
Active Testing: Sample-Efficient Model Evaluation

Jannik Kossen^{*1} Sebastian Farquhar^{*1} Yarin Gal¹ Tom Rainforth²

Abstract

We introduce a new framework for sample-efficient model evaluation that we call *active testing*. While approaches like active learning reduce the number of labels needed for model *training*, existing literature largely ignores the cost of labeling *test* data, typically unrealistically assuming large test sets for model evaluation. This creates a disconnect to real applications, where test labels are important and just as expensive, e.g. for optimizing hyperparameters. Active testing addresses this by carefully selecting the test points to label, ensuring model evaluation is sample-efficient. To this end, we derive theoretically-grounded and intuitive acquisition strategies that are specifically tailored to the goals of active testing, noting these are distinct to those of active learning. As actively selecting labels introduces a bias; we further show how to remove this bias while reducing the variance of the estimator at the same time. Active testing is easy to implement and can be applied to any supervised machine learning method. We demonstrate its effectiveness on models including WideResNets and Gaussian processes on datasets including Fashion-MNIST and CIFAR-100.

1. Introduction

Although unlabeled datapoints are often plentiful, labels can be expensive. For example, in scientific applications acquiring a single label can require expert researchers and weeks of lab time. However, some labels are more informative than others. In principle, this means that we can pick the most useful points to spend our budget wisely.

These ideas have motivated extensive research into actively selecting *training* labels (Atlas et al., 1990; Settles, 2010), but the cost of labeling *test* data has been largely ignored

^{*}Equal contribution ¹OATML, Department of Computer Science, ²Department of Statistics, Oxford. Correspondence to: Jannik Kossen <jannik.kossen@cs.ox.ac.uk>.




Figure 1. **Active testing** estimates the test loss much more precisely than **uniform sampling** for the same number of labeled test points. Active data selection during testing resolves a major barrier to sample-efficient machine learning, complementing prior work which has focused only on training. For details, see §5.1.

(Lowell et al., 2019). In artificial research settings, this is often not a problem: we can ‘cheat’ by using enormous test datasets if the goal is to see how good some sample-efficient training approach is. But for practitioners this creates a huge issue: in practice, one must evaluate model performance, both to choose the best model and to develop trust in individual models. Whenever labels are expensive enough that we need to carefully pick training data, we cannot afford to be wasteful with test data either.

To address this, we introduce a framework for actively selecting test points for efficient labeling that we call *active testing*. To this end, we derive acquisition functions with which to select test points to maximize the accuracy of the resulting empirical risk estimate. We find that the principles that make these acquisition functions effective are quite different from their active learning counterparts. Given a fixed budget, we can then estimate quantities like the test loss much more accurately than naively labeling points at random. An example of this is given in Fig. 1.

Starting with an idealized, but intractable, approach, we show how practical acquisition strategies for active testing can be derived. In particular, we derive specific principled acquisition frameworks for classification and regression. Each of these depends only on a predictive model for outputs, allowing for substantial customization of the framework to particular problems and computational

budgets. We realize the flexibility of this framework by showing how one can implement fast, simple, and surprisingly effective methods that rely only on the original model, as well as much more powerful techniques which use a ‘surrogate’ model that captures information from already acquired test labels and accounts for errors and potential overconfidence in the original model.

A difficulty in active testing is that choosing test data using information from the training data or the model being evaluated creates a sample-selection bias (MacKay, 1992; Dasgupta & Hsu, 2008). For example, acquiring points where the model is least certain (Houlsby et al., 2011) will likely overestimate the test loss: the least certain points will tend to be harder than average. Moreover, the effect will be stronger for overconfident models, undermining our ability to select models or optimize hyperparameters. We show how to remove this bias using the weighting scheme introduced by Farquhar et al. (2021). The recency of this technique is perhaps why the extensive literature on active learning has neglected actively selecting test data; this bias is far more harmful for testing than it is for training.

Our approach is general and applies to practically any machine learning model and task, not just settings where active learning is used. We show active testing of standard neural networks, Bayesian neural networks, Gaussian processes, and random forests from toy regression data to image classification problems like CIFAR-100. While our acquisition strategies provide a starting point for the field, we expect there to be considerable room for further innovation in active testing, much like the vast array of approaches developed for active learning.

