
Sparse Stochastic Zeroth-Order Optimization with an Application to Bandit Structured Prediction

Artem Sokolov
Amazon Research
& Heidelberg University, Germany
sokolov@cl.uni-heidelberg.de

Julian Hitschler
Computational Linguistics
Heidelberg University, Germany
hitschler@cl.uni-heidelberg.de

Stefan Riezler
Computational Linguistics & IWR
Heidelberg University, Germany
riezler@cl.uni-heidelberg.de

Abstract

Stochastic zeroth-order (SZO), or gradient-free, optimization allows to optimize arbitrary functions by relying only on function evaluations under parameter perturbations, however, the iteration complexity of SZO methods suffers a factor proportional to the dimensionality of the perturbed function. We show that in scenarios with natural sparsity patterns as in structured prediction applications, this factor can be reduced to the expected number of active features over input-output pairs. We give a general proof that applies sparse SZO optimization to Lipschitz-continuous, nonconvex, stochastic objectives, and present an experimental evaluation on linear bandit structured prediction tasks with sparse word-based feature representations that confirm our theoretical results.

1 Introduction

Random gradient-free methods [23] provide a simple approach to optimization by applying Gaussian smoothing to an arbitrary function, thus establishing Lipschitz-continuity of the gradient of the smoothed function, and allowing to approximate the gradient by comparisons of function values obtained at randomly perturbed parameter vectors. In optimization theory, related techniques are also named “zeroth-order” [6, 3] or “derivative-free” [7, 31] methods since rather than first- or second-order derivatives, only function values are accessible. The main advantage of stochastic zeroth-order (SZO) methods lie in their applicability to optimization of non-differentiable functions, or in black box situations where nothing but function values are available. The main disadvantage compared to stochastic first-order (SFO), or gradient-based, techniques is a dependency of the convergence speed on the dimensionality of the function to be evaluated.

In machine learning, obtaining a noisy realization of reward or loss function values at a proposed parameter perturbation corresponds to learning with “bandit feedback” [5, 1]. These methods have found renewed interest in the area of reinforcement learning, notably in applications where gradient information is available and SFO techniques would be applicable. The cited advantages of SZO methods for reinforcement learning are their simplicity and robustness against hyperparameter changes [18], the fact that they are highly parallelizable and do not require backpropagation [28], or the improved exploration behavior in parameter space (instead of in action space) [29].

The goal of this paper is to show that the bottleneck of SZO techniques—the dependency of convergence speed on the dimensionality of the parameter space being perturbed—can be reduced to the

expected number of active features over input-output pairs in structured prediction with sparse feature spaces. Such sparsity patterns can be found in sequence labeling tasks that represent sequences by context word representations. An example is the task of noun-phrase chunking where a natural sparsity pattern is defined by n-grams of input words and output chunk labels that can possibly occur in given input-output pairs. Our paper starts with a general proof that adapts the work of [23] to the case of sparse parameter perturbation for Lipschitz-continuous (but not necessarily Lipschitz-smooth), nonconvex, stochastic functions. We present three algorithms that instantiate our theory to structured prediction applications, one based on [23]’s standard two-point function evaluation method, and two methods that aim to reduce feedback complexity to a comparison of function values, or to a comparison of one-point feedback to a running average, with applications to human feedback in mind. In our experimental evaluation, we compare our SZO techniques to standard SFO techniques for bandit structured prediction tasks from natural language processing [33]. Our experimental results on the task of noun-phrase chunking show that the convergence speed of all compared SZO techniques is improved by applying sparse parameter perturbations, reaching a performance close to the standard SFO technique. Furthermore, we present an experimental evaluation on the task of statistical machine translation that investigates the use of a non-differentiable maximum-a-posteriori (MAP) criterion at training and test time in SZO techniques—something that is not possible in SFO methods. On this task our best SZO result outperforms the standard SFO technique, pointing to another possible advantage of SZO techniques.

2 Related Work

Zeroth-order, or gradient-free, stochastic optimization dates back to the finite-difference method for gradient estimation of [11] where the value of each component of w is perturbed separately while holding the other components at nominal value. This technique has since been replaced by more efficient methods based on simultaneous perturbation of all weight vector components [34, 35, 16]. More recent developments apply the simultaneous perturbation principle to bandit learning [5, 40, 1, 3] and extend it from (strongly) convex function to non-convex functions [6, 23]. To our knowledge, the application to structured prediction problems, especially regarding sparse perturbations in active feature space, is novel.

Connections of SZO methods to evolutionary algorithms and reinforcement learning have first been described in [35]. Recent work has applied SZO techniques successfully to policy gradient methods for deep reinforcement learning [29, 28, 26, 18]. The framework of bandit structured prediction [33, 15, 24, 2] is closely related to policy gradient techniques in reinforcement learning [39, 37, 14], for example, linear bandit structured prediction is termed “gradient bandits” in Sutton and Barto [36]. Earlier work in structured prediction has applied SZO techniques for expected loss minimization (a.k.a. minimum risk training) [19, 10].

The quality of SZO methods to provide improved exploration behavior at reduced variance is also appreciated in recent reparameterization approaches. Here a loss function is reparameterized via Gaussian smoothing, and loss values obtained under latent variable perturbations are combined with backpropagation of first-order derivatives [12, 8, 26]. Our work could provide an alternative to the combination of SZO and SFO techniques to an end-to-end application of SZO techniques in reparameterization scenarios.

3 Sparse Zeroth-Order Stochastic Optimization for Nonconvex Objectives

In the following, we give a theoretical analysis of gradient-free optimization by sparse parameter perturbation for Lipschitz-continuous (but not necessarily Lipschitz-smooth), nonconvex, and stochastic objectives. Our analysis builds on [23].

