
Bandit-Driven Batch Selection for Robust Learning under Label Noise

Michal Lisicki

University of Guelph, Vector Institute
mlisicki@uoguelph.ca

Mihai Nica

University of Guelph, Vector Institute
nicam@uoguelph.ca

Graham W. Taylor

University of Guelph, Vector Institute
gwtaylor@uoguelph.ca

Abstract

We introduce a novel approach for batch selection in Stochastic Gradient Descent (SGD) training, leveraging combinatorial bandit algorithms. Our methodology focuses on optimizing the learning process in the presence of label noise, a prevalent issue in real-world datasets. Experimental evaluations on the CIFAR-10 dataset reveal that our approach consistently outperforms existing methods across various levels of label corruption. Importantly, we achieve this superior performance without incurring the computational overhead commonly associated with auxiliary neural network models. This work presents a balanced trade-off between computational efficiency and model efficacy, offering a scalable solution for complex machine learning applications.

1 Introduction

As applications increasingly demand larger and more complex deep learning models, the need for efficient training strategies has become paramount. One way to accelerate training and potentially improve model performance is through the use of Curriculum Learning (CL) and adaptive batch selection. These techniques optimize learning by selectively focusing on data samples that are intrinsically rich and informative at the most appropriate stages of the learning process. Such strategies not only accelerate convergence but also enhance the model’s ability to generalize [19, 20, 27].

While many methods use difficulty metrics to select easy, hard, or uncertain instances for training [31], a key area lies in handling noisy or mislabeled datasets [28]. This domain is particularly important for two reasons: a) the impact of batch selection strategies is easily measured, leading to more insightful conclusions; and b) it addresses the prevalent real-world scenarios where data is often sourced from the web [18] or crowdsourced [7], and a large portion is considered “unclean”.

Sample selection strategies using auxiliary Deep Neural Networks (DNN) effectively mitigate the impact of noisy or mislabeled data. However, these approaches incur substantial computational overhead, limiting their scalability [9, 13, 17, 33]. While alternative methods like SELFIE [26] offer computational efficiency, they are under-explored and rely on steps like re-labeling for optimal performance. Meanwhile, the literature on CL and batch selection offers numerous methods for efficient sample selection across diverse domains [8, 19].

This paper introduces a novel approach that synergizes insights from the CL and batch selection literature to enhance efficient sampling schemes, specifically targeting scenarios with prevalent label noise. Our methodology aims to achieve superior performance without the computational burden

often associated with deploying additional DNNs, thereby striking a balance between efficacy and computational efficiency. Unlike traditional CL approaches that focus on individual instances or tasks, our method refines the feedback loop from each training iteration to optimize the *selected batch*. This approach is particularly relevant for tackling the challenges posed by the increasing computational complexity and diversity of machine learning applications across various domains.

2 Background

Batch selection and curriculum learning CL [3] and its variants like Self-Paced Learning (SPL) [15] and Hard-Example Mining (HEM) [6, 19] provide frameworks for adaptive instance, batch, or task selection based on difficulty or importance. Despite the efficacy of these strategies in enhancing stability and convergence, a universal solution remains elusive, prompting exploration into varied strategies, new importance metrics, and advanced re-weighting and sampling techniques.

Re-weighting the model’s loss by instance, akin to importance-sampling techniques, has been investigated by [6, 19, 24]. These studies have indicated that re-weighting can stabilize gradient estimates and reduce bias in the original objective function. However, both Loshchilov and Hutter [19] and Chang et al. [6] have argued that the impact of this strategy on performance is limited, and that comparable or superior results can be achieved by sampling from a weight-induced distribution. Matiisen et al. [20] compared the *sample selection* strategies, ϵ -greedy, Boltzmann, and Thompson sampling, and concluded that the optimal strategy hinges heavily on the sample weight metric. A novel metric not tied to difficulty was introduced by Chang et al. [6], emphasizing samples with high prediction uncertainty, and inspired by active learning. The authors demonstrated that, by avoiding overly easy or hard instances, their strategy surpassed SPL or HEM on datasets like MNIST and CIFAR with and without label noise. Song et al. [27] introduced Recency Bias to boost SGD’s convergence by combining principles of [19] and [6]. The technique centers on prediction uncertainty, measured by predicted label entropy, as its adaptive sample selection metric. It uses a Boltzmann distribution with energy based on prediction uncertainty and a pressure parameter. Leveraging a sliding window, it emphasizes recent scores, mitigating overfitting and slow convergence. Importantly, Active Bias, Recency Bias, and our approach add minimal computational load to model training.

