

---

# GPU-Accelerated Primal Learning for Extremely Fast Large-Scale Classification

---

**John T. Halloran**

Department of Public Health Sciences  
University of California, Davis  
jthalloran@ucdavis.edu

**David M. Rocke**

Department of Public Health Sciences  
University of California, Davis  
dmrocke@ucdavis.edu

## Abstract

One of the most efficient methods to solve  $L_2$ -regularized primal problems, such as logistic regression and linear support vector machine (SVM) classification, is the widely used trust region Newton algorithm, *TRON* [39]. While *TRON* has recently been shown to enjoy substantial speedups on shared-memory multi-core systems [36, 22], exploiting graphical processing units (GPUs) to speed up the method is significantly more difficult, owing to the highly complex and heavily sequential nature of the algorithm. In this work, we show that using judicious GPU-optimization principles, *TRON* training time for different losses and feature representations may be drastically reduced. For sparse feature sets, we show that using GPUs to train logistic regression classifiers in LIBLINEAR is up to an order-of-magnitude faster than solely using multithreading. For dense feature sets—which impose far more stringent memory constraints—we show that GPUs substantially reduce the lengthy SVM learning times required for state-of-the-art proteomics analysis, leading to dramatic improvements over recently proposed speedups. Furthermore, we show how GPU speedups may be mixed with multithreading to enable such speedups when the dataset is too large for GPU memory requirements; on a massive dense proteomics dataset of nearly a quarter-billion data instances, these mixed-architecture speedups reduce SVM analysis time from over half a week to less than a single day while using limited GPU memory.

## 1 Introduction

Over the past decade, GPUs have become valuable computing resources to accelerate the training of popular machine learning models, playing a key role in the widespread use of deep models and the growing ecosystem of deep learning packages [9, 11, 1, 30, 44]. When a training algorithm admits an efficient GPU implementation (such as gradient boosted trees [42], nonlinear kernel learning [8, 50], and primal methods like *L-BFGS* [40] and variants of gradient descent), the speedups gained using GPUs, as opposed to only CPUs, are often substantial. For instance, in PyTorch [44], training a logistic regression classifier on the `rcv1` [37] dataset with gradient descent is 14.6 times faster using a Tesla V100 GPU versus using 24 CPU threads with an Intel Xeon Gold 5118 (similarly, training with *L-BFGS* in this example is 13.1 times faster using the V100, detailed in [23]).

Specialized solvers commonly provide even more speed. For the previous logistic regression example, using just a single CPU thread with scikit-learn’s [45] *TRON* solver—the primal learning algorithm for logistic regression and SVM classification/regression adapted from LIBLINEAR [13]—is 94.7 and 10.9 times faster than GPU-accelerated gradient descent and *L-BFGS*, respectively, implemented in PyTorch. However, while significant work has been done to further accelerate *TRON* and many other extremely fast machine learning solvers [27, 54, 35, 31, 32] using multiple CPU cores [6, 29, 28, 36, 10, 46, 55, 22], analogous GPU speedups for such efficient algorithms are often lacking.

This lack of GPU exploitation is due to the specialized structure and complexity of these algorithms, which **naturally lend themselves to multithreaded speedups** on shared memory systems, **yet resist optimizations on GPU architectures**.

For example, *TRON* relies on random access to features for SVM losses, which is naturally supported in multithreaded systems, but prevents memory coalescing (and is thus deleterious) for GPU computation. Furthermore, large memory transfers between the CPU and GPU are expensive, so that the complex, sequential dependency of variables in specialized algorithms like *TRON* make optimal GPU use difficult. Indeed, we show that while most of the computational bottlenecks for logistic regression in *TRON* are linear algebra operations [36] (for which GPUs greatly outperform CPUs), using *TRON*'s original variable access pattern in LIBLINEAR results in poor GPU performance—performing even worse than using only a single CPU thread on one of the presented datasets.

Herein, we show that using just a single GPU, excellent training speedups are achievable by overly CPU-specialized machine learning algorithms such as *TRON*. In particular, for different feature representations and loss functions, we show that *TRON* training times may be drastically reduced using judicious GPU-optimization principles.

**Sparse Features.** For sparse feature representations, we successively optimize *TRON* for logistic regression (referred to as *TRON*-LR) in LIBLINEAR using several strategies to: a) decouple the sequential dependence of variables, b) minimize the number of large-memory transfers between GPU and CPU, and c) maximize parallelism between the CPU and GPU. We show that while *TRON*'s original variable access pattern limits the effectiveness of GPU computation, using a single CPU thread with our GPU optimizations results in a 70.8% improvement in training time (averaged over the presented datasets) over the single-threaded version of *TRON* in standard LIBLINEAR. In addition, we show that mixing our GPU optimizations with multithreading provides further speedups, resulting in an average **89.2% improvement over single-thread optimized *TRON*** and an average **65.2% improvement over *TRON* in the multithread-optimized version of LIBLINEAR** [36].

**Dense Features.** For dense feature representations, we show that GPUs substantially reduce SVM learning times for state-of-the-art analysis of dense proteomics datasets [33]. Overcoming the random access restrictions of *TRON* SVM learning (referred to as *TRON*-SVM), we show that using just a single GPU leads to an average as much as triples the performance of recently proposed speedups for this application [22]. **On a large-scale dataset of over 23 million data instances, these GPU speedups reduce SVM learning time from 14.4 hours down to just 1.9 hours.** Furthermore, dense feature sets impose stringent GPU memory constraints, particularly for the massive datasets regularly produced in biological experiments. Thus, we demonstrate how GPU optimizations may be mixed with multithreading to significantly reduce GPU memory constraints. On a massive proteomics dataset consisting of over 215 million data instances—which exceeds memory requirements for GPU-only speedups—these mixed-architecture speedups drastically outperform recent multithread-optimized solvers, **reducing standard analysis time from 4.4 days down to just 19.7 hours.**

The paper is organized as follows. In Section 2, we describe relevant previous work speeding up *TRON* for both sparse and dense feature sets on shared memory systems. In Section 3, we define the general *TRON* algorithm and computational bottlenecks encountered minimizing different loss functions. In Sections 4 and 5, we discuss how the computational bottlenecks in algorithms like *TRON* natively resist GPU speedups, and GPU-optimization principles to overcome these hurdles (providing the resulting GPU optimizations for the objectives and feature architectures under study). We demonstrate that the presented GPU-optimizations drastically outperform recent multithreaded speedups in Section 6, and conclude with future avenues extending the presented work to other high-performance GPU packages (such as PyTorch) in Section 7.

