

SUBMODULAR MUTUAL INFORMATION FOR TARGETED DATA SUBSET SELECTION

Anonymous authors

Paper under double-blind review

ABSTRACT

With the rapid growth of data, it is becoming increasingly difficult to train or improve deep learning models with the right subset of data. We show that this problem can be effectively solved at an additional labeling cost by *targeted data subset selection* (TSS) where a subset of unlabeled data points similar to an auxiliary set are added to the training data. We do so by using a rich class of Submodular Mutual Information (SMI) functions and demonstrate its effectiveness for image classification on CIFAR-10 and MNIST datasets. Lastly, we compare the performance of SMI functions for TSS with other state-of-the-art methods for closely related problems like active learning. Using SMI functions, we observe $\approx 30\%$ **gain over the model’s performance before re-training with added targeted subset**; $\approx 12\%$ **more than other methods**.

1 INTRODUCTION

Recent times have seen unprecedented growth in data across modalities such as text, images and videos. This has naturally given rise to techniques for finding effective smaller subsets of the data for a variety of end-tasks. An example of this is data subset selection for efficient and/or cost-effective training of machine learning models, wherein we need to select samples which are most informative for training a model. Training on such smaller subsets of data often entails significant speedups and reduction in labeling time/cost without sacrificing much on accuracy. Another flavor of this is targeted data subset selection which focuses on improving an existing model which is performing poorly in specific cases or improving a dataset which is imbalanced in certain attributes. Quite often, in these end-tasks, we want to be able to select *subsets that align well with a certain target set*.

1.1 TARGETED DATA SUBSET SELECTION

In real-world settings, there is often a distribution shift between training data and test data. Moreover, this distribution shift is reoccurring because new critical cases arise quite often. In such cases a model’s performance can be improved (at a given additional labeling cost) by augmenting the training data with some most informative samples matching the target distribution (hence called *targeted* subset) from a large pool of unlabeled data. One way of achieving this is by assuming access to a clean validation set matching the target set distribution and using it as a target. Another example is where the target set is a critical slice of the data (*e.g.*, indoor images of people in the dark or images from specific classes that the user might care about) and we want to improve the model’s performance on the target without sacrificing the overall accuracy and with minimum additional labeling costs (Fig. 1).

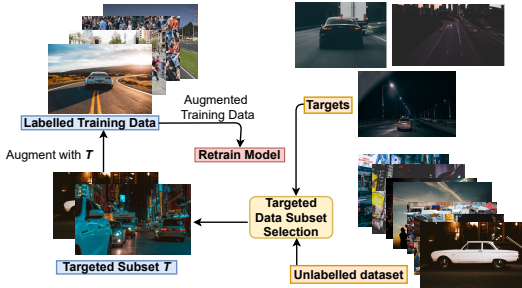


Figure 1: Motivating example for targeted data subset selection (TSS): the night images (target) are under-represented in training data. TSS mines for night images and augments the training data to improve the performance of the final model.

2 PRELIMINARIES

Submodular Functions: We let \mathcal{V} denote the *ground-set* of n data points $\mathcal{V} = \{1, 2, 3, \dots, n\}$ and a set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}$. The function f is submodular (Fujishige, 2005) if it satisfies the diminishing marginal returns, namely $f(j|\mathcal{X}) \geq f(j|\mathcal{Y})$ for all $\mathcal{X} \subseteq \mathcal{Y} \subseteq \mathcal{V}, j \notin \mathcal{Y}$. Facility location, set cover, log determinants, *etc.* are some examples (Iyer, 2015). Due to close connections between submodularity and entropy, submodular functions can also be viewed as *information functions* (Zhang & Yeung, 1998). Submodularity ensures that a greedy algorithm achieves bounded approximation factor when maximized (Nemhauser et al., 1978).

Submodular Mutual Information (MI): Given a set of items $\mathcal{A}, \mathcal{B} \subseteq \mathcal{V}$, the submodular mutual information (MI) (Gupta & Levin, 2020; Iyer et al., 2020) is defined as $I_f(\mathcal{A}; \mathcal{B}) = f(\mathcal{A}) + f(\mathcal{B}) - f(\mathcal{A} \cup \mathcal{B})$. Intuitively, this measures the similarity between \mathcal{B} and \mathcal{A} and we refer to \mathcal{B} as the query set.