Automated curriculum learning (ACL) distinguishes itself among CL approaches by the degree of control over the learning process. In ACL, the selection of tasks is determined dynamically using an algorithm, typically an RL or a bandit method. Graves et al. [8] and Matiisen et al. [20] have proposed utilizing a non-stationary bandit (Exp3). They demonstrated that when an agent lacks prior knowledge of its tasks, ACL can significantly boost training efficiency relative to uniform sampling. Moreover, a bandit algorithm can discover complex orderings and opportunities for efficient knowledge transfer in an unsorted curriculum. Although prior literature has focused primarily on task-based ACL, the same principles can be utilized in instance and batch selection.

This paper aims to build upon the mentioned foundational CL techniques by introducing efficient batch selection methods, particularly in the context of learning with noisy labels.

Efficient learning with noisy labels Learning with Noisy Labels (LNL) shares a connection with batch selection but has a different objective. While batch selection picks instances that inherently aid training, LNL focuses on distinguishing between those with clean and noisy labels. Both fields converge when the right batch selection strategy is used to isolate the “clean” instances.

LNL is challenging due to DNNs’ tendency to memorize complex and possibly incorrect instances, after initially learning simpler patterns [1, 34]. This can lead to memorization of inaccurate labels, compromising model generalization. Mislabeled data can also cause confirmation bias, which arises when models overfit to early-selected instances. Multi-network and co-training [9, 33] address these issues but add computational overhead and complexity. Both challenges highlight the need for robust training methods.

The “small-loss trick” is a commonly utilized tool for filtering out noisy labels by deeming instances with smaller losses as likely clean. While prevalent in DNN approaches (e.g., [9, 13, 17, 25, 33]), this method is not optimal when noisy and clean example distributions overlap significantly. Alternative methods, like measuring prediction uncertainty over time, have been explored as indicators of label corruption [6, 22, 26].

Similarly to batch selection, LNL methods can be broadly categorized into: loss correction and sample selection [28]. Loss correction includes re-weighting or re-labeling. Active Bias [6] re-weights instances based on prediction variance, bridging LNL and batch selection. A method by Ren et al. [24] uses a clean validation set to dynamically assign weights. Re-labeling refines labels from a mix of noisy labels and DNN predictions, as seen in [23, 29, 32], effectively providing data augmentation. While innovative, these methods pose risks like overfitting to noisy labels and add complexity. Filtering out noisy labels, by contrast, provides a simpler approach but sacrifices information from label refinement. Sample selection filters out mislabeled data during training, often using an auxiliary DNN. MentorNet [13] stands out as a pivotal multi-network approach. It supervises a StudentNet by emphasizing “clean” instances and refining the learning trajectory based on feedback. Han et al. [9] proposed a “Co-teaching” paradigm, an alternative, where two DNNs are trained simultaneously and share insights on small-loss instances to diminish errors from noisy labels. “Co-teaching+” [33] tackles the risk of two networks reaching a consensus with an ‘update by disagreement’ strategy. The “deep abstaining classifier” [30] is a feature-based multi-network approach, which is particularly effective against structured noise. DivideMix [17] has demonstrated the state-of-the-art LNL performance by employing a semi-supervised (SSL) approach. It dynamically segregates training data into clean and noisy sets using a Gaussian Mixture Model (GMM) and the small-loss trick. To avoid confirmation bias, it utilizes co-teaching. The mislabeled instances are stripped of their labels and refined using SSL [4]. Despite their ability to counteract confirmation bias, the multi-network training approaches often come with significant computational overhead.

SELFIE [26] is a hybrid approach, combining both loss correction and sample selection. SELFIE seeks to refurbish labels of unclean samples selectively, based on uncertainty, and leverage them along with clean samples, to further reduce false corrections while fully exploiting the current training data. While methods like SELFIE are computationally efficient and provide increase in performance, we argue that their selection strategies can be made better. In this paper we focus on improving the batch selection methodology, and compare performance to the pure selection method, Active Bias [6], and two bandit approaches — the *Exponential-weight algorithm for Exploration and Exploitation* (Exp3) [2] and *Follow the Perturbed Leader* (FPL) [16, 21]. Details on these algorithms can be found in Sec. 3 and in Appendix, Sec. A.