3.1 Problem Statement

We would like to solve a stochastic optimization problem

$$\min_w f(w), \text{ where } f(w) := \mathbb{E}_x[F(w, x)], \tag{1}$$

and \mathbb{E}_x denotes the expectation over inputs $x \in \mathcal{X}$, and $w \in \mathbb{R}^n$ is the dimensionality of the weight vector parameterizing the objective function. We address the the general case of non-convex functions

F for which we furthermore assume Lipschitz-continuity¹, i.e., $F(w, x) \in C^{0,0}$ iff

$$|F(w, x) - F(w', x)| \leq L_0(F(\cdot, x)) \|w - w'\|, \forall w, w', x. \quad (2)$$

[23] show how to achieve a smooth version of an arbitrary function $f(w)$ by Gaussian blurring that assures continuous derivatives everywhere in its domain. In their work, random perturbation of parameters is based on n -dimensional Gaussian random vectors u from a zero-mean isotropic multivariate Gaussian with unit $n \times n$ covariance matrix $\Sigma = I$ s.t.

$$\mathcal{N}(u) = \frac{1}{\sqrt{(2\pi)^n \cdot \det \Sigma}} e^{-\frac{1}{2} u^\top \Sigma^{-1} u}, \quad (3)$$

and a Gaussian approximation of $f(w)$ is defined as $f_\mu(w) = \mathbb{E}_u[f(w + \mu u)]$, where $\mu \geq 0$ is a smoothing parameter. Furthermore, a Lipschitz-continuous gradient even for a non-differentiable original function f can be given by applying standard differentiation rules to $f_\mu(w)$, yielding $\nabla_w f_\mu(w) = \mathbb{E}_u[\frac{f(w + \mu u) - f(w)}{\mu} u]$ (see [23], eq. (21)).

In the case of linear stochastic structured prediction, the functional $f(w)$ is an expectation over inputs x , and the function $F(w, x)$ is defined with respect to a linear model $w^\top \phi(x, y)$ where $\phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^n$ is a joint, possibly sparse, feature representation on inputs and outputs. We express sparsity in parameter perturbation by restricting the Gaussian random vector $u \in \mathbb{R}^n$ to the active, i.e. non-zero, features for each input x , where \upharpoonright is the restriction operator s.t.

$$\mathcal{N}(u, x) = \mathcal{N}(u) \upharpoonright_{u_i \text{ s.t. } \exists y \in Y(x), \phi_i(x, y) \neq 0}. \quad (4)$$

We denote by $\bar{u} \in \mathbb{R}^n$ the Gaussian random vector resulting from a sparse perturbation, and by

$$\bar{n}(x) := \|\bar{u}\|_0 = |\bar{u}_1|^0 + \dots + |\bar{u}_n|^0, \text{ where } 0^0 = 0, \quad (5)$$

the effective number of parameter perturbations for an input x . Based on the notion of sparse perturbation vectors $\bar{u} \in \mathbb{R}^n$, we can redefine the Gaussian approximation as

$$f_\mu(w) = \mathbb{E}_{\bar{u}}[f(w + \mu \bar{u})]. \quad (6)$$

3.2 Convergence Analysis

Our first Lemma applies standard differentiation rules to the the continuous function $f_\mu(w)$ in eq. (6), yielding a Lipschitz-continuous gradient. A full proof, adapting the calculations given in [23], eq. (21), to our case is given in the supplementary material.

Lemma 1.

$$\nabla_w f_\mu(w) = \mathbb{E}_{\bar{u}, x} \left[\frac{F(w + \mu \bar{u}, x) - F(w, x)}{\mu} \bar{u} \right]. \quad (7)$$

[23] show that for Lipschitz-continuous functions F , the variance of the gradient approximation can be bounded by the Lipschitz constant and by the norm of the random perturbation. The term $\mathbb{E}_u[\|u\|^p]$ can itself be bounded by a function of the exponent p and the dimensionality n of the function space. This is how the dependency on n enters iteration complexity bounds and where an adaptation to sparse perturbations has to chime in. In order to adapt [23]'s bounds to sparse gradient-free optimization, we first need to match their bounds on the norms of the random perturbation. The simple case of the squared norm of the random perturbation given below illustrates the idea. If a coordinate i in feature space is not perturbed, no variance $\mathbb{V}_{\bar{u}}[\bar{u}_i]$ is incurred. The smaller the variance, the smaller the factor $\bar{n}(x)$ that directly influences iteration complexity bounds:

$$\begin{aligned} \mathbb{E}_{\bar{u}}[\|\bar{u}\|^2] &= \mathbb{E}_{\bar{u}}[\bar{u}_1^2 + \bar{u}_2^2 + \dots + \bar{u}_n^2] = \mathbb{E}_{\bar{u}}[\bar{u}_1^2] + \mathbb{E}_{\bar{u}}[\bar{u}_2^2] + \dots + \mathbb{E}_{\bar{u}}[\bar{u}_n^2] \\ &= \mathbb{V}_{\bar{u}}[\bar{u}_1] + \mathbb{V}_{\bar{u}}[\bar{u}_2] + \dots + \mathbb{V}_{\bar{u}}[\bar{u}_n] = \bar{n}(x). \end{aligned}$$

For the general case of $\mathbb{E}_{\bar{u}}[\|\bar{u}\|^p]$, $p \geq 2$, we match [23]'s equation (17). We shorten the proof to the parts relevant to sparse perturbations.

¹We use [22]'s notation of function classes where $C^{k,p}$ denotes the class of k times differentiable functions whose p th derivative is Lipschitz continuous.