Kaushal et al. (2020) extend MI to handle the case when the *target* can come from an auxiliary set \mathcal{V}' different from the ground set \mathcal{V} . For targeted data subset selection, \mathcal{V} is the source set of data instances and the target is a subset of data points (validation set or the specific set of examples of interest). Let $\Omega = \mathcal{V} \cup \mathcal{V}'$. We define a set function $f : 2^{\Omega} \rightarrow \mathbb{R}$. Although f is defined on Ω , the discrete optimization problem will only be defined on subsets $\mathcal{A} \subseteq \mathcal{V}$. To find an optimal subset given a query set $\mathcal{Q} \subseteq \mathcal{V}'$, we can define $g_{\mathcal{Q}}(\mathcal{A}) = I_f(\mathcal{A}; \mathcal{Q})$, $\mathcal{A} \subseteq \mathcal{V}$ and maximize the same.

2.1 EXAMPLES OF SMI FUNCTIONS

We use the MI functions recently introduced in Iyer et al. (2020); Gupta & Levin (2020) and their extensions introduced in Kaushal et al. (2020). For any two data points $i \in \mathcal{V}$ and $j \in \mathcal{Q}$, let sq_{ij} denote the similarity between them.

Graph Cut MI: The submodular mutual information (SMI) instantiation of graph-cut (GCMI) is defined as:

$$I_f(\mathcal{A}; \mathcal{Q}) = 2 \sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{Q}} sq_{ij} \quad (1)$$

Since maximizing GCMI maximizes the joint pairwise sum with the query set, it will lead to a summary similar to the query set \mathcal{Q} . In fact, specific instantiations of GCMI have been intuitively used for query-focused summarization for videos Vasudevan et al. (2017) and documents Lin (2012); Li et al. (2012).

Facility Location MI - V1: In the first variant of FL, we set D to be V . The SMI instantiation of FL1MI can be defined as:

$$I_f(\mathcal{A}; \mathcal{Q}) = \sum_{i \in \mathcal{V}} \min(\max_{j \in \mathcal{A}} sq_{ij}, \eta \max_{j \in \mathcal{Q}} sq_{ij}) \quad (2)$$

The first term in the $\min(\cdot)$ of FL1MI models diversity, and the second term models query relevance. An increase in the value of η causes the resulting summary to become more relevant to the query.

Facility Location MI - V2: In the V2 variant, we set D to be $V \cup \mathcal{Q}$. The SMI instantiation of FL2MI can be defined as:

$$I_f(\mathcal{A}; \mathcal{Q}) = \sum_{i \in \mathcal{Q}} \max_{j \in \mathcal{A}} sq_{ij} + \eta \sum_{i \in \mathcal{A}} \max_{j \in \mathcal{Q}} sq_{ij} \quad (3)$$

FL2MI is very intuitive for query relevance as well. It measures the representation of data points that are the most relevant to the query set and vice versa. It can also be thought of as a bidirectional representation score.

Log Determinant MI: The SMI instantiation of LogDetMI can be defined as:

$$I_f(\mathcal{A}; \mathcal{Q}) = \log \det(S_{\mathcal{A}}) - \log \det(S_{\mathcal{A}} - \eta^2 S_{\mathcal{A}, \mathcal{Q}} S_{\mathcal{Q}}^{-1} S_{\mathcal{A}, \mathcal{Q}}^T) \quad (4)$$

$S_{\mathcal{A}, \mathcal{B}}$ denotes the cross-similarity matrix between the items in sets \mathcal{A} and \mathcal{B} . The similarity matrix is constructed in such a way that the cross-similarity between \mathcal{A} and \mathcal{Q} is multiplied by η to control the trade-off between query-relevance and diversity.

3 A FRAMEWORK FOR TARGETED DATA SUBSET SELECTION

We apply SMI functions to a simple setting of targeted data subset selection for improving a model’s accuracy on some target classes/instances at a given additional labeling cost (k instances) and without compromising on the overall accuracy. Let \mathcal{E} be an initial training set of labeled instances and \mathcal{T} be the set of examples that the user cares about and desires better performance on. Let \mathcal{U} be a large unlabeled dataset. We maximize a MI function $I_f(\mathcal{A}; \mathcal{T})$ to compute an optimal subset $\mathcal{A}^* \subseteq \mathcal{U}$ of size k given \mathcal{T} as target (query) set. We then augment \mathcal{E} with labeled \mathcal{A}^* and re-train the model to achieve better accuracy without compromising on the accuracy of other classes/instances. Through instantiating a rich class of MI functions including GCMI, FL1MI, FL2MI, COM and LogDetMI, TSS offers a rich treatment to targeted subset selection. Our framework allows for adding an explicit diversity term ($\gamma g(\mathcal{A})$) helpful in cases such as GCMI that do not model diversity. The algorithm is summarized in Algorithm 1.