## 2 Previous Work

Serving as the primal solver in the popular package LIBLINEAR [13], *TRON* has been extensively tested and shown to enjoy superior speed and convergence compared to other second-order solvers, such as the widely-used quasi-Newton algorithm *L-BFGS* [40] and the modified Newton algorithm *L2-SVM-MFN* [35] (one of the fastest algorithms for large-scale primal SVM learning). As a Newton method, the algorithm enjoys general quadratic convergence without loading the entire Hessian into memory, thus only using linear memory. For logistic and SVM losses in LIBLINEAR, *TRON*'s convergence speed has further been theoretically improved by refining trust-region update rules [26] and applying a preconditioner matrix to help stabilize optimization [25]. In [36], multithreaded

optimizations in shared-memory multi-core systems were extensively explored to speed up *TRON*'s computational bottlenecks (further described in Section 3) for logistic regression. Evaluating several multithreading libraries (i.e., OpenMP, Intel's Math Kernel Library, and the sparse matrix multiplication package *librsb*) over a large number of datasets, OpenMP was found to provide the best multithreaded performance and was subsequently integrated into the multi-core release of LIBLINEAR.

## 2.1 SVM Classification Using *TRON* for Fast Large-Scale Proteomics Analysis

In proteomic analysis pipelines, SVM classification using Percolator [33] is a critical step towards accurately analyzing protein data collected via tandem mass spectrometry (MS/MS). Given a collection of MS/MS spectra representing the protein subsequences (called *peptides*) present in a biological sample, the first stage of proteomics analysis typically consists of identifying the input spectra by searching (i.e., scoring and ranking) a database of peptides. This first stage thus results in a list of *peptide-spectrum-matches* (PSMs) and their respective scores. In practice, however, database-search scoring functions are often poorly calibrated, making PSMs from different spectra difficult to compare and diminishing overall identification accuracy. To correct for this, the list of PSMs, as well as dense feature vectors describing each match, are fed into Percolator for *recalibration*.

Percolator first estimates PSM labels (i.e., correct versus incorrect) using false discovery rate analysis [34], then trains a linear SVM to classify correct and incorrect identifications. These two steps are iterated until convergence and the input PSM scores are subsequently recalibrated using the final learned SVM parameters. Furthermore, to prevent overfitting and improve generalizability within each iteration, three-fold cross-validation is carried out over three disjoint partitions of the original dataset, followed by further nested cross-validation within each fold [16].

The accuracy improvements of Percolator recalibration have been well demonstrated for a wide variety of PSM scoring functions—e.g., linear [33, 7, 52], *p*-value based [15, 24, 38], and dynamic Bayesian networks [18, 19, 17]—and complex PSM feature sets—e.g., Fisher kernels [20, 21], subscores of linear functions [47], ensembles of scoring functions [49], and features derived using deep models [14]. However, due to the iterative training of many SVMs during cross-validation, Percolator requires substantial analysis times for large-scale datasets commonly collected in MS/MS experiments. Initial work sought to speed up Percolator runtimes by randomly sampling a small portion of the data to train over [41], but this was subsequently shown to unpredictably diminish the performance of learned parameters [22]. Thus, to combat these lengthy analysis times without affecting learned SVM parameters, recent work [22] applied extensive systems-level speedups and multithreading in both Percolator's original primal solver, *L2-SVM-MFN*, and *TRON* (heavily optimized to utilize dense feature vectors). While both optimized solvers were shown to significantly improve Percolator training times for large-scale data, *TRON* displayed markedly superior performance.

## 3 Trust Region Newton Methods for Primal Classification

Consider feature vectors  $\mathbf{x}_i \in \mathbb{R}^n, i = 1, \dots, l$  and label vector  $\mathbf{y} \in \{-1, 1\}^l$ , and let  $X = [\mathbf{x}_1 \dots \mathbf{x}_l]^T$  be the feature matrix. For vectors, index-set subscripts denote subvectors and for matrices, pairs of index-set subscripts denote submatrices. The general objective, which we wish to minimize w.r.t.  $\mathbf{w}$ , is

$$f(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|_2^2 + C \sum_{i=1}^l \ell(\mathbf{w}; \mathbf{x}_i, y_i), \tag{1}$$

where  $\frac{1}{2} \|\mathbf{w}\|_2^2$  is the regularization term,  $C > 0$  is a regularization hyperparameter, and  $\ell(\mathbf{w}; \mathbf{x}_i, y_i)$  is a loss function.

When  $\ell(\mathbf{w}; \mathbf{x}_i, y_i) = \log(1 + \exp(-y_i \mathbf{w}^T \mathbf{x}_i))$ , commonly referred to as the *logistic loss*, minimizing Equation 1 corresponds to learning a classifier using logistic regression. Similarly, minimizing Equation 1 when  $\ell(\mathbf{w}; \mathbf{x}_i, y_i) = (\max(0, 1 - y_i \mathbf{w}^T \mathbf{x}_i))^2$ , commonly referred to as the *quadratic SVM* or *L2-SVM loss*, corresponds to learning a linear SVM classifier. We denote Equation 1 under the logistic loss as  $f_{LR}(\mathbf{w})$  and, under the L2-SVM loss, as  $f_{L2}(\mathbf{w})$ .