Lemma 2.

$$\mathbb{E}_{\bar{u}}[\|\bar{u}\|^p] \leq (p + \bar{n}(x))^{p/2}. \quad (8)$$

Proof.

$$\begin{aligned} \mathbb{E}_{\bar{u}}[\|\bar{u}\|^p] &= (2\pi)^{-n/2} \int_{\bar{u}} \|\bar{u}\|^p e^{-\frac{1}{2}\bar{u}^\top \bar{u}} \\ &= (2\pi)^{-n/2} \int_{\bar{u}} \|\bar{u}\|^p e^{-\frac{\tau}{2}\bar{u}^\top \bar{u}} e^{-\frac{1-\tau}{2}\bar{u}^\top \bar{u}} \\ &\leq (2\pi)^{-n/2} \int_{\bar{u}} \left(\frac{p}{\tau e}\right)^{p/2} e^{-\frac{1-\tau}{2}\bar{u}^\top \bar{u}} \\ &= \left(\frac{p}{\tau e}\right)^{p/2} (2\pi)^{-n/2} \int_{\bar{u}} e^{-\frac{1}{2}\bar{u}^\top \left(\frac{I}{1-\tau}\right)^{-1} \bar{u}} \\ &= \left(\frac{p}{\tau e}\right)^{p/2} (2\pi)^{-n/2} \sqrt{(2\pi)^n \cdot \det \frac{I}{1-\tau}} \\ &= \left(\frac{p}{\tau e}\right)^{p/2} \frac{1}{(1-\tau)^{\bar{n}(x)/2}} \\ &\leq (p + \bar{n}(x))^{p/2}. \end{aligned}$$

The first inequality follows from $t^p e^{-\frac{\tau}{2}t^2} \leq \left(\frac{p}{\tau e}\right)^{p/2}$, for $t \geq 0$ (see [23], eq. (80)). The second inequality follows by minimizing the right-hand side in $\tau \in (0, 1)$ (see [23], Lemma 1). \square

Lemma 2 applies the idea illustrated above to higher order norms of random perturbations: If a coordinate is not perturbed, the determinant of the covariance matrix reduces to a product of variances of the active features. This allows us to bound the perturbation factor for each input by $\bar{n}(x) \ll n$. Our convergence theorem analyzes the following SZO algorithm with Sparse Perturbations (SZO-SP).

Algorithm 1 SZO-SP

Input: sequence of learning rates h_k , smoothing parameter $\mu > 0$.
Initialize: $w_0 = 0$
for $k \geq 0$ **do**
 For each w_k , sample x_k and \bar{u}_k .
 Compute $s_\mu(w_k) := \frac{F(w_k + \mu \bar{u}_k, x_k) - F(w_k, x_k)}{\mu} \bar{u}_k$.
 Update $w_{k+1} = w_k - h_k s_\mu(w_k)$.
end for

We present an analysis for nonconvex functionals $F(w, x)$ and $f(w) = \mathbb{E}_x[F(w, x)]$. Furthermore, we assume that each $F(\cdot, x) \in C^{0,0}$ with $L_0(F(\cdot, x)) \leq L_0$, and that $f \in C^{0,0}$.

Theorem 1. Assume a sequence $\{w_k\}_{k \geq 0}$ be generated by Algorithm 1. Let $f(w) \geq f^*$, $\forall w \in \mathbb{R}^n$, and define $\bar{n} \geq \mathbb{E}_{x_k}[\bar{n}(x_k)]$, $\forall k \geq 0$ and $S_N := \sum_{k=0}^N h_k$. Furthermore, let L_1 denote the Lipschitz constant of $\nabla f_\mu(w)$, and let $\bar{\mathcal{U}}_k = (\bar{u}_0, \dots, \bar{u}_k)$ and $\mathcal{X}_k = (x_0, \dots, x_k)$. Then for any $N > 0$, we have

$$\frac{1}{S_N} \sum_{k=0}^N h_k \mathbb{E}_{\bar{\mathcal{U}}_k, \mathcal{X}_k} [\|\nabla f_\mu(w_k)\|^2] \leq \frac{1}{S_N} \left((f_\mu(w_0) - f^*) + \left(\frac{1}{2} L_1 (\bar{n} + 4)^2 L_0^2\right) \sum_{k=0}^N h_k^2 \right). \quad (9)$$

Proof. The proof uses the fact that for a Gaussian approximation $f_\mu(w)$, its gradient is Lipschitz-continuous even if the gradient of $f(w)$ is not. For Lipschitz constant L_1 of $\nabla f_\mu(w)$, we have

$$f_\mu(w_{k+1}) - f_\mu(w_k) - \langle \nabla f_\mu(w_k), w_{k+1} - w_k \rangle \leq \frac{1}{2} L_1 \|w_k - w_{k+1}\|^2.$$

Applying an update step of Algorithm 1 lets us rewrite w_{k+1} as $w_k - h_k s_\mu(w_k)$, leading to

$$f_\mu(w_{k+1}) \leq f_\mu(w_k) - h_k \langle \nabla f_\mu(w_k), s_\mu(w_k) \rangle + \frac{1}{2} L_1 h_k^2 \|s_\mu(w_k)\|^2.$$

Taking expectations in \bar{u}_k and x_k , we can apply equation (13), and get

$$\mathbb{E}_{\bar{u}_k, x_k} [f_\mu(w_{k+1})] \leq f_\mu(w_k) - h_k \|\nabla f_\mu(x_k)\|^2 + \frac{1}{2} L_1 h_k^2 \|s_\mu(w_k)\|^2.$$

Furthermore, the expected squared norm of $s_\mu(w_k)$ can be bounded by

$$\begin{aligned} \mathbb{E}_{\bar{u}_k, x_k} [\|s_\mu(w_k)\|^2] &= \frac{1}{\mu^2} \mathbb{E}_{\bar{u}_k, x_k} [(F(w_k + \mu \bar{u}_k, x_k) - F(w_k, x_k))^2 \|\bar{u}_k\|^2] \\ &\leq \frac{1}{\mu^2} \mathbb{E}_{\bar{u}_k, x_k} [\|w_k + \mu \bar{u}_k - w_k\|^2 L_0^2 \|\bar{u}_k\|^2] \\ &= \mathbb{E}_{\bar{u}_k, x_k} [L_0^2 \|\bar{u}_k\|^4] \leq L_0^2 (\bar{n} + 4)^2. \end{aligned}$$

