations, either in the FFT algorithm or in power series algorithms.

Overall, AD-LNS is the only one of the four inference algorithms that is fast, stable (does not overflow or underflow), and accurate over a wide range of parameter settings.

Learning: Speed and Accuracy. We now evaluate the speed and accuracy of learning. We generate data from an integer HMM and use the L-BFGS (Liu & Nocedal, 1989) algorithm to recover parameters by maximizing the log-likelihood. Based on the inference results, we consider as a baseline the optimizer that uses AD-LNS to compute the likelihood and *numerically* estimates gradients by finite differences. The per-iteration cost to compute numerical gradients is the number of parameters times the cost to compute the log-likelihood (or twice this if central differences are used). For our approach, we used nested forward-over-reverse AD algorithm to compute exact gradients. In this case, the running time to compute the gradient is a constant factor greater than the time to compute the log-likelihood. We predict based on this that as the number of parameters increases, exact gradients will be faster.

Figure 3 shows the results. For each trial, we simulate

20 independent realizations of the integer HMM for up to $K = 10$ time steps, with immigration distribution $m_k \sim \text{Poisson}(5)$ and $\rho = 0.6$. For the left plot, the offspring distribution $z_{k,i} \sim \text{Poisson}(1.2)$ is not time-varying, and we fit only the single parameter δ . In this case, the cost of numerical gradient is only twice that of computing the objective, and the two methods behave similarly. Note that the running time is shown on a log scale. In the right plot, the offspring distribution $z_{k,i} \sim \text{Poisson}(\delta_k)$ is time-varying, with $\delta_k \sim \text{Exp}(1)$, and we fit each of these values for a total of K parameters. In this case, we see that the exact gradient method is several times faster for $K \geq 4$. The running time of both methods increases with K due to the increased cost of inference. We saw similar results across a wide range of models, and expect the benefits of using exact gradients to be even more significant for models with more parameters.

In most cases, the optimizer followed nearly the exact same sequence of parameter values regardless of the gradient computation, but was faster with exact gradients; in some cases the algorithm with exact gradients took more iterations and achieved a better objective value. See Figure 4.

6. Discussion

We introduce new AD techniques for inference in integer HMMs. By cleanly incorporating nested derivatives into the AD computation model, we greatly simplify the algorithms for inference using PGFs. By implementing forward-mode AD operations using the logarithmic number system and using fast power series algorithms, we achieve fast, accurate, and stable algorithms for very high-order derivatives. Our new techniques for nested forward-over-reverse AD allow us to compute exact gradients in integer HMMs for the first time, which leads to significantly faster, and sometimes more accurate, learning procedures.

Related work. An alternate approach to compute nested derivatives is to repeatedly apply AD at each level of nesting starting with the innermost, for example, using source transformation. Pearlmutter & Siskind (2007) and Siskind & Pearlmutter (2008) introduced a tagging mechanism for forward AD to properly handle nesting and avoid “perturbation confusion” (Siskind & Pearlmutter, 2005); this idea is implemented in functional AD tools such as DiffSharp (Baydin et al., 2015). However, since each application of AD increases running time by a constant factor, and the innermost function is “transformed” each time (either through source transformation or by adding a new tagged perturbation), the running time of these approaches is exponential in the number of levels of nesting, while our running time is polynomial. But note that these methods can handle more general functions than ours (i.e., multivariate, not subject to Assumption 1). An interesting avenue of future work is to extend our approach to multivariate nested derivatives.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grants No. 1617533 and 1522054.

References

- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M. Automatic differentiation in machine learning: a survey. *arXiv preprint arXiv:1502.05767*, 2015.
- Brent, R. P. and Kung, H. T. Fast algorithms for manipulating formal power series. *Journal of the ACM (JACM)*, 25(4):581–595, 1978.
- Coleman, J. N., Chester, E. I., Softley, C. I., and Kadlec, J. Arithmetic on the European logarithmic microprocessor. *IEEE Transactions on Computers*, 49(7):702–715, 2000.
- Dail, D. and Madsen, L. Models for estimating abundance from repeated counts of an open metapopulation. *Biometrics*, 67(2):577–587, 2011.
- Domke, J. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*, pp. 318–326, 2012.
- Foo, C.-s., Do, C. B., and Ng, A. Y. Efficient multiple hyperparameter learning for log-linear models. In *Advances in neural information processing systems*, pp. 377–384, 2008.
- Griewank, A. and Walther, A. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008.
- Heathcote, C. R. A branching process allowing immigration. *Journal of the Royal Statistical Society. Series B (Methodological)*, 27(1):138–143, 1965.
- Johansson, F. A fast algorithm for reversion of power series. *Mathematics of Computation*, 84(291):475–484, 2015.
- Liu, D. C. and Nocedal, J. On the limited memory BFGS method for large scale optimization. *Math. Program.*, 45(3):503–528, December 1989.
- Maclaurin, D., Duvenaud, D., and Adams, R. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning (ICML)*, pp. 2113–2122, 2015.
- Pearlmutter, B. A. and Siskind, J. M. Lazy multivariate higher-order forward-mode AD. In *Symposium on Principles of Programming Languages (POPL)*, pp. 155–160, 2007.
- Rabiner, L. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- Siskind, J. M. and Pearlmutter, B. A. Perturbation confusion and referential transparency: Correct functional implementation of forward-mode AD. 2005.
- Siskind, J. M. and Pearlmutter, B. A. Nesting forward-mode AD in a functional framework. *Higher-Order and Symbolic Computation*, 21(4):361–376, Dec 2008.
- Swartzlander, E. E. and Alexopoulos, A. G. The sign/logarithm number system. *IEEE Transactions on Computers*, 100(12):1238–1242, 1975.
- Swartzlander, E. E. J., Chandra, D. V. S., Nagle, H. T. J., and Starks, S. A. Sign/logarithm arithmetic for FFT implementation. *IEEE Transactions on Computers*, C-32(6):526–534, 1983.
- Wilson, H. and Keich, U. Accurate pairwise convolutions of non-negative vectors via FFT. *Computational Statistics and Data Analysis*, 101:300–315, 2016.
- Winner, K. and Sheldon, D. Probabilistic inference with generating functions for Poisson latent variable models. In *Advances in Neural Information Processing Systems 29*, 2016.
- Winner, K., Sujono, D., and Sheldon, D. Exact inference for integer latent-variable models. In *International Conference on Machine Learning (ICML)*, pp. 3761–3770, 2017.