---

**Algorithm 1** The *TRON* algorithm
 

---

```

1: Given  $w$ ,  $\Delta$ , and  $\sigma_0$ 
2: Calculate  $f(w)$  // Critically depends on  $z = X^T w$ 
3: while Not converged do
4:   Find  $d = \operatorname{argmin}_v q(v)$  s.t.  $\|v\|_2 \leq \Delta$ . // Critically depends on  $\nabla f(w)$ ,  $\nabla^2 f(w)v$ 
5:   Calculate  $f(w + d)$ ,  $\sigma = \frac{f(w+d) - f(w)}{q(d)}$  // Critically depends on  $z = X^T (w + d)$ 
6:   if  $\sigma > \sigma_0$  then
7:      $w \leftarrow w + d$ , increase trust region  $\Delta$ .
8:   else
9:     Shrink  $\Delta$ .
10:  end if
11: end while

```

---

*TRON* is detailed in Algorithm 1. At each iteration, given the current parameters  $w$  and trust region interval  $\Delta$ , *TRON* considers the following quadratic approximation between function parameters,

$$f(w + d) - f(w) \approx q(d) \equiv \nabla f(w)^T d + \frac{1}{2} d^T \nabla^2 f(w) d. \quad (2)$$

A truncated Newton step ( $d$  on line 4 in Algorithm 1), confined in the trust region, is then found using a conjugate gradient procedure. If  $q(d)$  is close to  $f(w + d) - f(w)$ ,  $w$  is updated to  $w + d$  and the trust region interval is increased for the subsequent iteration. Otherwise,  $w$  remains unchanged and the trust region interval is shrunk.

Note that the function evaluation  $f(w)$ —which critically depends on computing  $z = X^T w$  for both losses—must be computed for each new iteration, as well as the gradient and Hessian for Equation 2. However, computing only the Hessian-vector product in Equation 2 avoids loading the entire Hessian into memory (which would be intractable for large datasets). Thus, the most intensive portions of *TRON* are the computations of  $z = X^T w$ ,  $\nabla f(w)$ , and  $\nabla^2 f(w)v$  (where  $v$  is the optimization variable in line 4 of Algorithm 1), summarized for both losses in Table 1. Further derivation of these quantities is available in [23].

We note that arbitrary loss functions (and combinations thereof) may be used in Equation 1, thus allowing future work utilizing the highly efficient *TRON* in popular automatic differentiation [5] packages [44, 1, 48, 43]. However, these packages rely on GPUs for optimal performance, the use of which *TRON* natively resists (as we’ll see, and rectify, for the two loss functions considered).

Logistic Loss	L2-SVM Loss
$z = X^T w$ , to compute $f_{\text{LR}}(w)$	$z = X^T w$ , to compute $f_{\text{L2}}(w)$
$\nabla f_{\text{LR}}(w) = w + C \sum_{i=1}^l (h(y_i z_i) - 1) y_i x_i$ , where $h(y_i z_i) = (1 + e^{-y_i z_i})^{-1}$	$\nabla f_{\text{L2}}(w) = w + 2C X_{I,:}^T (z_I - y_I)$ , where $I \equiv \{i   1 - y_i z_i > 0\}$ is an index set and the operator $:$ denotes all elements along the corresponding dimension (i.e., all columns in this case)
$\nabla^2 f_{\text{LR}}(w)v = v + CX^T(D(Xv))$ , where $D$ is a diagonal matrix with elements $D_{i,i} = h(y_i z_i)(1 - h(y_i z_i))$	$\nabla^2 f_{\text{L2}}(w)v = v + 2CX_{I,:}^T(X_{I,:}v)$

Table 1: *TRON* major bottleneck computations for logistic and L2-SVM losses.

## 4 Accelerating *TRON*-LR training using GPUs

Assume a shared-memory multi-core system and a single GPU with sufficient memory for the variables in Table 1 (this is later relaxed in Section 5). Herein, the CPU is referred to as the *host* and the GPU is referred to as the *device*.

*TRON*-LR runtime is dominated by three major matrix-vector multiplications in the bottleneck computations listed in Table 1:  $z = X^T w$ ,  $\nabla^2 f_{\text{LR}}(w)v$ , and  $\nabla f_{\text{LR}}(w) = w + CX\hat{z}$ , where  $\hat{z}_i = (h(y_i z_i) - 1)y_i$ . For instance, profiling *TRON*-LR in LIBLINEAR training on the large-scale

SUSY [4] dataset, these **three matrix-vector multiplications account for 82.3% of total training time**. We thus first attempt to accelerate *TRON*-LR by computing these quantities quickly on the device (as was similarly done in [36] using multithreading).

In LIBLINEAR, this first attempt at GPU acceleration (called *TRON*-LR-GPU<sup>0</sup>) is implemented using cuSPARSE to perform sparse linear algebra operations as efficiently as possible for LIBLINEAR’s sparse feature representation. Compared to the standard single-threaded implementation of LIBLINEAR on the SUSY dataset, *TRON*-LR-GPU<sup>0</sup> achieves a speedup of 0.65—*TRON*-LR-GPU<sup>0</sup> is actually slower than the single-threaded LIBLINEAR! *TRON*-LR-GPU<sup>0</sup> fares better on other presented datasets, but performs poorly overall (displayed in Figure 1).