Improving sampling efficiency by exploration The challenge of balancing exploration and exploitation is inherent in the process of batch selection. This balance is crucial for identifying new instances that can enhance training efficiency and subsequently leverage them for optimal training outcomes. In addition, it is essential to account for the high degree of non-stationarity in neural network training. Specifically, a network can typically be trained on a particular data instance for only a few iterations before it risks overfitting. To address this issue, our approach aims to manage non-stationarity by dynamically adapting weight estimates. This adaptation can be achieved either through periodic reevaluation, which may be computationally expensive, or by employing a discounted moving average.

So far we are aware of only one study that has directly compared various exploration-exploitation strategies in the context of instance selection: its mixed results suggest such strategies depend highly on the researcher’s choice of a sample weight metric [20]. Although ϵ -greedy and UCB bandit methods have demonstrated effective performance in instance selection [6, 20], the Boltzmann exploration strategy has recently gained prominence in this subfield [5, 19]. In particular, the adversarial bandit — Exp3 [2], which uses Boltzmann exploration, is commonly utilized as a baseline in non-stationary environments, and has been shown to be particularly effective in automated curriculum learning [8, 20].

Our work diverges from prior studies focused on choosing instances [6, 27] or tasks [8, 20], and targets batch selection instead. We utilized the FPL strategy, which can be thought of as a natural extension of Exp3 into combinatorial (batch), rather than individual (instance) action selection.

3 Methods

Adversarial multi-armed bandit problem A classic baseline approach for non-stationary environments is the adversarial bandit, in particular, the Exp3 algorithm and its variants. In an adversarial K -armed bandit problem, at each time step $t \in \{1, 2, \dots, T\}$, the player selects an action $a_t \in \{1, 2, \dots, K\}$ and then an adversary, with full knowledge of the player’s previous actions, assigns

a reward vector $\mathbf{r}_t = (r_{t,1}, r_{t,2}, \dots, r_{t,K}) \in [0, 1]^K$ across all actions. The player receives a reward r_{t,a_t} corresponding to the selected action a_t . There is typically almost no restrictions on how the adversary can choose the reward vectors, as long as the sequence of reward vectors $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_T$ is fixed in advance or chosen based on the player’s past actions. The player’s goal remains to maximize the total collected reward or equivalently, to minimize regret.

Combinatorial bandits for batch selection In order to select a full batch of instances at once we need to utilize the combinatorial bandit paradigm, which considers the joint utility of combinations of “basic arms”. Formally, combinatorial bandits can be considered a type of bandit where a subset of arms is selected in a form of a binary vector $\mathbf{a} \in \{0, 1\}^d$, and the final reward is derived from either a Hadamard or a dot product of that vector with the reward vector \mathbf{r} . In this work we consider only the subset of a pre-specified batch size m , s.t. $\|\mathbf{a}\|_1 = m$, and a semi-bandit reward model (see Appendix, Sec. A). A direct application of Exp3 to the semi-bandit problem would entail monitoring the sequence of estimates for $\binom{K}{m}$ arms, a task that is computationally infeasible. The state-of-the-art approach to semi-bandits is *Follow the Perturbed Leader* (FPL) [10], which mimics Exp3, but estimates probabilities using reward perturbations, rather than storing them directly. FPL was originally introduced by Hannan [10] and Kalai and Vempala [14], with an efficient version operating on a principle of geometric re-sampling (GR) proposed by Neu and Bartók [21]. In this work, we adapt the FPL algorithm to batch selection.

FPL operates over n rounds, maintaining a vector of weights $w_{t,i}$ for each action \mathbf{a}_i in the action set \mathcal{A} . In our case \mathbf{a} is a binary vector, such that setting an i -th action $a_i = 1$ corresponds to selecting an i -th instance \mathbf{x}_i . Each round the algorithm perturbs the weights with noise $\boldsymbol{\rho}_t$ from distribution Q , selecting the action \mathbf{a}_t that maximizes the inner product with the perturbed weight vector. While Neu and Bartók [21] used the Exp(1) distribution for Q , recent work by Honda et al. [11] suggests the Fréchet(2) (also known as inverse Weibull) distribution yields optimal regret in adversarial settings. As opposed to the algorithms presented in literature [11, 21] we estimate reward, rather than loss associated with each arm. We have found this adaptation to significantly improve performance in our application, however we acknowledge that while this algorithm remains in line with reward estimation done in Exp3, the original theoretical performance guarantees for combinatorial arm selection may no longer apply.