The first inequality follows by the assumption of Lipschitz continuity of all $F(\cdot, x) \in C^{0,0}$ with $L_0(F(\cdot, x)) \leq L_0$. The second inequality follows by applying Lemma 2 for $p = 4$ and taking the expectation $\mathbb{E}_{x_k} [\bar{n}(x_k)]$ whose upper bound is denoted by \bar{n} . This yields the following inequality

$$\mathbb{E}_{\bar{u}_k, x_k} [f_\mu(w_{k+1})] \leq f_\mu(w_k) - h_k \|\nabla f_\mu(x_k)\|^2 + \frac{1}{2} L_1 L_0^2 (\bar{n} + 4)^2 h_k^2.$$

Taking expectations over $\bar{\mathcal{U}}_k$ and \mathcal{X}_k , and summing up over $k = 0, \dots, N$ yields the result. \square

Theorem 1 gives a non-asymptotic bound on the expected squared gradient norm for any sequence of iterates of Algorithm 1. [23], Section 7, furthermore show that by an appropriate choice of learning rates h_k and smoothing parameters μ , the iteration complexity for nonconvex zeroth-order optimization, i.e., the number of iterations necessary to guarantee an accuracy of ϵ for the expected squared norm of the gradient of f_μ , can be analyzed as $\mathcal{O}(\frac{n^3}{\epsilon^2})$. The same algebraic manipulations can be applied to result (9) that is adapted to sparse perturbations, leading to the following Corollary:

Corollary 1.

$$\mathbb{E}_{\bar{\mathcal{U}}_N, \mathcal{X}_N} [\|\nabla f_\mu(w_N)\|^2] \leq \epsilon \text{ if } N \geq \mathcal{O}\left(\frac{\bar{n}^3}{\epsilon^2}\right). \quad (10)$$

Corollary 1 shows that the factor n^3 that the ϵ -accuracy of SZO methods for nonconvex optimization suffers in comparison to nonconvex SFO optimization can be reduced to the factor \bar{n}^3 that can benefit from strong sparsity patterns. A full proof, adapting the calculations given in [23], Section 7, to our case is given in the supplementary material.

4 Algorithms for Bandit Structured Prediction

Update Rules. Algorithm 1 defines an update rule by a *two-point function evaluation*. Two-point update rules have been introduced as simultaneous perturbation gradient approximation by [34], and have later become standard in gradient-free optimization [6, 3].

A possibility to reduce feedback complexity, with applications that obtain feedback from humans in mind, is to ask for a boolean-valued, relative comparison of function values. Algorithm 1 can be modified easily to use the following *function comparison* update rule:

$$\text{If } F(w_k + \mu \bar{u}_k, x_k) < F(w_k, x_k), \text{ Update } w_{k+1} = w_k + \frac{h_k}{\mu} \bar{u}_k.$$

This rule can be seen as an SZO alternative to the dueling bandits algorithm of [40].

A similar effect as a two-point function evaluation can be achieved by comparing a one-point function evaluation against a running average of function evaluations performed so far. This technique is known as control variates in Monte Carlo simulation [27]. The idea is to augment a random variable X whose expectation is sought, by another random variable Y to which X is highly correlated. Let

Y denote the control variate, and let \bar{Y} denote its expectation. Then the quantity $X - Y + \bar{Y}$ is an unbiased estimator of $\mathbb{E}[X]$. The variance reduction effect of control variates can be seen by computing the variance of this quantity: $\text{Var}(X - Y) = \text{Var}(X) + \text{Var}(Y) - 2\text{Cov}(X, Y)$. Choosing a control variate such that $\text{Cov}(X, Y)$ is positive and large enough, the variance of the gradient estimate will be reduced. In our case, the random variable is a one-point gradient approximation evaluated at a sampled x_k and \bar{u}_k . A well-known control variate from reinforcement learning [39] incorporates the average cumulative reward (or loss) as so-called baseline into the estimator, yielding a *baseline comparison* update rule:

$$\text{Compute } Y_k = \frac{1}{k} \sum_{j=1}^k F(w_j + \mu \bar{u}_j, x_j) \bar{u}_k, \text{ Update } w_{k+1} = w_k - \frac{h_k}{\mu} (F(w_k + \mu \bar{u}_k, x_k) - Y_k) \bar{u}_k.$$

Note the similarity of the above rule to the two-point feedback rule where Y_k plays the role of a slowly changing unperturbed function value. A similar rule has been used in gradient-free reinforcement learning [29].

Linear Structured Prediction Models. One possibility to instantiate the algorithms described above to structured prediction is to encode a task loss evaluation under MAP prediction as initial function F . In this work, we assume MAP prediction under a linear model $\hat{y}(x, w) = \text{argmax}_{y \in \mathcal{Y}(x)} w^\top \phi(x, y)$. Let $\Delta : \mathcal{Y} \rightarrow [0, 1]$ be a task loss function for structured prediction, e.g. 1-BLEU score for machine translation, then the initial function F can be defined as

$$F(w, x) := \Delta(\hat{y}(x, w)). \quad (11)$$

This criterion has been used in [19, 10]. A similar deterministic criterion has been used for SZO optimization for deep reinforcement learning by [28, 18].