In summary, our main contributions are:

- We formalize active testing as a framework for sample-efficient model evaluation (§2).
- We derive principled acquisition strategies for regression and classification (§3).
- We empirically show that our method yields sample-efficient and unbiased model evaluations, significantly reducing the number of labels required for testing (§5).

2. Active Testing

In this section, we introduce the active testing framework. For now, we set aside the question of how to design the acquisition scheme itself, which is covered in §3. We start with a model which we wish to evaluate, $f : \mathcal{X} \rightarrow \mathcal{Y}$, with inputs $\mathbf{x} \in \mathcal{X}$. Note that f is *fixed as given*: we will evaluate it, and it will not change during evaluation. We make very few assumptions about the model. It could be parametric/non-parametric, stochastic/deterministic and could be applied to any supervised machine learning task.

Our goal is to estimate some model evaluation statistic in a sample-efficient way. Depending on the setting, this could be a test accuracy, mean squared error, negative log-likelihood, or something else. For sake of generality, we can write the ‘test loss’ for arbitrary loss functions \mathcal{L} evaluated over a test set $\mathcal{D}_{\text{test}}$ of size N as

$$\hat{R} = \frac{1}{N} \sum_{i_n \in \mathcal{D}_{\text{test}}} \mathcal{L}(f(\mathbf{x}_{i_n}), y_{i_n}). \quad (1)$$

This test loss is an unbiased estimate for the true risk $R = \mathbb{E}[\mathcal{L}(f(\mathbf{x}), y)]$ and is what we would be able to calculate if we possessed labels for every point in the test set. However, for active testing, we cannot afford to label all the test data. Instead, we can label only a subset $\mathcal{D}_{\text{test}}^{\text{observed}} \subseteq \mathcal{D}_{\text{test}}$. Although we could choose all elements of $\mathcal{D}_{\text{test}}^{\text{observed}}$ in a single step, doing so is sub-optimal as information garnered from previously acquired labels can be used to select future points more effectively. Thus, we will pre-emptively introduce an index m tracking the labeling order: at each step m , we acquire the label y_{i_m} for the point with index i_m and add this point to $\mathcal{D}_{\text{test}}^{\text{observed}}$.

2.1. A Naive Baseline

Standard practice ignores the cost of labeling the test set—it does not actively pick the test points. For a labeling budget M , this is equivalent to uniformly sampling a subset of the test data and then calculating the subsample empirical risk

$$\hat{R}_{\text{iid}} = \frac{1}{M} \sum_{i_m \in \mathcal{D}_{\text{test}}^{\text{observed}}} \mathcal{L}(f(\mathbf{x}_{i_m}), y_{i_m}). \quad (2)$$

Uniform sampling guarantees the data are independently and identically distributed (i.i.d.) so that the estimate is unbiased, $\mathbb{E}[\hat{R}_{\text{iid}}] = \hat{R}$, and converges to the empirical test risk, $\hat{R}_{\text{iid}} \rightarrow \hat{R}$ as $M \rightarrow N$. However, although this estimator is unbiased its *variance* can be high in the typical setting of $M \ll N$. That is, on any given run the test loss estimated according to this method might be very different from \hat{R} , even though they will be equal in expectation.

2.2. Actively Sampling Test Points

To improve on this naive baseline, we need to reduce the variance of the estimator. A key idea of this work is that this can be done by *actively selecting* the most useful test points to label. Unfortunately, doing this naively will introduce unwanted and highly problematic bias into our estimates.

In the context of pool-based active learning, Farquhar et al. (2021) showed that biases from active selection can be corrected by using a stochastic acquisition process and formulating an importance sampling estimator. Namely, they introduce an *acquisition distribution* $q(i_m)$ that denotes the probability of selecting index i_m to be labeled. They then compute the Monte Carlo estimator \hat{R}_{LURE} which, in

active testing setting, takes the form

$$\hat{R}_{\text{LURE}} = \frac{1}{M} \sum_{m=1}^M v_m \mathcal{L}(f(\mathbf{x}_{i_m}), y_{i_m}), \quad (3)$$

where M is the size of $\mathcal{D}_{\text{test}}^{\text{observed}}$, N is the size of $\mathcal{D}_{\text{test}}$, and

$$v_m = 1 + \frac{N - M}{N - m} \left(\frac{1}{(N - m + 1)q(i_m)} - 1 \right). \quad (4)$$

Not only does \hat{R}_{LURE} correct the bias of active sampling, if the proposal $q(i_m)$ is suitable it can also (drastically) reduce the variance of the resulting risk estimates compared to both \hat{R}_{iid} as well as naively applying active sampling without bias correction. This is because \hat{R}_{LURE} is based on importance sampling: a technique designed precisely for reducing variance through appropriate proposals (Kahn & Marshall, 1953; Kahn, 1955; Owen, 2013).