Algorithm 1: Follow The Perturbed Leader (Reward-guided)

Data: $\mathcal{A}, n, \eta, M, Q$

```

1 for  $i = 1$  to  $d$  do
2    $w_{0,i} \leftarrow 0$  // Initialize weight vector
3 for  $t = 1$  to  $n$  do
4   Sample  $\boldsymbol{\rho}_t \sim Q$  // Sample weight perturbations
5   Compute  $\mathbf{a}_t = \arg \max_{\mathbf{a} \in \mathcal{A}} \langle \mathbf{a}, \eta \mathbf{w}_{t-1} + \boldsymbol{\rho}_t \rangle$  // Choose combinatorial action
6    $\mathbf{r}_t \sim \nu_{\mathbf{a}_t}$  // Draw reward vector from arm  $\mathbf{a}_t$  of MAB  $\nu$ 
7   foreach  $i$  with  $a_{t,i} = 1$  do
8     // Geometric Re-sampling
9     Sample  $\sigma_{t,i} \sim \text{Geometric}(p_{t,i})$ 
10     $\hat{r}_{t,i} = \min\{M, \sigma_{t,i}\} a_{t,i} r_{t,i}$  // Compute bounded reward estimate
11     $w_{t,i} = w_{t-1,i} + \hat{r}_{t,i}$  // Update weight of chosen action

```

Following action selection, for each i where $a_{t,i} = 1$, the algorithm proceeds with a geometric re-sampling (GR) step. Sampling from the geometric distribution estimates $1/p_i$ and in practice is not done directly, but rather by sampling arms from $Q + \eta \mathbf{w}_{t-1}$ and counting the number of iterations to re-occurrence. M is the cap on sampling size, to trade off computational efficiency with estimation accuracy. The algorithm draws a sample $\sigma_{t,i}$ from the approximated geometric distribution, and computes a bounded reward estimate $\hat{r}_{t,i}$ in the same way as Exp3, as an importance-weighted estimate, by $a_{t,i} \sigma_{t,i} r_{t,i}$.¹

Finally, the algorithm updates the weight $w_{t,i}$ of the chosen action a_i by adding the reward estimate $\hat{r}_{t,i}$ to the previous weight $w_{t-1,i}$. This process continues for n rounds, enabling the algorithm to

¹ $\sigma_{t,i}$ estimates $1/p_i$




Figure 1: Test error over the course of training with confidence intervals (CI) over 5 runs for Uniform, Active Bias (weighted), Exp3 and FPL, for label noise ratio $\in \{0.1, 0.3, 0.5\}$.

effectively explore and exploit the action space by balancing the current estimated rewards and the exploration noise introduced by perturbations. While FPL may require more computational resources compared to Exp3, it offers the advantage of reducing dependency on the combinatorial action space. This makes FPL a practical choice for real-world sequential decision-making tasks.

Label noise According to Song et al. [28], label noise can be either instance-independent, characterized by constant rates and probabilities, or instance-dependent, where corruption probabilities vary with data features and true labels. This study concentrates on symmetric, instance-independent noise to provide a baseline in a controlled setting.

LNL weight metric Choosing the right metric to select informative instances is still an open problem. In the field of LNL, metrics based on prediction loss [9, 13, 17, 33] and prediction uncertainty [6, 22, 26] have shown particular promise. The following metric was proposed by Chang et al. [6] as part of the Active Bias method:

$$w_i \propto \widehat{\text{var}}(p_{\mathcal{H}_i^{t-1}}(y_i | \mathbf{x}_i)),$$

where, for each instance \mathbf{x}_i , it saves prediction probabilities for their target class over time in a history buffer \mathcal{H}_i , and then computes their variance.

We employed this metric in our study, as it was shown suitable both for LNL and for batch selection in general. Unlike the metrics that are derived from the change in the state of the model, the probability-based metrics reflect the model’s current confidence in its predictions, rendering them independent of the target solutions, and therefore, consistent across instances. This property makes them inherently balanced for problems such as LNL. However, it should be noted that while these metrics offer advantages, they do not directly track the progression of training. Therefore, following Song et al. [27], we limit the size of the history to 10 predictions.