Note that the MAP criterion (11) is not only not differentiable, but also not Lipschitz-continuous. This discontinuous criterion can be replaced by the smooth criterion (12) that is obtained by computing an annealed expected loss criterion under temperature parameter $\gamma \geq 0$:

$$F(w, x) := \mathbb{E}_{p_{w, \gamma}(y|x)}[\Delta(y)] = \sum_{y \in \mathcal{Y}(x)} \Delta(y) \frac{(\exp w^\top \phi(x, y))^\gamma}{\sum_{y \in \mathcal{Y}(x)} (\exp w^\top \phi(x, y))^\gamma}. \quad (12)$$

This criterion approaches criterion (11) for $\gamma \rightarrow \infty$ (see [32]), thus both criteria can be used interchangeably in experiments.

Stochastic First-Order Optimization. In our experiments, we will compare SZO approaches to gradient-based SFO algorithms. The latter will serve as upper bound in terms of convergence speed. SFO algorithms for linear bandit structured prediction have been introduced in [33]. We employ an algorithm that optimizes an expected loss objective $\mathbb{E}_{p(x)p_w(y|x)}[\Delta(y)]$ by performing gradient descent on the stochastic gradient $s(w) = \Delta(y) \frac{\partial \log p_w(y|x)}{\partial w}$.

5 Experiments

Experimental Design. The structured prediction tasks in our experiments have been established for bandit structured prediction using SFO optimization in [33]. The task of noun-phrase chunking uses high-dimensional, but very sparse, word-based feature representations. The task of statistical machine translation (SMT) is a sequence-to-sequence prediction problem using a small, dense feature representation. The goal of the latter task is to investigate potential gains of using the same, non-differentiable, criteria at training and test time. All tasks are based on linear models.

Training for all tasks was done by supervised-to-bandit conversion where bandit feedback is simulated by evaluating Δ against gold standard structures which are never revealed to the learner. Δ is a loss function obtained from a task reward, namely 1-BLEU score at sentence level [21] for SMT, and 1-F1 score for chunking.

Convergence speed is evaluated by plotting the average cumulative loss against iterations. In our experiments, we use the MAP criterion (11) and define average cumulative loss/reward at iteration t as $\bar{\Delta}_t = \frac{1}{t} \sum_{k=1}^t \Delta(\hat{y}_k(x_k, w_k + \mu \bar{u}_k))$. This criterion corresponds to regret under the assumption of zero loss for the optimal model.




Figure 1: Average cumulative loss on training data for sparse noun phrase chunking (left) and dense statistical machine translation (right). All results were obtained with fixed hyperparameters for each task, averaged over 3 runs with different random seeds for each algorithm, showing mean results in bold lines, 2 standard deviations in filled areas.

Test set evaluation is done by the standard offline evaluation for the respective tasks. Machine translation is evaluated by measuring the corpus-based BLEU score [25] against an unseen test set of human reference translations. Chunking is evaluated by F1 score on an unseen test set. The early stopping point for test set evaluation is chosen according to a standard online-to-batch conversion by selecting the model that performs best on development data for final evaluation on test data. All evaluation results are averaged over three runs with different random seeds. Further comparison points are out-of-domain lower bounds, in-domain upper bounds, and the SFO algorithm (Expected Loss Minimization) of [33].

Sparse Models for Noun-Phrase Chunking. We followed [30] in applying a linear conditional random field (CRF) model to the noun phrase chunking task on the CoNLL-2000 dataset. The original training set was split into a development set (top 1,000 sent.) and a training set (7,936 sent.); the test set was kept intact (2,012 sent.). Training for bandit learning on the chunking task is done by cold starting the models from $w_0 = \mathbf{0}$.

For an input sentence x , each CRF node x^i carries an observable word and its part-of-speech tag, and has to be assigned a chunk tag c^i out of 3 labels: **Beginning**, **Inside**, or **Outside** (of a noun phrase). Chunk labels are not nested. As in [30], we use second order Markov dependencies (bi-gram chunk tags), such that for sentence position i , the state is $y^i = c^{i-1}c^i$, increasing the label set size from 3 to 9. The model uses feature templates that combine these labels with uni-, bi-, and tri-grams of Part-of-Speech tags, and with uni- and bi-grams of words, leading to high sparsity pattern of on average 0.25% active features for over 1.5M features on the training set.

Dense Feature Models for SMT. The learning goal in our SMT experiment is re-ranking of n -best translation lists of size 5,000 using a linear combination of 14 dense features. The experiments are based on the *cdec* [4] framework. The experimental setup is French-to-English domain adaptation from Europarl to NewsCommentary domains using the data of [13].

The bandit learning algorithms were initialized with the learned weights of the out-of-domain median model and used 40,444 parallel in-domain sentence pairs. Bandit feedback was simulated by evaluating the sampled translation against the reference using as loss function Δ a smoothed per-sentence $1 - \text{BLEU}$ (zero n -gram counts being replaced with 0.01). The possible range of improvements is given by the difference in performance of 0.257 BLEU for the out-of-domain model, compared to 0.284 BLEU for an in-domain model, evaluated on a separate in-domain test set of 2,007 parallel sentence pairs.

Table 1: Test set evaluation for chunking under F1 score, and for machine translation under BLEU. Results for stochastic learners are shown for SFO, and SZO algorithms with perturbation of all parameters (ALL) or for sparse perturbations of parameters for active features only (SPARSE). All results are averaged over three runs with different random seeds for each algorithm. Best SZO results are shown in bold face; higher results are better.

Update Rule	Test F1 Chunking	Test BLEU SMT
Function Comparison (ALL)	0.810	0.265
Baseline Comparison (ALL)	0.819	0.266
Two-Point Evaluation (ALL)	0.841	0.273
Function Comparison (SPARSE)	0.842	-
Baseline Comparison (SPARSE)	0.869	-
Two-Point Evaluation (SPARSE)	0.888	-
SFO	0.908	0.263

Experimental Results for Chunking. The left plot in Figure 1 confirms our theoretical findings by showing faster convergence for SPARSE perturbations (dashed curves) over perturbation of ALL parameters (solid curves) for each update rule. SPARSE perturbations for *two-point* updates lead to a similar convergence speed as the SFO algorithm, followed by SPARSE perturbations for *baseline comparison* updates, and SPARSE perturbation for *function comparison*. The hyperparameters for all algorithms were kept fixed at constant learning rate $h = 0.01$ and exploration parameter $\mu = 0.01$.