#### 4.1 Sequentially Dependent Variables

The critical issue encountered by *TRON*-LR-GPU<sup>0</sup> is *TRON*’s overly sequential dependency of variables; once variable vectors are computed on the GPU, they are immediately needed on the host CPU to proceed with the next step of the algorithm. For instance, computing the bottleneck  $z = X^T(\mathbf{w} + \mathbf{d})$  using cuSPARSE is fast, but  $z$  must immediately be transferred back to the host to compute  $f_{\text{LR}}(\mathbf{w} + \mathbf{d})$  (in line 5 of Algorithm 1). However, large-memory transfers between the host and device are expensive, especially when either the host or device are waiting idle for the transaction to complete (steps to conceal transfer latency are discussed in [23]).

Furthermore, all other major operations in Algorithm 1 are locked in the same manner as the previous bottleneck example: the trust region update (lines 6-10) can not proceed without the value of  $f_{\text{LR}}(\mathbf{w} + \mathbf{d})$ , and, without either the updated  $\mathbf{w}$  or trust region, operations for the next iteration’s truncated Newton step (line 4) are unable to run concurrently in an attempt to conceal transfer latency. Clearly, this pattern of variable access is suboptimal for GPU use (best evidenced by *TRON*-LR-GPU<sup>0</sup>’s performance in Section 6).

#### 4.2 Decoupling Dependencies to Maximize Host and Device Parallelism

To optimally use the GPU, we must first decouple the sequential dependency of variables discussed in Section 4.1. Recall that, for  $\nabla f_{\text{LR}}(\mathbf{w})$ , the vector  $\hat{z}$  is such that  $\hat{z}_i = (h(y_i z_i) - 1)y_i$ . To decrease sequential dependencies on the computational bottleneck  $z = X^T(\mathbf{w} + \mathbf{d})$ , we first note that calculation of  $f_{\text{LR}}(\mathbf{w} + \mathbf{d})$  always precedes  $\nabla f_{\text{LR}}(\mathbf{w} + \mathbf{d})$ . Thus, to decouple gradient variables, once  $z$  is calculated on the device, we prepare all device-side variables needed to compute  $X^T \hat{z}$  in the event that  $\sigma > \sigma_0$ . Specifically, after computing  $z = X^T(\mathbf{w} + \mathbf{d})$  on the device, we use a custom CUDA kernel to compute  $\hat{z}$  followed by a Thrust reduction to compute  $f_{\text{LR}}(\mathbf{w} + \mathbf{d}) = \frac{1}{2}\|\mathbf{w} + \mathbf{d}\|_2^2 + C \sum_{i=1}^l \log(1 + e^{-y_i z_i})$  (note that the scalar output of the reduction, i.e.  $f_{\text{LR}}(\mathbf{w} + \mathbf{d})$ , is immediately available to the host). The computation of  $\hat{z}$  is massively parallelizable, so the grid-stride loop in the custom kernel is extremely efficient. Thus, if  $\sigma > \sigma_0$ , the variable  $\hat{z}$  is already in device memory and the gradient is quickly calculated using cuSPARSE on the device as  $\nabla f_{\text{LR}}(\mathbf{w} + \mathbf{d}) = \mathbf{w} + \mathbf{d} + X^T \hat{z}$ . Finally,  $\nabla f_{\text{LR}}(\mathbf{w} + \mathbf{d})$  is transferred from device to host, which is notably efficient when  $l \gg n$  (i.e., the optimal setting for primal learning).

This set of operations accomplishes several optimizations simultaneously:

- **Decoupling dependencies, avoiding large transfers:**  $z$  and  $\hat{z}$  are completely decoupled of any dependency for host-side computation, thanks to the custom reduction and kernel. This saves several large transfers of  $z$ ,  $\hat{z}$  from (and to) the device, and avoids the need to conceal transfer latency.
- **Coalesced memory:** the device performs optimally as all operations allow memory coalescing.
- **Device saturation:** an uninterrupted series of intensive computation is performed on the device (i.e., no device-side stalls due to host dependencies).
- **Host and device parallelism:** the complete decoupling of  $z$ ,  $\hat{z}$  allows more independent operations to be run on the host while the device runs concurrently.

We complete the total GPU optimization of *TRON*-LR by speeding up the remaining bottleneck, the Hessian-vector product  $\nabla^2 f_{\text{LR}}(\mathbf{w})v$ . As with the previous optimizations, device variables are maximally decoupled from host-side dependencies, while using device-side functions which allow

<i>TRON-SVM-GPU</i>	<i>TRON-SVM-MIX</i>
$\mathbf{z} = X\mathbf{w}$ is calculated and stored on the device.	$\mathbf{z} = X\mathbf{w}$ is calculated and stored on the device, then transferred to the host.
$I = \{i : y_i z_i < 1\}$ is calculated on the device, then $f_{L2}(\mathbf{w}) = \frac{1}{2}\ \mathbf{w}\ _2^2 + C \sum_{i=1}^l (1 - y_i z_i > 0)^2$ is computed on the device while the host runs independent, sequential operations.	$I$ is calculated on the device, then transferred to the host. The device-side computation of $f_{L2}(\mathbf{w})$ is run concurrently with this transfer.
On the device, $\hat{\mathbf{z}} = (\mathbf{z}_I - \mathbf{y}_I)$ and $\hat{X} = X_{I,:}$ are computed. The gradient $\nabla f_{L2}(\mathbf{w}) = \mathbf{w} + 2C\hat{X}^T \hat{\mathbf{z}}$ is then computed and transferred to the host.	With $I$ and $\mathbf{z}$ on the host, $\nabla f_{L2}(\mathbf{w})$ is computed using multithreading.
The Hessian-product is computed on the device as $\nabla^2 f_{L2}(\mathbf{w})\mathbf{v} = \mathbf{v} + 2C\hat{X}^T(\hat{X}\mathbf{v})$ and transferred to the host.	Using multithreading, the Hessian-product is calculated on the host as $\nabla^2 f_{L2}(\mathbf{w})\mathbf{v} = \mathbf{v} + 2CX_{I,:}^T(X_{I,:}\mathbf{v}) = \mathbf{v} + 2C \sum_{i \in I} (\mathbf{x}_i^T \mathbf{v}) \mathbf{x}_i$ .