Algorithm 1 TSS

Require: Initial Labeled set of Examples: \mathcal{E} , large unlabeled dataset: \mathcal{U} , A target subset/slice where we want to improve accuracy: \mathcal{T} , Loss function \mathcal{L} for learning

- 1: Train model with loss \mathcal{L} on labeled set \mathcal{E} and obtain parameters θ_E
 - 2: Compute the gradients $\{\nabla_{\theta_E} \mathcal{L}(x_i, y_i), i \in \mathcal{U}\}$ (using hypothesized labels) and $\{\nabla_{\theta_E} \mathcal{L}(x_i, y_i), i \in \mathcal{T}\}$.
 - 3: Compute the similarity kernels S (this includes kernel of the elements within \mathcal{U} , within \mathcal{T} and between \mathcal{U} and \mathcal{T}) and define a submodular function f and diversity function g
 - 4: $\hat{\mathcal{A}} \leftarrow \max_{\mathcal{A} \subseteq \mathcal{U}, |\mathcal{A}| \leq k} I_f(\mathcal{A}; \mathcal{T}) + \gamma g(\mathcal{A})$
 - 5: Obtain the labels of the elements in \mathcal{A}^* : $L(\hat{\mathcal{A}})$
 - 6: Train a model on the combined labeled set $\mathcal{E} \cup L(\hat{\mathcal{A}})$
-

4 EFFECTIVENESS OF SMI FOR TSS

Dataset, Baselines and Implementation details: We demonstrate the effectiveness of TSS in obtaining a targeted subset for improving image classification accuracy for some target classes on CIFAR-10 and MNIST datasets. To simulate a real-world setting, we split the available train set further into train, validate and a data lake such that (i) the train set has few labeled instances and poorly represents two randomly picked classes (target), and (ii) data lake is a large set whose labels we do not use (resembling a large pool of unlabeled data in real-world). The poorly represented classes do not perform well on the validation set and hold clue to picking up the target of interest. Performance is measured on the test set from the respective datasets. We then apply TSS (Algorithm 1) comparing MI functions with other existing approaches. Specifically, for MI functions we use LogDetMI, GCMI, FL1MI, FL2MI, and GCMI + Diversity (equivalent to an intuitive approach of minimizing average gradient difference with the target) For existing approaches, we compare with three active learning baselines (uncertainty sampling (US), BADGE, and GLISTER-ACTIVE (GLISTER)) running them only once as per our setting (i.e. we select the unlabeled subset only once). Since these active learning baselines do not explicitly have information of the target set, to further strengthen them we also compare against two variants which are target-aware. The first is ‘targeted uncertainty sampling’ (TUS) where a product of the uncertainty and the similarity with the target is used to identify the subset, and second is GLISTER-TSS where the target set is used in the bi-level optimization. Finally, we also compare with pure diversity/representation functions (Facility Location (FL), Graph Cut (GC), Log Determinant (LogDet), Disparity-Sum (DSUM)) and random sampling. We train the model (ResNet-18 (He et al., 2016) for CIFAR-10, LeNet (LeCun et al., 1989) for MNIST) using cross-entropy loss and SGD optimizer until training accuracy exceeds 99% (Base model). After augmenting the train set with the labeled version of the selected subset and re-training the model, we report the average gain in accuracy for the target classes and overall gain in accuracy across all classes, averaged across 10 runs of randomly picking any two classes as target. We run TSS for different budgets and also study the effect of budget on the performance. Wherever applicable, we keep the internal parameters at their default values of 1.

Results: In Table 1, we report the results for a budget of 400 for CIFAR-10 and 70 for MNIST. To keep the setting as realistic as possible, we set the target set to be much smaller than the budget (around 10% of the budget – 44 for CIFAR-10 and 6 for MNIST). We report the effect of budget on the gain in accuracy of the target classes in Fig. 2. On both datasets, MI functions yield the best

improvement in accuracy on the target classes ($\approx 30\%$ gain over the model’s performance before re-training with added targeted subset; $\approx 12\%$ more than other methods) while also simultaneously increasing the overall accuracy by $\approx 6\%$. They consistently outperform BADGE, GLISTER-TSS, US and TUS across all budgets. Since the SMI functions (LogDetMI, FL2MI and GCMI+DIV) model both query-relevance and diversity, they perform better than both a) functions which tend to prefer relevance (GCMI, TUS) and b) functions which tend to prefer diversity/representation (BADGE, FL, GC, DSUM, LogDet). Also, we observe that, as the budget is increased, the MI functions outperform other methods by greater margins on the target class accuracy (Fig. 2). This is expected, as other methods are not effective in considering the target.