Test set results for chunking are shown Table 1. The results were obtained by tuning hyperparameters on development sets for constant learning rate h in the range of 10^{-2} to 10^{-3} , and exploration parameter μ in the range of 10^{-1} to 10^{-2} . Best development settings for all algorithms were obtained close to the time horizon of 4M iterations, showing an undertraining behavior in all cases. However, we find the same relations in test set performance as were obtained for convergence speed, with SPARSE *two-point* updates coming close to the SFO result.

Experimental Results for Machine Translation. Figure 1 depicts the convergence behavior of SZO optimization for the dense SMT re-ranking task, using perturbations of ALL features. Out of the SZO update rules, *two-point* update lead to fastest convergence, followed by *baseline comparison* and *function comparison* update rules. All average cumulative loss results were obtained under the same hyperparameter settings of $h = 0.001$ and $\mu = 0.001$ for all algorithms.

Test set results are shown in 1. Here optimal hyperparameter settings were adjusted on a development set of 1,064 parallel in-domain sentence pairs for constant learning rate h in the range of 10^0 to 10^{-5} , and exploration parameter μ in the range of 10^{-2} to 10^{-6} . Final results are obtained by averaging three independent runs using the hyperparameters found on the development set. Best development settings for all algorithms are obtained at 1-4M iterations, except for two-point feedback that reached an optimum already around 100k iterations. The goal of the SMT experiment was to show a possible advantage of using the same non-differentiable MAP criterion at training and test time for SZO algorithms. As we can see, all SZO results outperform the SFO result. In the range of 2.7 BLEU points between out-of-domain and in-domain models, we can achieve an improvement of 1.6 BLEU points by SZO with *two-point* updates, which is also an improvement of 1.0 BLEU points over the SFO Algorithm.

Further Experiments. Further experiments using sparse, word-based models for multiclass text classification (which can be viewed as a degenerate structured prediction task) are given in the supplementary material.

6 Conclusion

The theoretical contribution of this paper is to show that the main bottleneck in SZO optimization—the dependency of the iteration complexity on the dimensionality of the function to be perturbed—can be reduced to the expected number of active features in sparse structured prediction scenarios. We presented experimental results on linear structured prediction tasks that confirm our theoretical results.

Furthermore, we showed that it can be advantageous to use the same criterion at train and test time, something that is impossible for SFO algorithms if this criterion is non-differentiable.

The experiments in this paper were obtained by perturbing the discontinuous deterministic MAP criterion (11). Using a smooth annealed criterion (12) yields similar results, and allows to match our experiments with our (and other existing) theory that assumes at least Lipschitz continuity of perturbed functions. To our knowledge, existing work on SZO optimization for discontinuous functions is concerned with a theory of differentiation and asymptotic convergence results [20, 38], thus non-asymptotic convergence analysis for discontinuous SZO is an interesting open problem.

Acknowledgments.

This research was supported in part by the German research foundation (DFG), and in part by a research cooperation grant with the Amazon Development Center Germany.

References

- [1] Agarwal, A., Dekel, O., and Xiao, L. (2010). Optimal algorithms for online convex optimization with multi-point bandit feedback. In *COLT*, Haifa, Israel.
- [2] Bahdanau, D., Brakel, P., Xu, K., Goyal, A., Lowe, R., Pineau, J., Courville, A., and Bengio, Y. (2017). An actor-critic algorithm for sequence prediction. In *ICLR*, Toulon, France.
- [3] Duchi, J. C., Jordan, M. I., Wainwright, M. J., and Wibisono, A. (2015). Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806.
- [4] Dyer, C., Lopez, A., Ganitkevitch, J., Weese, J., Ture, F., Blunsom, P., Setiawan, H., Eidelman, V., and Resnik, P. (2010). cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models. In *ACL Demo*, Uppsala, Sweden.
- [5] Flaxman, A. D., Kalai, A. T., and McMahan, H. B. (2005). Online convex optimization in the bandit setting: gradient descent without a gradient. In *SODA*, Philadelphia, PA.
- [6] Ghadimi, S. and Lan, G. (2012). Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 4(23):2342–2368.
- [7] Jamieson, K. G., Nowak, R. D., and Recht, B. (2012). Query complexity of derivative-free optimization. In *NIPS*, Lake Tahoe, CA.
- [8] Jang, E., Gu, S., and Poole, B. (2017). Categorical reparameterization with gumbel-softmax. In *ICLR*, Toulon, France.
- [9] Kakade, S. M., Shalev-Shwartz, S., and Tewari, A. (2008). Efficient bandit algorithms for online multiclass prediction. In *ICML*, Helsinki, Finland.
- [10] Keshet, J., McAllester, D., and Hazan, T. (2011). PAC-Bayesian approach for minimization of phoneme error rate. In *ICASSP*, Prague, Czech Republic.
- [11] Kiefer, J. and Wolfowitz, J. (1952). Stochastic estimation of the maximum of a regression function. *Annals of Mathematical Statistics*, 23(3):462–466.
- [12] Kingma, D. P. and Welling, M. (2014). Auto-encoding variational bayes. In *ICLR*, Banff, Canada.
- [13] Koehn, P. and Schroeder, J. (2007). Experiments in domain adaptation for statistical machine translation. In *WMT*, Prague, Czech Republic.
- [14] Konda, V. R. and Tsitsiklis, J. N. (2000). Actor-critic algorithms. In *NIPS*, Vancouver, Canada.
- [15] Kreutzer, J., Sokolov, A., and Riezler, S. (2017). Bandit structured prediction for neural sequence-to-sequence learning. In *ACL*, Vancouver, Canada.