Table 2: Major operations of the *TRON-SVM* solvers designed for GPU compute.

peak performance. Further details are in [23], including a comprehensive summary of the previously described *TRON-LR* GPU-optimizations.

**Decreasing runtimes via mixed-architecture speedups.** While computations remain which may be accelerated using the same GPU-optimization principles, allocating additional device vectors becomes problematic for large-scale datasets and current GPU memory ranges. Thus, in addition to the previously described GPU optimizations, we accelerate remaining bottleneck areas using multithreading. In particular, multithreading using *OpenMP* is used to accelerate vector-matrix-vector multiplications in the conjugate gradient procedure (previously optimized using loop unrolling) and application of the preconditioner matrix [25] (which is jointly sped up using existing device-side computation during the Hessian-vector product optimizations).

## 5 Accelerating *TRON-SVM* training using GPUs

Focusing on speeding up SVM learning in the state-of-the-art software *Percolator* [33]—which uses dense feature vectors to analyze large-scale proteomics datasets—the GPU-optimization principles from Section 4.2 are applied to *TRON-SVM*: device-side variables are decoupled from dependent host-side computations, necessary transfers are run asynchronously in parallel with the maximum number of host/device operations, and linear algebra operations are extensively and efficiently carried out using *cuBLAS*. However, speed ups in *TRON-SVM* possess a key difficulty for GPU computation; for  $\mathbf{z} = X^T \mathbf{w}$ , the active set  $I \equiv \{i | 1 - y_i z_i > 0\}$  is recomputed every iteration. Thus, computation of both  $\nabla f_{L2}(\mathbf{w})$  and  $\nabla^2 f_{L2}(\mathbf{w})$  require the submatrix  $X_{I,:}$  (as seen in Table 1).

While accessing  $X_{I,:}$  is naturally supported through shared-memory random access for multithreaded speedups (e.g., in [22], which used *OpenMP* to speed up the *TRON-SVM* quantities in Table 1 within *Percolator*), the non-contiguous nature of this operation leads to misaligned (i.e., not coalesced) device memory which prevents optimal GPU use. Furthermore, as noted in Section 4.2, large memory transfers between host and device are expensive, hindering approaches where  $I$  is first computed then a randomly accessed submatrix is created on the host and transferred to the device.

To overcome this challenge, we first make use of the insight that, prior to computing  $f_{L2}(\mathbf{w})$ , the active set  $I$  may be computed and stored entirely on the device. With  $I$  on the device, the submatrix  $X_{I,:}$  may be efficiently computed *within device memory*. Computing  $I$  and  $X_{I,:}$  on the device entirely decouples these variables from host-side compute and accomplishes all simultaneous optimizations listed in Section 4.2. The major operations of the resulting GPU-optimized solver, called *TRON-SVM-GPU*, are listed in Table 2.

**Decreasing GPU-memory utilization via mixed-architecture speedups.** In *TRON-SVM-GPU*, the device memory required to decouple  $X_{I,:}$  from host-side compute proves prohibitive for extremely large-scale proteomics datasets. To remedy this, the mixed-architecture solver, *TRON-SVM-MIX*, utilizes the GPU for heavy lifting before using multithreading for efficient random access to  $X_{I,:}$  (after  $I$  is computed) during *TRON*’s conjugate gradient procedure. Thus, *TRON-SVM-MIX* utilizes much less GPU memory than *TRON-SVM-GPU*, at the expense of some speed due to fewer operations being run on the device. The major operations of *TRON-SVM-MIX* are listed in Table 2.




Figure 1: Factor of speedup for several optimized versions of *TRON-LR* in *LIBLINEAR*. The  $x$ -axis displays the number of threads used for multithreaded methods. The  $y$ -axis denotes the multiplicative factor of training speedup for each method relative to the single-threaded version of *TRON-LR* in the standard *LIBLINEAR*. Training times were measured within *LIBLINEAR* as the time elapsed calling `tron_obj.tron()`. As is standard practice, wallclock times were measured as the minimum reported times over ten runs.


## 6 Results and Discussion

All experiments were run on a dual Intel Xeon Gold 5118 compute node with 48 computational threads, an NVIDIA Tesla V100 GPU, and 768 GB of memory.

**Speedups for sparse features.** The *TRON-LR* GPU-optimized and mixed-architecture solvers (described in Section 4.2) are referred to as *TRON-LR-GPU* and *TRON-LR-MIX*, respectively. *TRON-LR-GPU*, *TRON-LR-MIX*, and *TRON-LR-GPU<sup>0</sup>* were all developed based on *LIBLINEAR* v2.30. Single-threaded *LIBLINEAR* tests were run using v2.30. The multithread-optimized version of *TRON-LR* described in [36], referred to herein as *TRON-LR-CPU*, was tested using multi-core *LIBLINEAR* v2.30. All single-threaded *TRON-LR* implementations (i.e., *TRON-LR-GPU<sup>0</sup>*, *TRON-LR-GPU*, and the single-threaded optimized version of *TRON-LR* in standard *LIBLINEAR*) were run with the same command line parameters: `-c 4 -e 0.1 -s 0`. Multithreaded implementations were run with the additional flag `-nr i`, specifying the use of  $i$  compute threads. As is standard practice, wallclock times were measured as the minimum reported times over ten runs. Training times were measured within *LIBLINEAR* as the time elapsed calling `tron_obj.tron()`. Six datasets of varying statistics (i.e., number of features, instances, and nonzero elements) were downloaded from <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/> and used to benchmark the *TRON-LR* solvers (statistics for each dataset are listed in [23]).