Figure 2: Comparison of different methods for targeted subset selection for different budgets on CIFAR-10 and MNIST. X-axis: budgets, Y-axis: gain in model accuracy for target classes. MI based approaches (lines in red) significantly outperform others across all subset sizes. (Section 4).

Method	CIFAR-10		MNIST	
	Target	Overall	Target	Overall
Base	11.2	52.34	52.76	88.26
Random	+2.75	+1.43	+1.08	-0.03
BADGE (Ash et al., 2020)	+8.8	+0.66	+6.7	+1.66
GLISTER (Killamsetty et al., 2020)	+14.15	+1.25	+18.36*	+2.56
GLISTER-TSS	+18.3*	+1.53	+11.19	+2.2
US (Settles, 2009)	+3.95	+2.03	+7.56	+1.18
TUS	+10.45	+2.99*	+6.21	+1.61
LogDet	+12.3	+2.8	+14.2	+2.96
FL	+6.8	+1.14	+17.05	+3.4*
GC	-1.3	-1.1	+14.97	+3.26
DSUM	+3.8	+1.66	+14.39	+2.79
LogDetMI	+29.05	+4.95	+32.97	+6.84
FL2MI	+33.7	+6.59	+33.57	+6.69
FL1MI	+17.65	+4.2	+27.51	+5.57
GCMI	+26.55	+4.5	+30.79	+6.12
GCMI+DIV	+31.65	+6.01	+34.59	+7.10

Table 1: Comparison of TSS (MI functions) with other methods for a budget of 400 (CIFAR-10) and 70 (MNIST). The numbers are the gain in % accuracy of the target classes (Target) and all classes (Overall) over the Base model after re-training the model (see text). Best among existing approaches is indicated with *, highest in blue, 2nd and 3rd highest in red and green respectively.

5 CONCLUSION

We demonstrate the effectiveness of SMI functions for improving a model’s performance by augmenting the training data with samples that match a target distribution (targeted data subset selection). Through experiments on CIFAR-10 and MNIST datasets, we empirically verify the superiority of SMI functions over existing methods.

REFERENCES

- Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *ICLR*, 2020.
- Satoru Fujishige. *Submodular functions and optimization*. Elsevier, 2005.
- Anupam Gupta and Roie Levin. The online submodular cover problem. In *ACM-SIAM Symposium on Discrete Algorithms*, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Rishabh Iyer, Ninad Khargoankar, Jeff Bilmes, and Himanshu Asnani. Submodular combinatorial information measures with applications in machine learning. *arXiv preprint arXiv:2006.15412*, 2020.
- Rishabh Krishnan Iyer. *Submodular optimization and machine learning: Theoretical results, unifying and scalable algorithms, and applications*. PhD thesis, 2015.
- Vishal Kaushal, Suraj Kothawade, Ganesh Ramakrishnan, Jeff Bilmes, Himanshu Asnani, and Rishabh Iyer. A unified framework for generic, query-focused, privacy preserving and update summarization using submodular information measures. *arXiv preprint arXiv:2010.05631*, 2020.
- Krishnateja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, and Rishabh Iyer. Glister: Generalization based data subset selection for efficient and robust learning. *arXiv preprint arXiv:2012.10630*, 2020.
- Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- Jingxuan Li, Lei Li, and Tao Li. Multi-document summarization via submodularity. *Applied Intelligence*, 37(3):420–430, 2012.
- Hui Lin. *Submodularity in natural language processing: algorithms and applications*. PhD thesis, 2012.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.
- Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- Arun Balajee Vasudevan, Michael Gygli, Anna Volokitin, and Luc Van Gool. Query-adaptive video summarization via quality-aware relevance estimation. In *Proceedings of the 25th ACM international conference on Multimedia*, pp. 582–590, 2017.
- Zhen Zhang and Raymond W Yeung. On characterization of entropy function via information inequalities. *Information Theory, IEEE Transactions on*, 44(4):1440–1452, 1998.