- [16] Kushner, H. J. and Yin, G. G. (2003). *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, second edition.
- [17] Lewis, D. D., Yang, Y., Rose, T. G., and Li, F. (2004). RCV1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, 5:361–397.
- [18] Mania, H., Guy, A., and Recht, B. (2018). Simple random search provides a competitive approach to reinforcement learning. *CoRR*, abs/1803.07055.
- [19] McAllester, D. and Keshet, J. (2011). Generalization bounds and consistency for latent structural probit and ramp loss. In *NIPS*, Granada, Spain.
- [20] Moreau, L. and Aeyels, D. (2000). Optimization of discontinuous functions: A generalized theory of differentiation. *SIAM Journal on Optimization*, 11(1):53–69.
- [21] Nakov, P., Guzmán, F., and Vogel, S. (2012). Optimizing for sentence-level bleu+1 yields short translations. In *COLING*, Bombay, India.
- [22] Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*. Springer.
- [23] Nesterov, Y. and Spokoiny, V. (2015). Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, (17):527–566.
- [24] Nguyen, K., Daumé, H., and Boyd-Graber, J. (2017). Reinforcement learning for bandit neural machine translation with simulated feedback. In *EMNLP*, Copenhagen, Denmark.
- [25] Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *ACL*, Philadelphia, PA.
- [26] Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. (2018). Parameter space noise for exploration. In *ICLR*, Vancouver, Canada.
- [27] Ross, S. M. (2013). *Simulation*. Elsevier, fifth edition.
- [28] Salimans, T., Ho, J., Chen, X., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning. *CoRR*, abs/1703.03864.
- [29] Sehnke, F., Osendorfer, C., Rückstieß, T., Graves, A., Peters, J., and Schmidhuber, J. (2010). Parameter-exploring policy gradients. *Neural Networks*, 23(4):551–559.
- [30] Sha, F. and Pereira, F. (2003). Shallow parsing with conditional random fields. In *HLT-NAACL*, Edmonton, Canada.
- [31] Shamir, O. (2013). On the complexity of bandit and derivative-free stochastic convex optimization. In *COLT*, Princeton, NJ.
- [32] Smith, D. A. and Eisner, J. (2006). Minimum risk annealing for training log-linear models. In *COLING-ACL*, Sydney, Australia.
- [33] Sokolov, A., Kreutzer, J., Lo, C., and Riezler, S. (2016). Stochastic structured prediction under bandit feedback. In *NIPS*, Barcelona, Spain.
- [34] Spall, J. C. (1992). Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3):332–341.
- [35] Spall, J. C. (2003). *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*. Wiley.
- [36] Sutton, R. S. and Barto, A. G. (2017). *Reinforcement Learning. An Introduction*. The MIT Press, second edition.
- [37] Sutton, R. S., McAllester, D., Singh, S., and Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, Vancouver, Canada.

- [38] Vicente, L. and Custodio, A. (2012). Analysis of direct searches for discontinuous functions. *Mathematical Programming*, 133(1-2):229–325.
- [39] Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 20:229–256.
- [40] Yue, Y. and Joachims, T. (2009). Interactively optimizing information retrieval systems as a dueling bandits problem. In *ICML*, Montreal, Canada.

Supplementary Material

A Theorems and Proofs

Lemma 1.

$$\nabla_w f_\mu(w) = \mathbb{E}_{\bar{u}, x} \left[\frac{F(w + \mu\bar{u}, x) - F(w, x)}{\mu} \bar{u} \right]. \quad (13)$$

Proof.

$$\begin{aligned} \nabla_w f_\mu(w) &= \nabla_w \mathbb{E}_x \left[\int_{\bar{u}} \mathcal{N}(\bar{u}) F(w + \mu\bar{u}, x) d\bar{u} \right] \\ &= \mathbb{E}_x \left[\int_{\bar{u}} \nabla_w \mathcal{N}(\bar{u}) F(w + \mu\bar{u}, x) d\bar{u} \right] \\ &= \mathbb{E}_x \left[\int_{\bar{u}} \nabla_w \frac{1}{\sqrt{(2\pi)^n}} e^{-\frac{1}{2}\|\bar{u}\|^2} F(w + \mu\bar{u}, x) d\bar{u} \right] \\ &= \mathbb{E}_x \left[\int_y \nabla_w \frac{1}{\sqrt{(2\pi)^n}} e^{-\frac{1}{2}\|\frac{y-w}{\mu}\|^2} F(y, x) \cdot \frac{1}{\mu^n} dy \right] \quad \text{where } y = w + \mu\bar{u} \\ &= \mathbb{E}_x \left[\int_y 1 \cdot \frac{y-w}{\mu^2} \frac{1}{\sqrt{(2\pi)^n}} e^{-\frac{1}{2}\|\frac{y-w}{\mu}\|^2} F(y, x) \cdot \frac{1}{\mu^n} dy \right] \\ &= \mathbb{E}_x \left[\int_{\bar{u}} \frac{\bar{u}}{\mu} \frac{1}{\sqrt{(2\pi)^n}} e^{-\frac{1}{2}\|\bar{u}\|^2} F(w + \mu\bar{u}, x) d\bar{u} \right] \\ &= \mathbb{E}_x \left[\int_{\bar{u}} \frac{\bar{u}}{\mu} \mathcal{N}(\bar{u}) F(w + \mu\bar{u}, x) d\bar{u} \right] \\ &= \mathbb{E}_x \left[\int_{\bar{u}} \frac{\bar{u}}{\mu} \mathcal{N}(\bar{u}) F(w + \mu\bar{u}, x) - F(w, x) d\bar{u} \right]. \end{aligned}$$