The speedups for all methods are displayed in Figure 1. *TRON-LR-GPU* significantly outperforms the multithread-optimized *TRON-LR-CPU* and the direct GPU implementation, *TRON-LR-GPU<sup>0</sup>*, across all datasets and threads. The mixed-architecture *TRON-LR-MIX* further improves upon *TRON-LR-GPU* performance in each dataset for all threads used, leading to over tenfold speedups in training time on half of the presented datasets. We note that, due to thread scheduling overhead, multithreaded methods experience diminished performance for large numbers of threads in Figures 1a, 1b, 1f. However, the presented GPU optimizations consistently provide the best speedups when both multithreading is not used and when multithreading is overutilized.

**Speedups for dense features.** The *TRON* GPU solvers described in Section 5—the GPU-optimized *TRON-SVM-GPU* and the mixed-architecture *TRON-SVM-MIX*—are compared against the multithread-optimized versions of *TRON* (referred to as *TRON-SVM-CPU*) and *L2-SVM-MFN* from [22]. The methods are tested using two extremely large datasets. The first dataset, referred to as the Kim dataset, is a larger version of the benchmark dataset used in [22], consisting of 23,330,311 PSMs (i.e., proteomics data instances, described in 2.1). The second dataset, referred to as the



(a) SVM speedups for a large, dense proteomics dataset containing 23,330,311 PSMs.

(b) SVM speedups for a massive dense dataset containing 215,282,771 PSMs, too large to be analyzed by “TRON-SVM-GPU.”

Figure 2: Factor of speedup for SVM learning in Percolator for dense large- and massive-scale datasets. Speedup factor is calculated as the original Percolator SVM learning time divided by the speed up learning time. The x-axis displays the number of threads utilized by multithreaded methods “L2-SVM-MFN,” “TRON-SVM-CPU,” and “TRON-SVM-MIX.”

Wilhelm dataset, was collected from a map of the human proteome [51] and contains 215,282,771 PSMs. All multithreaded solvers were tested using 8, 16, 24, 32, 40, and 48 threads. As in [22], to effectively measure the runtime of multithreaded methods without excess thread-scheduling overhead, parallelization of Percolator’s outermost cross-validation was disabled.

Reported runtimes are the minimum wall-clock times measured over five runs for the Kim dataset and three runs for the Wilhelm dataset. The original Percolator SVM learning runtimes (collected using Percolator v3.04.0) were 14.4 hours and 4.4 days for the Kim and Wilhelm datasets, respectively. Speedups for both datasets are illustrated in Figure 2. For the Kim dataset, speedup results for all discussed methods are illustrated in Figure 2a. For the Wilhelm dataset, total Tesla V100 memory (16 GB) is exceeded for TRON-SVM-GPU. However, the reduced memory requirements of TRON-SVM-MIX allow GPU speedups for this massive dataset (illustrated in Figure 2b).

Both GPU solvers greatly accelerate Percolator SVM learning while dominating previously proposed multithreaded speedups. For the Kim dataset, TRON-SVM-MIX and TRON-SVM-GPU achieve 6.6 and 7.4 fold speedups, respectively, over Percolator’s current SVM learning engine. For the Wilhelm dataset, TRON-SVM-MIX achieves a 5.4 fold speedup while being notably efficient using few system threads—with at most 16 threads, TRON-SVM-MIX improves the average training time of TRON-SVM-CPU and L2-SVM-MFN by 50% and 70%, respectively. Together, these two solvers present versatile trade-offs for different compute environments; when the dataset does not exceed the GPU memory, TRON-SVM-GPU offers superior performance. However, when onboard GPU memory is limited, a small portion of speed may be traded for much less memory consumption by using TRON-SVM-MIX. Furthermore, when the number of computational threads is also limited, TRON-SVM-MIX offers significantly better (and more stable) performance at low numbers of utilized threads compared to the purely multithreaded solvers TRON-SVM-CPU and L2-SVM-MFN.

## 7 Conclusions and Future Work

In this work, we’ve shown that by using general GPU-optimization principles, excellent speedups may be enjoyed by algorithms which natively resist GPU optimization. For the widely used TRON algorithm, we’ve presented several GPU-optimized solvers for both sparse and dense feature sets of  $L_2$ -regularized primal problems. Using a single GPU, these solvers were shown to dominate recently proposed speedups for logistic regression (within LIBLINEAR) and SVM classification for state-of-the-art proteomics analysis (within Percolator). Furthermore, for sparse features, we’ve shown how multithreading may compliment GPU optimizations and, for memory-restrictive dense features, how multithreading may relieve device-memory requirements while allowing substantial GPU speedups. The former optimizations achieve over an order-of-magnitude speedup on half of the



presented datasets (and an average 9.3 fold speedup on all datasets), while the latter optimizations decrease massive-scale biological analysis time from 4.4 days down to just 19.7 hours.

There are significant avenues for future work. We plan to extend GPU-optimized *TRON* implementations to use general gradient and Hessian-vector product information computed in automatic differentiation [5] packages such as PyTorch [44] and TensorFlow [1], which utilize second-order primal solvers (such as *L-BFGS*) to optimize losses while relying on GPU compute for optimal performance. Furthermore, we plan to apply the presented GPU-optimization principles to speed up other fast machine learning solvers [27, 2, 54, 35, 31, 32] which, like *TRON*, are natively designed to rely on sequential dependencies of variables.