The estimated weights serve two main purposes: either to re-weight the loss as in [24] or to parameterize the probability distribution over data instances. The latter often employs a Boltzmann distribution (e.g. [8, 20]): $P_s(i|H, S_e, D) = e^{w_i/\tau}/Z$ where Z is the normalization constant, H denotes the history of scores (e.g. instance losses or prediction probabilities), S_e is the set of samples used in the current epoch, and $D = \{(\mathbf{x}_i, y_i) \mid i = 1, 2, \dots, N\}$ represents the dataset. Given our interest in the role of exploration in sample efficiency, we primarily focus on sampling methods underpinned by bandit algorithms such as Exp3, which also employs a Boltzmann-like distribution.

4 Results

Experimental Setup We evaluated the performance of various sampling methods including Uniform Sampling, Active Bias, Exp3, and FPL on the CIFAR-10 dataset using a DenseNet model [12] with 40 layers. We used the Adam optimizer with momentum 0.9 and an initial learning rate of 0.1 that is decayed by multiplicative factor of 0.1 after 40k and 60k iterations. The batch size was set to 128 and we ran 200 epochs, consisting of 391 batches each. All methods were repeated 5 times with different seeds, under varying label corruption percentages ranging from 0% to 50%. We report the mean and 95% confidence intervals (CI) of test accuracy achieved by each method. We will release a PyTorch implementation to reproduce our experiments upon paper acceptance.

Results Under all label corruption scenarios, FPL exhibited significantly reduced noise and superior performance compared to the other methods, with Exp3 outperforming Active Bias, and Active Bias performing better than Uniform Sampling. In conditions with no label noise, no significant improvement was observed across methods, revealing a potential limitation in sensitivity to “hard to classify” instances and an overfocus on mislabeling.




Figure 2: Left: Lowest and final epoch test errors (%) for each method on CIFAR-10 by noise ratio. Right: Visualizing last epoch performance.

Discussion The methods maintained a consistent ranking across noise levels, with the performance gap widening as noise increased (see Fig. 1 and 2). FPL consistently yielded smooth and stable convergence, due to its ability to choose informative instances. When adopting the same weight metric and neural network architecture as Active Bias, our results show that implementing a bandit strategy can lead to significant performance gains. This underscores the importance of not just selecting an optimal weight metric, but also employing a beneficial exploration policy.




Figure 3: Analysis of instance selection for Uniform, Exp3, and FPL¹ with 20% label noise, showing selection occurrence during the initial (a) and final (b) 1000 iterations and total proportion of mislabeled instances (c; front) over total occurrences (c; background). The order of curves at index 0 aligns well with the overall performance of the methods, revealing a concentration of selection in Exp3 and FPL, particularly pronounced in FPL, with Exp3 demonstrating overfitting to a limited set of instances.

Analysis of instance occurrences (Fig. 3) reveals insights into the differences in sampling strategies, addressing our initial inquiry into performance gain from utilizing batch- as opposed to instance-based feedback. Initially (Fig. 3a), all methods show similar selection frequencies, but distinctions emerge as training concludes (Fig. 3b), especially for Exp3 and FPL. Notably, the algorithm’s pattern of concentration on specific instances correlates well with its performance. While Exp3 resembles an exponential distribution, FPL produces a threshold at about 20k instances, filtering 30k of the remaining images. The preference for ‘clean’ instances between the 10k and 20k sorted index intensifies towards the end of training, indicating the algorithm’s inclination to retain instances

initially deemed ‘clean’. These insights emphasize FPL’s efficiency as an m -set combinatorial bandit method, and highlight its suitability for batch selection.

In Fig. 3c, we display curves representing total counts throughout the run, providing a holistic view of each method’s sample selection strategy. Over these, with solid lines, we superimpose the percentage of mislabeled instances within a sliding window of 1000 sorted instances. Each point on the overlay represents the mislabeling percentage within that window, revealing a trend: instances sampled less frequently (toward the right) have higher mislabeling percentages. This visualization supports our hypothesis that bandit methods with uncertainty-based metrics, like Exp3 and FPL, enhance performance by focusing on and filtering out mislabeled instances.

While Exp3 effectively identifies mislabeled instances like FPL (Fig. 3c), it tends to overfit to a narrow set and over-explore the rest (Fig. 3b). This is expected as Exp3 is an instance-based algorithm. However, this overfitting poses risks to its efficacy, as consistently selecting the same subset of impactful instances, combined with a broad array of less pertinent ones, leads to lower performance. Conversely, FPL, by adjusting weights in accordance with other instances, revisits a larger, more balanced subset regularly, forming more informative and consequential batches, ensuring optimal selection in batch training scenarios.