Importantly, there are no restrictions on how $q(i_m)$ can depend on the data, and in our context $q(i_m)$ is actually shorthand for $q(i_m; i_{1:m-1}, \mathcal{D}_{\text{test}}, \mathcal{D}_{\text{train}})$. This means that we will be able use proposals that depend on the already acquired test data, as well as the training data and the trained model, as we explain in the next section.

3. Acquisition Functions for Active Testing

In the last section, we showed how to construct an unbiased estimator of the test risk using actively sampled test data. This is exactly the quantity that the practitioner cares about for evaluating a model. For an estimator to be sample-efficient, its variance should be as small as possible for any given number of labels, M . We now use this principle to derive acquisition proposal distributions (i.e. acquisition functions) for active testing by constructing an idealized proposal and then showing how it can be approximated.

3.1. General Framework

As shown by Farquhar et al. (2021), the optimal oracle proposal for \hat{R}_{LURE} is to sample in proportion to the true loss of each data point, resulting in a single-sample zero-variance estimate of the risk. In practice, we cannot know the true loss before we have access to the actual label. In particular, the true distribution of outputs for a given input is typically noisy and we can never know this noise without evaluating the label. In the context of deriving an unbiased Monte Carlo estimator, the best we can ever hope to achieve is to sample from the expected loss over the true $y \mid \mathbf{x}_{i_m}$,¹

$$q^*(i_m) \propto \mathbb{E}_{p(y|\mathbf{x}_{i_m})} [\mathcal{L}(f(\mathbf{x}_{i_m}), y)]. \quad (5)$$

Note that as i_m can only take on a finite set of values, the required normalization can be performed straightforwardly.

¹For estimators other than \hat{R}_{LURE} , e.g. ones based on quadrature schemes or direct regression of the loss, this may no longer be true.

Of course, $q^*(i_m)$ remains intractable because we do not know the true distribution for $y \mid \mathbf{x}_{i_m}$. We need to approximate it for unlabeled \mathbf{x}_{i_m} in a way that captures regions where $f(\mathbf{x})$ is a poor predictive model as these will contribute the most to the loss. This can be hard as $f(\mathbf{x})$ itself has already been designed to approximate $y \mid \mathbf{x}$.

Thankfully, we have the following tools at our disposal to deal with this: (a) We can incorporate *uncertainty* to identify regions with a lack of available information (e.g. regions far from any of the training data); (b) We can introduce *diversity* in our predictions compared to $f(\mathbf{x})$ (thereby ensuring that mistakes we make are as distinct as possible to those of $f(\mathbf{x})$); and (c) as we label new points in the test set, we can obtain *more accurate* predictions than $f(\mathbf{x})$ by incorporating these additional points. These essential strategies will help us identify regions where $f(\mathbf{x})$ provides a poor fit. We give examples of how we incorporate them in practice in §3.4.

We now introduce a general framework for approximating $q^*(i_m)$ that allows us to use these mechanisms as best as possible. The starting point for this is to consider the concept of a *surrogate* for $y \mid \mathbf{x}$, where we introduce some potentially infinite set of parameters θ , a corresponding generative model $\pi(\theta)\pi(y|\mathbf{x}, \theta)$, and then approximate the true $p(y|\mathbf{x})$ using the marginal distribution $\pi(y|\mathbf{x}) = \mathbb{E}_{\pi(\theta)} [\pi(y|\mathbf{x}, \theta)]$ of the surrogate. We can now approximate $q^*(i_m)$ as

$$q(i_m) \propto \mathbb{E}_{\pi(\theta)\pi(y|\mathbf{x}_{i_m}, \theta)} [\mathcal{L}(f(\mathbf{x}_{i_m}), y)]. \quad (6)$$

With θ we represent our subjective uncertainty over the outcomes in a principled way. However, our derivations will lead to acquisition strategies also compatible with discriminative, rather than generative, surrogates, for which θ will be implicit.