## Broader Impact

This paper solely focuses on speeding up machine learning software, and thus impacts machine learning packages or applications which use either the included software or the paper’s GPU optimization principles (to speed up an algorithm not discussed). Benefits include faster software, with specific applications including real-time classification for self-driving cars [53], flagging credit card fraud [12], water monitoring to preserve ecosystems in maritime and archipelagic countries [3]), etc. Machine learning companies/researchers/practitioners who do not use GPU resources may be put at a disadvantage from this research, but any advantage/disadvantage is defined solely in terms of training time speed.

**Acknowledgments:** This work was supported by the National Center for Advancing Translational Sciences (NCATS), National Institutes of Health, through grant UL1 TR001860 and a GPU donation from the NVIDIA Corporation.

## References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Galen Andrew and Jianfeng Gao. Scalable training of  $l_1$ -regularized log-linear models. In *Proceedings of the 24th international conference on Machine learning*, pages 33–40, 2007.
- [3] Riyadh Arridha, Sritrusta Sukaridhoto, Dadet Pramadihanto, and Nobuo Funabiki. Classification extension based on iot-big data analytic for smart environment monitoring and analytic in real-time system. *International Journal of Space-Based and Situated Computing*, 7(2):82–93, 2017.
- [4] Pierre Baldi, Peter Sadowski, and Daniel Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5:4308, 2014.
- [5] Atılım Günes Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *The Journal of Machine Learning Research*, 18(1):5595–5637, 2017.
- [6] Joseph K Bradley, Aapo Kyrola, Danny Bickson, and Carlos Guestrin. Parallel coordinate descent for  $l_1$ -regularized loss minimization. In *Proceedings of the 28th international conference on Machine learning*, 2011.
- [7] M. Brosch, L. Yu, T. Hubbard, and J. Choudhary. Accurate and sensitive peptide identification with Mascot Percolator. *J. Proteome Res.*, 8(6):3176–3181, 2009.
- [8] Bryan Catanzaro, Narayanan Sundaram, and Kurt Keutzer. Fast support vector machine training and classification on graphics processors. In *Proceedings of the 25th international conference on Machine learning*, pages 104–111, 2008.

- [9] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*, 2015.
- [10] Wei-Lin Chiang, Mu-Chu Lee, and Chih-Jen Lin. Parallel dual coordinate descent method for large-scale linear classification in multi-core environments. In *KDD*, pages 1485–1494, 2016.
- [11] François Chollet et al. Keras. <https://keras.io>, 2015.
- [12] Sahil Dhankhad, Emad Mohammed, and Behrouz Far. Supervised machine learning algorithms for credit card fraudulent transaction detection: a comparative study. In *2018 IEEE International Conference on Information Reuse and Integration (IRI)*, pages 122–125. IEEE, 2018.
- [13] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. Liblinear: A library for large linear classification. *Journal of machine learning research*, 9(Aug):1871–1874, 2008.
- [14] Siegfried Gessulat, Tobias Schmidt, Daniel Paul Zolg, Patroklos Samaras, Karsten Schnatbaum, Johannes Zerweck, Tobias Knaute, Julia Rechenberger, Bernard Delanghe, Andreas Huhmer, et al. Prosit: proteome-wide prediction of peptide tandem mass spectra by deep learning. *Nature methods*, 16(6):509, 2019.
- [15] Viktor Granholm, Sangtae Kim, Jose CF Navarro, Erik Sjolund, Richard D Smith, and Lukas Kall. Fast and accurate database searches with ms-gf+ percolator. *J. Proteome Res.*, pages 890–897, 2013.
- [16] Viktor Granholm, William Stafford Noble, and Lukas Kall. A cross-validation scheme for machine learning algorithms in shotgun proteomics. *BMC bioinformatics*, 13(16):S3, 2012.
- [17] John T Halloran. Analyzing tandem mass spectra using the drip toolkit: Training, searching, and post-processing. In *Data Mining for Systems Biology*, pages 163–180. Springer, 2018.
- [18] John T. Halloran, Jeff A. Bilmes, and William S. Noble. Learning peptide-spectrum alignment models for tandem mass spectrometry. In *Uncertainty in Artificial Intelligence (UAI)*, Quebec City, Quebec Canada, July 2014. AUAI.
- [19] John T Halloran, Jeff A Bilmes, and William S Noble. Dynamic bayesian network for accurate detection of peptides from tandem mass spectra. *Journal of Proteome Research*, 15(8):2749–2759, 2016.
- [20] John T Halloran and David M Rocke. Gradients of generative models for improved discriminative analysis of tandem mass spectra. In *Advances in Neural Information Processing Systems*, pages 5728–5737, 2017.
- [21] John T Halloran and David M Rocke. Learning concave conditional likelihood models for improved analysis of tandem mass spectra. In *Advances in Neural Information Processing Systems*, pages 5420–5430, 2018.
- [22] John T Halloran and David M Rocke. A matter of time: Faster percolator analysis via efficient svm learning for large-scale proteomics. *Journal of proteome research*, 17(5):1978–1982, 2018.
- [23] John T Halloran and David M Rocke. Gpu-accelerated primal learning for extremely fast large-scale classification: Supplementary material. In *Advances in Neural Information Processing Systems*, 2020.
- [24] J Jeffrey Howbert and William S Noble. Computing exact p-values for a cross-correlation shotgun proteomics score function. *Molecular & Cellular Proteomics*, pages mcp–O113, 2014.
- [25] Chih-Yang Hsia, Wei-Lin Chiang, and Chih-Jen Lin. Preconditioned conjugate gradient methods in truncated newton frameworks for large-scale linear classification. In *Asian Conference on Machine Learning*, pages 312–326, 2018.
- [26] Chih-Yang Hsia, Ya Zhu, and Chih-Jen Lin. A study on trust region update rules in newton methods for large-scale linear classification. In *Asian Conference on Machine Learning*, pages 33–48, 2017.