As weights play a crucial role in understanding the learning process in depth, we further analyze the weight dynamics of FPL in Fig. 4. Initially set to 0, the weights adjust smoothly throughout training, maintaining a balance in instance selection without anomalies or overfitting. There is a 40%-60% split in instance selection (Fig. 3b) that is clearly reflected in the weights, with those corresponding to highly informative instances increasing rapidly, whereas those consistently labeled (either correctly or as mislabeled) remaining closer to zero. It is worth noting here that all instances were selected at least once, with all weights turning strictly positive by the end of the run. Entropy visualization (Fig. 4a) further emphasizes effective convergence on a well-sized subset of instances, reinforcing selection for better exploitation without excessive exploration.




Figure 4: Instance weight visualization. Entropy-aggregated weights are visualized over time in (a). Once all instances are selected by epoch 40, the entropy gradually declines as certain instances gain importance. This pattern indicates effective diversification without overfitting. In (b), a random subset of weights is displayed over time to further validate their individual trajectories. In (c) the lowest weights are shown to dynamics of weights that remain close to 0.

Scalability and hyperparameter sensitivity We ran our experiments using $\eta = 0.3$ and $\gamma = 0.1$ for Exp3 and $\eta \approx 18$, $\beta \approx 20$, and the Fréchet(0.45) distribution for FPL. To show that FPL has small sensitivity to these hyperparameters, we ran a grid search in their vicinity (see Fig. 5). We found the number of GR samples to be optimal between 500 and 1000. In that range GR introduces an additional computational overhead of 20%-40%. This may seem alarming at first, however, we point out that GR is *embarrassingly parallelizable* and instance-based, which makes it scalable in practical applications.

Limitations and future directions Exp3’s slower adaptation and potential benefits of its variants like Exp3.P or Exp3.IX warrant further consideration. FPL excels in balancing exploration and exploitation but shows limited improvement in noise-free scenarios, suggesting a potential overfocus on mislabeled instances. While all methods generalize well, tests were limited in scope. Future work will include naturally noisy sets like WebVision [18], as well as metrics like area under the margin (AUM) [22], to deepen insights and enhance results.

¹Active bias method is excluded here as we use its loss re-weighting variant, and so its sampling distribution is the same as uniform.

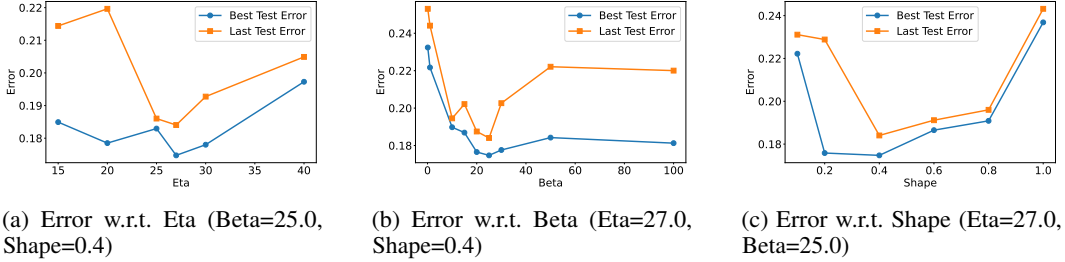


Figure 5: FPL hyperparameter sensitivity analysis for 50% label noise.

5 Conclusions

This investigation into the performance of sampling methods under different noise conditions has revealed key insights into their adaptability, stability, and algorithmic nuances. FPL’s effective balance between exploration and exploitation, particularly its focus on uncertain instances, underscores its superior performance. Nonetheless, the absence of marked improvement in noise-free settings and the limited scope of our experiments highlight avenues for future research and refinement.