3.2. Illustrative Example

Figure 2 shows how active testing chooses the **next** test point among all available **test** data. The model, here a Gaussian process (Rasmussen, 2003), has been trained using the **training** data and we have already acquired some test points (**crosses**). Figure 2 (b) shows the **true** loss known only to an oracle. Our **approximate** expected loss is a good proxy in some parts of the input space and worse elsewhere. The next point is selected by sampling proportionately to the approximate expected loss. In this example, the surrogate is a Gaussian process that is retrained whenever new labels are observed. The closer the approximate expected loss is to the true loss, the lower the variance of the estimator \hat{R}_{LURE} will be; the estimator will always be unbiased.

3.3. Deriving Acquisition Functions

We now give principled derivations leading to acquisition functions for a variety of widely-used loss functions.




Figure 2. Illustration of a single active testing step. (a) The model has been trained on five points and we currently have observed four test points. (b) We assign acquisition probabilities using the estimated loss of potential test points. Because we do not have access to the true labels, these estimates are different from the true loss. Our next acquisition is then sampled from this distribution.

Regression. Substituting the squared error loss $\mathcal{L}(f(\mathbf{x}), y) = (f(\mathbf{x}) - y)^2$ into (6) yields

$$q(i_m) \propto \mathbb{E}_{\pi(y|\mathbf{x}_{i_m})} [(f(\mathbf{x}_{i_m}) - y)^2], \quad (7)$$

and if we apply a bias-variance decomposition this becomes

$$q(i_m) \propto \underbrace{(f(\mathbf{x}_{i_m}) - \mathbb{E}_{\pi(y|\mathbf{x}_{i_m})} [y])^2}_{\textcircled{1}} + \underbrace{\mathbb{V}_{\pi(y|\mathbf{x}_{i_m})} [y]}_{\textcircled{2}}. \quad (8)$$

Here, $\textcircled{1}$ is the squared difference between our model prediction $f(\mathbf{x}_{i_m})$ and the mean prediction of the surrogate: it measures how wrong the surrogate believes the prediction to be. $\textcircled{2}$ is the predictive variance: the uncertainty that the surrogate has about y at \mathbf{x}_{i_m} .

Both $\textcircled{1}$ and $\textcircled{2}$ are readily accessible in models such as Gaussian processes or Bayesian neural networks. However, we actually do not need an explicit $\pi(\theta)$ to acquire with (8): we only need to provide approximations for the mean prediction $\mathbb{E}_{\pi(y|\mathbf{x}_{i_m})} [y]$ and predictive variance $\mathbb{V}_{\pi(y|\mathbf{x}_{i_m})} [y]$. For example, these exist for ‘deep ensembles’ (Lakshminarayanan et al., 2017) that compute mean and variance predictions from a set of standard neural networks.

A critical subtlety to appreciate is that $\textcircled{2}$ incorporates both aleatoric and epistemic uncertainty (Kendall & Gal, 2017). It is not our estimate for the level of noise in $y|\mathbf{x}_{i_m}$ but the variance of our subjective beliefs for what the value of y could be at \mathbf{x}_{i_m} . This is perhaps easiest to see by noting that

$$\mathbb{V}_{\pi(y|\mathbf{x}_{i_m})} [y] = \mathbb{V}_{\pi(\theta)} [\mathbb{E}_{\pi(y|\mathbf{x}_{i_m}, \theta)} [y]] + \mathbb{E}_{\pi(\theta)} [\mathbb{V}_{\pi(y|\mathbf{x}_{i_m}, \theta)} [y]], \quad (9)$$

where the first term is the variance in our mean prediction and represents our epistemic uncertainty and the latter is our mean prediction of the ‘aleatoric variance’ (label noise). This is why our construction using θ is crucial: it stresses that (8) should also take epistemic uncertainty into account.

For regression models with Gaussian outputs $\mathcal{N}(f(\mathbf{x}), \sigma^2)$, the negative log-likelihood loss function and the squared error are related by affine transformation—and following (6) so are the acquisition functions.