The last line follows since $\mathbb{E}_{x, \bar{u}} \left[\frac{F(w, x)}{\mu} \bar{u} \right] = \frac{F(w, x)}{\mu} \mathbb{E}_{\bar{u}, x} [\bar{u}] = 0$. \square

Corollary 1.

$$\mathbb{E}_{\bar{u}_N, \mathcal{X}_N} [\|\nabla f_\mu(w_N)\|^2] \leq \epsilon \text{ if } N \geq \mathcal{O}\left(\frac{\bar{n}^3}{\epsilon^2}\right). \quad (14)$$

Proof. Our goal is to bound the terms on the righthandside of Theorem 1 in order to classify Algorithm 1 by the number of iterations necessary to guarantee an accuracy of ϵ for the expected squared norm of the gradient of f_μ . The inequality is repeated here:

$$\frac{1}{S_N} \sum_{k=0}^N h_k \mathbb{E}_{\bar{u}_k, \mathcal{X}_k} [\|\nabla f_\mu(w_k)\|^2] \leq \frac{1}{S_N} \left((f_\mu(w_0) - f^*) + \left(\frac{1}{2} L_1 (\bar{n} + 4)^2 L_0^2 \right) \sum_{k=0}^N h_k^2 \right). \quad (15)$$

We follow [23] in assuming a constant learning rate $h_k := h$, $k \geq 0$.




Figure 2: Average cumulative loss on training data for sparse multiclass text classification. All results are averaged over 3 runs, with mean results in bold lines, 2 standard deviations in filled areas.

We bound the approximation gap between the original function f and the smoothed approximation f_μ by α choosing $\mu \leq \tilde{\mu} = \frac{\alpha}{\bar{n}^{1/2} L_0}$. The latter is possible by applying [23], eq. (18), to our case s.t. $|f_\mu(w) - f(w)| \leq \mu L_0 \bar{n}^{1/2}$.

Furthermore, we apply [23], eq. (22) to our case s.t. $L_1 = \frac{\bar{n}^{1/2}}{\mu} L_0$ where $L_0(F(\cdot, x)) \leq L_0$ for all $F(\cdot, x), x \in \mathcal{X}$.

Let $S_N = \sum_{k=0}^N h_k = (N+1)h$, then the right-hand side of eq. (15) becomes

$$\frac{1}{(N+1)h} \left((f_{\tilde{\mu}}(w_0) - f^*) + \frac{1}{\tilde{\mu}} \bar{n}^{1/2} (\bar{n}+4)^2 L_0^3 (N+1) h^2 \right) \quad (16)$$

$$= \frac{1}{(N+1)h} \left((f_{\tilde{\mu}}(w_0) - f^*) + \frac{\bar{n}^{1/2} L_0}{\alpha} \bar{n}^{1/2} (\bar{n}+4)^2 L_0^3 (N+1) h^2 \right) \quad (17)$$

$$= \frac{f_{\tilde{\mu}}(w_0) - f^*}{(N+1)h} + \frac{h}{\alpha} \bar{n} (\bar{n}+4)^2 L_0^4 \quad (18)$$

$$\leq \frac{L_0 R}{(N+1)h} + \frac{h}{\alpha} \bar{n} (\bar{n}+4)^2 L_0^4. \quad (19)$$

The last inequality follows from Lipschitz continuity s.t. $f_{\tilde{\mu}}(w_0) - f^* \leq L_0 \|w_0 - w^*\|$ and the additional assumption of a bound $\|w_0 - w^*\| \leq R$. Minimizing this upper bound in h , by taking the first derivative and setting it to zero, gives

$$h^* = \left(\frac{\alpha R}{\bar{n} (\bar{n}+4)^2 L_0^3 (N+1)} \right)^{1/2}. \quad (20)$$

Plugging this back into the upper bound gives

$$\frac{L_0 R}{(N+1)h^*} + \frac{h^*}{\alpha} \bar{n} (\bar{n}+4)^2 L_0^4 \quad (21)$$

$$= 2 \left(\frac{\bar{n} (\bar{n}+4)^2 L_0^5 R}{\alpha (N+1)} \right)^{1/2}. \quad (22)$$

Thus, in order to guarantee an ϵ -accuracy on the left-hand side of eq. (15), we need $N \geq \mathcal{O}\left(\frac{\bar{n}^3}{\epsilon^2}\right)$ iterations. \square

B Further Experiments

Sparse Models for Multiclass Text Classification. Multiclass text classification uses a sparse linear feature representation on the Reuters RCV1 dataset [17]. This dataset a standard benchmark for (simplified) structured prediction that has been used in a bandit setup by [9]. The simplified problem uses a binary Δ function indicating incorrect assignment of one out of 4 classes. The data were split into a training set (509,381 documents from original `test_pt [0-2] .dat` files), a development set (19,486 docs: every 8th entry from `test_pt3.dat` and a test set (19,806 docs from `train.dat`). Training for bandit learning is done by cold starting the models from $w_0 = \mathbf{0}$. Meta-parameter settings were determined on development sets for constant learning rate h in the range of 10^0 to 10^{-5} , and exploration parameter μ in the range of 10^{-2} to 10^{-6} .

Following [9], we used documents with exactly one label from the set of labels {CCAT, ECAT, GCAT, MCAT} and converted them to *tfidf* word vectors of dimension 227,903 on the training set. This *tfidf* conversion yields very sparse features with a sparsity pattern of on average 0.5%.

As shown in Figure 2, best convergence behavior for sparse multiclass text classification is obtained by the SFO method which functions as upper bound for the SZO methods. Among SZO methods we see a clear grouping of algorithms with SPARSE perturbation of active features only (dashed curves) and standard SZO methods where ALL parameters are perturbed (dense curves), with a clear advantage in convergence speed for the former. A comparison of the different SZO update rules defined in Section 4 shows a similar ranking to the experiments described above, with updates based on (SPARSE) two-point function evaluation converging fastest.