- [27] Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S Sathya Keerthi, and Sellamanickam Sundararajan. A dual coordinate descent method for large-scale linear svm. In *Proceedings of the 25th international conference on Machine learning*, pages 408–415, 2008.
- [28] Cho-Jui Hsieh, Hsiang-Fu Yu, and Inderjit S Dhillon. Passcode: Parallel asynchronous stochastic dual co-ordinate descent. In *ICML*, volume 15, pages 2370–2379, 2015.
- [29] Martin Jaggi, Virginia Smith, Martin Takác, Jonathan Terhorst, Sanjay Krishnan, Thomas Hofmann, and Michael I Jordan. Communication-efficient distributed dual coordinate ascent. In *Advances in neural information processing systems*, pages 3068–3076, 2014.
- [30] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [31] Tyler Johnson and Carlos Guestrin. Blitz: A principled meta-algorithm for scaling sparse optimization. In *International Conference on Machine Learning*, pages 1171–1179, 2015.
- [32] Tyler B Johnson and Carlos Guestrin. A fast, principled working set algorithm for exploiting piecewise linear structure in convex problems. *arXiv preprint arXiv:1807.08046*, 2018.
- [33] L. Kall, J. Canterbury, J. Weston, W. S. Noble, and M. J. MacCoss. A semi-supervised machine learning technique for peptide identification from shotgun proteomics datasets. *Nat. Methods*, 4:923–25, 2007.
- [34] Lukas Käll, John D Storey, Michael J MacCoss, and William Stafford Noble. Assigning significance to peptides identified by tandem mass spectrometry using decoy databases. *Journal of proteome research*, 7(01):29–34, 2008.
- [35] S. S. Keerthi and D. DeCoste. A modified finite newton method for fast solution of large scale linear svms. *Journal of Machine Learning Research*, 6(Mar):341–361, 2005.
- [36] Mu-Chu Lee, Wei-Lin Chiang, and Chih-Jen Lin. Fast matrix-vector multiplications for large-scale logistic regression on shared-memory systems. In *Data Mining (ICDM), 2015 IEEE International Conference on*, pages 835–840. IEEE, 2015.
- [37] David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397, 2004.
- [38] Andy Lin, J Jeffrey Howbert, and William Stafford Noble. Combining high-resolution and exact calibration to boost statistical power: A well-calibrated score function for high-resolution ms2 data. *Journal of proteome research*, 17(11):3644–3656, 2018.
- [39] C. J. Lin, R. C. Weng, and S. S. Keerthi. Trust region newton method for large-scale logistic regression. *Journal of Machine Learning Research*, 9(Apr):627–650, 2008.
- [40] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- [41] The Matthew, Michael J MacCoss, William S Noble, and Lukas Käll. Fast and accurate protein false discovery rates on large-scale proteomics data sets with percolator 3.0. *Journal of The American Society for Mass Spectrometry*, 27(11):1719–1727, 2016.
- [42] Rory Mitchell and Eibe Frank. Accelerating the xgboost algorithm using gpu computing. *PeerJ Computer Science*, 3:e127, 2017.
- [43] Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, et al. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*, 2017.
- [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,

- Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [45] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [46] Virginia Smith, Simone Forte, Chenxin Ma, Martin Takáč, Michael I Jordan, and Martin Jaggi. Cocoa: A general framework for communication-efficient distributed optimization. *The Journal of Machine Learning Research*, 18(1):8590–8638, 2017.
- [47] M. Spivak and W. S. Noble. Learning score function parameters for improved spectrum identification in tandem mass spectrometry experiments. *J. Proteome Res.*, 11(9):4499–4508, 2012. PMC3436966.
- [48] Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. Chainer: a next-generation open source framework for deep learning. In *Proceedings of workshop on machine learning systems (LearningSys) in the twenty-ninth annual conference on neural information processing systems (NIPS)*, volume 5, pages 1–6, 2015.
- [49] Bo Wen, Chaoqin Du, Guilin Li, Fawaz Ghali, Andrew R Jones, Lukas Käll, Shaohang Xu, Ruo Zhou, Zhe Ren, Qiang Feng, et al. Ipeak: An open source tool to combine results from multiple ms/ms search engines. *Proteomics*, 15(17):2916–2920, 2015.
- [50] Zeyi Wen, Jiashuai Shi, Qinbin Li, Bingsheng He, and Jian Chen. Thundersvm: A fast svm library on gpus and cpus. *The Journal of Machine Learning Research*, 19(1):797–801, 2018.
- [51] Mathias Wilhelm, Judith Schlegl, Hannes Hahne, Amin Moghaddas Gholami, Marcus Lieberenz, Mikhail M Savitski, Emanuel Ziegler, Lars Butzmann, Siegfried Gessulat, Harald Marx, et al. Mass-spectrometry-based draft of the human proteome. *Nature*, 509(7502):582–587, 2014.
- [52] Mingguo Xu, Zhendong Li, and Liang Li. Combining percolator with x! tandem for accurate and sensitive peptide identification. *Journal of proteome research*, 12(6):3026–3033, 2013.
- [53] Masaru Yoshioka, Naoki Suganuma, Keisuke Yoneda, and Mohammad Aldibaja. Real-time object classification for autonomous vehicle using lidar. In *2017 International Conference on Intelligent Informatics and Biomedical Sciences (ICIIBMS)*, pages 210–211. IEEE, 2017.
- [54] Guo-Xun Yuan, Chia-Hua Ho, and Chih-Jen Lin. An improved glmnet for l1-regularized logistic regression. *Journal of Machine Learning Research*, 13(Jun):1999–2030, 2012.
- [55] Yong Zhuang, Yuchin Juan, Guo-Xun Yuan, and Chih-Jen Lin. Naive parallelization of coordinate descent methods and an application on multi-core l1-regularized classification. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1103–1112, 2018.