Classification. For classification, predictions $f(\mathbf{x})$ generally take the form of conditional probabilities over outcomes $y \in \{1, \dots, C\}$. First, we study cross-entropy

$$\mathcal{L}(f(\mathbf{x}), y) = -\log f(\mathbf{x})_y. \quad (10)$$

Here, we again introduce a surrogate, and, using (6), obtain

$$q(i_m) \propto \mathbb{E}_{\pi(y|\mathbf{x}_{i_m})} [-\log f(\mathbf{x}_{i_m})_y]. \quad (11)$$

Now expanding the expectation over y yields

$$q(i_m) \propto -\sum_y \pi(y | \mathbf{x}_{i_m}) \log f(\mathbf{x}_{i_m})_y, \quad (12)$$

which is the cross-entropy between the marginal predictive distribution of our surrogate, $\pi(y | \mathbf{x}_{i_m})$, and our model.

We can also derive acquisition strategies based on accuracy. Namely, writing one minus accuracy to obtain a loss,

$$\mathcal{L}(f(\mathbf{x}), y) = 1 - \mathbb{1}[y = \arg \max_{y'} f(\mathbf{x})_{y'}], \quad (13)$$

and substituting into (6) yields

$$q(i_m) \propto 1 - \pi(y = y^*(\mathbf{x}_{i_m}) | \mathbf{x}_{i_m}), \quad (14)$$

where $y^*(\mathbf{x}_{i_m}) = \arg \max_{y'} f(\mathbf{x}_{i_m})_{y'}$.

3.4. Tactics for Obtaining Good Surrogates

In §3.1 we introduced three ways for the surrogate to assist in finding high-loss regions of $f(\mathbf{x})$: we want it to (a) account for uncertainty over the outcomes, (b) make predictions that are diverse to $f(\mathbf{x})$, and (c) incorporate information from all available data. Motivated by this, we apply the following tactics to obtain good surrogates:

Uncertainty. We should use surrogates that incorporate both epistemic and aleatoric uncertainty effectively, and further ensure that these are well-calibrated. Capturing epistemic uncertainty is essential to predicting regions of high loss, while aleatoric uncertainty still contributes and cannot be ignored, particularly if heteroscedastic. A variety of different approaches can be effective in this regard and thus provide successful surrogates. For example, Bayesian neural networks, deep ensembles, and Gaussian processes.

Fidelity. In real-world settings, f may be constrained to be memory-efficient, fast, or interpretable. If labels are

Algorithm 1 Active Testing

Input: Model f trained on data $\mathcal{D}_{\text{train}}$

- 1: Train surrogate π & choose acquisition proposal form q
- 2: **for** $m = 1$ to M **do**
- 3: $i_m \sim q(i_m; \pi)$, observe y_{i_m} , add to $\mathcal{D}_{\text{test}}^{\text{observed}}$
- 4: Calculate $\mathcal{L}(f(\mathbf{x}_{i_m}), y_{i_m})$ and v_m \triangleright Eq. (4)
- 5: Update π , e.g. retraining on $\mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{test}}^{\text{observed}}$
- 6: **end for**
- 7: Return \hat{R}_{LURE} \triangleright Eq. (3)

expensive enough, we can relax these constraints at test time and construct a more capable surrogate. In fact, we practically find that using an ensemble of models like f is a robust way of achieving sample-efficiency.

Diversity. By choosing the surrogate from a different model family or adjusting its hyperparameters, we can decorrelate the errors of the surrogate and f , resulting in better exploration. For example, we find that random forests (Breiman, 2001) can help evaluate neural networks.

Extra data. If our computational constraints are not critical, we should retrain the surrogate on $\mathcal{D}_{\text{test}}^{\text{observed}} \cup \mathcal{D}_{\text{train}}$ after each step. The exposure to additional data will make the surrogate a better approximation of the true outcomes.

Thompson-Ensemble. Retraining the surrogate can also create diversity in predictions due to stochasticity in the training process, the addition of new data, or even deliberate randomization. In fact, we can view retraining the surrogate at regular intervals as implicitly defining an ensemble of surrogates, with the surrogate used at any given iteration forming a Thompson-sample (Thompson, 1933) from this ensemble. This will generally be more powerful and more diverse than a single surrogate, providing further motivation for retraining and potentially even deliberate variations in surrogates/hyperparameters between iterations.