References

- [1] Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A Closer Look at Memorization in Deep Networks, July 2017. URL <http://arxiv.org/abs/1706.05394>. arXiv:1706.05394 [cs, stat].
- [2] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- [3] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML ’09*, pages 41–48, New York, NY, USA, June 2009. Association for Computing Machinery. ISBN 978-1-60558-516-1. doi: 10.1145/1553374.1553380. URL <https://doi.org/10.1145/1553374.1553380>.
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32, 2019.
- [5] Nicolò Cesa-Bianchi, Claudio Gentile, Gábor Lugosi, and Gergely Neu. Boltzmann Exploration Done Right, November 2017. URL <http://arxiv.org/abs/1705.10257>.
- [6] Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. Active Bias: Training More Accurate Neural Networks by Emphasizing High Variance Samples. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/2f37d10131f2a483a8dd005b3d14b0d9-Abstract.html>.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [8] Alex Graves, Marc G. Bellemare, Jacob Menick, Remi Munos, and Koray Kavukcuoglu. Automated curriculum learning for neural networks. In *international conference on machine learning*, pages 1311–1320. PMLR, 2017.
- [9] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/hash/a19744e268754fb0148b017647355b7b-Abstract.html.

- [10] James Hannan. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.
- [11] Junya Honda, Shinji Ito, and Taira Tsuchiya. Follow-the-Perturbed-Leader Achieves Best-of-Both-Worlds for Bandit Problems. In *Proceedings of The 34th International Conference on Algorithmic Learning Theory*, pages 726–754. PMLR, February 2023. URL <https://proceedings.mlr.press/v201/honda23a.html>.
- [12] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks, January 2018. URL <http://arxiv.org/abs/1608.06993>. arXiv:1608.06993 [cs].
- [13] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. MentorNet: Learning Data-Driven Curriculum for Very Deep Neural Networks on Corrupted Labels, August 2018. URL <http://arxiv.org/abs/1712.05055>. arXiv:1712.05055 [cs].
- [14] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. In *Proceedings of the 16th Annual Conference on Computational Learning Theory*. Citeseer, 2003.
- [15] M. Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 1*, NIPS’10, pages 1189–1197, Red Hook, NY, USA, December 2010. Curran Associates Inc.
- [16] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020. URL <https://tor-lattimore.com/downloads/book/>.
- [17] Junnan Li, Richard Socher, and Steven C. H. Hoi. DivideMix: Learning with Noisy Labels as Semi-supervised Learning, February 2020. URL <http://arxiv.org/abs/2002.07394>. arXiv:2002.07394 [cs].
- [18] Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*, 2017.
- [19] Ilya Loshchilov and Frank Hutter. Online Batch Selection for Faster Training of Neural Networks, April 2016. URL <http://arxiv.org/abs/1511.06343>.
- [20] Tamber Matisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher-student curriculum learning. *IEEE transactions on neural networks and learning systems*, 2019.
- [21] Gergely Neu and Gábor Bartók. Importance weighting without importance weights: An efficient algorithm for combinatorial semi-bandits. *Journal of Machine Learning Research*, 17(154): 1–21, 2016. URL <http://jmlr.org/papers/v17/15-091.html>.
- [22] Geoff Pleiss, Tianyi Zhang, Ethan R. Elenberg, and Kilian Q. Weinberger. Identifying Mislabeled Data using the Area Under the Margin Ranking, December 2020. URL <http://arxiv.org/abs/2001.10528>. arXiv:2001.10528 [cs, stat].
- [23] Scott Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training Deep Neural Networks on Noisy Labels with Bootstrapping, April 2015. URL <http://arxiv.org/abs/1412.6596>. arXiv:1412.6596 [cs].
- [24] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to Reweight Examples for Robust Deep Learning. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4334–4343. PMLR, July 2018. URL <https://proceedings.mlr.press/v80/ren18a.html>.
- [25] Yanyao Shen and Sujay Sanghavi. Learning with Bad Training Data via Iterative Trimmed Loss Minimization, February 2019. URL <http://arxiv.org/abs/1810.11874>. arXiv:1810.11874 [cs, stat].
- [26] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. SELFIE: Refurbishing Unclean Samples for Robust Deep Learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5907–5915. PMLR, May 2019. URL <https://proceedings.mlr.press/v97/song19b.html>. ISSN: 2640-3498.

- [27] Hwanjun Song, Minseok Kim, Sundong Kim, and Jae-Gil Lee. Carpe Diem, Seize the Samples Uncertain “at the Moment” for Adaptive Batch Selection. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM '20*, pages 1385–1394, New York, NY, USA, October 2020. Association for Computing Machinery. ISBN 978-1-4503-6859-9. doi: 10.1145/3340531.3411898. URL <https://dl.acm.org/doi/10.1145/3340531.3411898>.
- [28] Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [29] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint Optimization Framework for Learning with Noisy Labels, March 2018. URL <http://arxiv.org/abs/1803.11364>. arXiv:1803.11364 [cs, stat].
- [30] Sunil Thulasidasan, Tanmoy Bhattacharya, Jeff Bilmes, Gopinath Chennupati, and Jamal Mohd-Yusof. Combating Label Noise in Deep Learning Using Abstention, August 2019. URL <http://arxiv.org/abs/1905.10964>. arXiv:1905.10964 [cs, stat].
- [31] Xin Wang, Yudong Chen, and Wenwu Zhu. A Survey on Curriculum Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4555–4576, September 2022. ISSN 1939-3539. doi: 10.1109/TPAMI.2021.3069908.
- [32] Kun Yi and Jianxin Wu. Probabilistic End-To-End Noise Correction for Learning With Noisy Labels. pages 7017–7025, 2019. URL https://openaccess.thecvf.com/content_CVPR_2019/html/Yi_Probabilistic_End-To-End_Noise_Correction_for_Learning_With_Noisy_Labels_CVPR_2019_paper.html.
- [33] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W. Tsang, and Masashi Sugiyama. How does Disagreement Help Generalization against Label Corruption?, May 2019. URL <http://arxiv.org/abs/1901.04215>. arXiv:1901.04215 [cs, stat].
- [34] Chiyan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization, February 2017. URL <http://arxiv.org/abs/1611.03530>. arXiv:1611.03530 [cs].

A Bandit algorithms

Classic adversarial bandits The classic adversarial algorithm, Exp3 (Alg. 8), employs a multi-step process for importance adjustments of reward estimates. First, it adjusts the reward r_j for arm j at time t using the formula $\hat{r}_j = \frac{r_{t,j} \mathbb{I}_{j=i_t}}{p_{t,j}}$, where $p_{t,j}$ is the probability of choosing arm j . These adjusted rewards are then used to estimate sample weights w_j . Finally, these weights w_j are utilized in the Boltzmann distribution for sampling instances.

Algorithm 2: Exp3 Algorithm

Data: $\gamma \in (0, 1]$, K

- 1 **for** $i = 1$ **to** K **do**
- 2 $w_{1,i} \leftarrow 1$ // Initialize weights
- 3 **for each round** $t = 1, 2, \dots$ **do**
- 4 $p_{t,j} \leftarrow (1 - \gamma) \frac{w_{t,j}}{\sum_{k=1}^K w_{t,k}} + \frac{\gamma}{K}$ // Compute pmf
- 5 $i_t \sim p_t$ // Sample action
- 6 $r_{t,i_t} \sim \nu_{i_t}$ // Draw reward from arm i_t of MAB ν
- 7 $\hat{r}_j \leftarrow \frac{r_{t,j} \mathbb{I}_{\{j=i_t\}}}{p_{t,j}}$ // Compute reward estimate
- 8 $w_{t+1,j} \leftarrow w_{t,j} \exp(\gamma \hat{r}_{t,j})$ // Update weights

The Exp3 algorithm is very efficient computationally and is suitable for task or individual instance selection, but it doesn’t take into account an impact of a full batch of instances on the performance

of the neural network, which may result in suboptimal performance when competing instances are present in the same batch.

Analysis and adoption of semi-bandit feedback in combinatorial bandits Combinatorial bandits can be categorized as full-information, semi-bandits, or full-bandit feedback. In the full-bandit feedback scenario we observe just one reward per batch. While maximizing this reward is our ultimate objective, ignoring the available information about rewards received for individual arms makes the decision process suboptimal. The full information setup, where rewards from all the arms are observed is computationally infeasible, as it requires re-evaluating the network on all the instances. In our research we adopt the semi-bandit, in which the rewards are observed only for the basic arms selected in the current round. This setup aligns well with our problem, where we observe the rewards for instances that the neural network was trained on in current iteration. As this information is readily available, no additional passes through the network are required.

Efficient variant of FPL was proposed by Neu and Bartók [21], who deployed Geometric Resampling to estimate probabilities for importance-weighted reward estimates (re-weighting step), making FPL the first computationally feasible solution to semi-bandits with strong guarantees. While other methods may offer comparable or superior theoretical performance in terms of upper regret bounds, they frequently suffer from computational inefficiency or require additional optimization steps, rendering them impractical for real-world applications. The primary appeal of FPL lies in its computational efficiency.