In §5 we empirically assess the relative importance of these considerations, which depends heavily on the situation. For example, the benefit of retraining using the labels acquired at test-time is especially large in very low-data settings, while the benefit of ensembling can be large even when there is more data available. Putting everything together, Algorithm 1 provides a summary of our general framework.

If compute is at a premium for acquisitions, a simple alternative heuristic is to use our original model for the surrogate. This avoids learning a new predictive model, but it suffers because now the surrogate can never disagree with f . Instead, we have to rely entirely on uncertainties for approximating (6): for regression, (1) in (8) is zero, and for classification, (12) reduces to the predictive entropy. In general, we do not recommend this strategy, *unless* computational constraints are substantial *and* there is reason to believe that the epistemic and aleatoric uncertainties from

Why are acquisition strategies different for active learning and active testing?

Researchers have already investigated acquisition functions for active learning, and it would be helpful if we could just apply these here. However, active testing is a different problem conceptually because we are not trying to use the data to fit a model.

First, popular approaches for active learning avoid areas with high aleatoric uncertainty while seeking out high epistemic uncertainty. This motivates acquisition functions like BALD (Houlsby et al., 2011) or BatchBALD (Kirsch et al., 2019). For active testing, however, areas of high aleatoric uncertainty can be critical to the estimate.

Second, as Imberg et al. (2020) point out, the optimal acquisition scheme for active learning will minimize the expected generalization error at the end of training. They show how this motivates additional terms beyond what one would get from minimizing the variance of the loss estimator.

Third, as Farquhar et al. (2021) show, a biased loss estimator can be helpful during training because it often partially cancels the natural bias of the training loss. This is no longer true at test-time, where we want to minimize bias as much as possible.

f represent the true loss well. If the latter is true, this simplistic approach can perform surprisingly well, although it is always outperformed by more complex strategies. In particular, training a single fixed surrogate that is distinct from f will still typically provide noticeable benefits.

4. Related Work

Efficient use of labels is a major aim in machine learning and it is often important to use large pools of unlabeled data through unsupervised or semi-supervised methods (Chapelle et al., 2009; Erhan et al., 2010; Kingma et al., 2014).

An even more efficient strategy is to collect only data that is likely to be particularly informative in the first place. Such approaches are known as optimal or adaptive *experimental design* (Lindley, 1956; Chaloner & Verdinelli, 1995; Sebastiani & Wynn, 2000; Foster et al., 2020; 2021) and are typically formalized through optimizing the (expected) information gained during an experiment.

Perhaps the best-known instance of adaptive experimental design is *active learning*, wherein the designs to be chosen are the data points for which to acquire labels (Atlas et al., 1990; Settles, 2010; Houlsby et al., 2011; Sener & Savarese, 2018). This is typically done by optimizing, or sampling

from, an acquisition function, with much discussion in the literature on the form this should take (Imberg et al., 2020).

What most of this work neglects is the wasteful acquisition of data for *testing*. Lowell et al. (2019) acknowledge this and describe it as a major barrier to the adoption of active learning methods in practice. The potential for ‘active testing’ was raised by Nguyen et al. (2018), but they focused on the special case of noisily annotated labels that must be vetted and did not acknowledge the substantial bias that their method introduces. Farquhar et al. (2021) introduce the variance-reducing unbiased estimator for active sampling which we apply. However, their focus is mostly on correcting the bias of active *learning* (Bach, 2007; Sugiyama, 2006; Beygelzimer et al., 2009; Ganti & Gray, 2012) and they do not consider appropriate acquisition strategies for active testing. Note that their theoretical results about the properties of \hat{R}_{LURE} carry over to active setting.

Other methods like Bayesian Quadrature (Rasmussen & Ghahramani, 2003; Osborne, 2010) and kernel herding (Chen et al., 2012) can also sometimes employ active selection of points. Of particular note, Osborne et al. (2012); Chai et al. (2019) study active learning of model evidence in the context of Bayesian Quadrature.

Bennett & Carvalho (2010); Katariya et al. (2012); Kumar & Raj (2018); Ji et al. (2021) explore the efficient evaluation of classifiers based on stratification, rather than active selection of individual labels. Namely, they divide the test pool into strata according to simple metrics such as classifier confidence. Test data are then acquired by first sampling a stratum and then selecting data *uniformly* within. Sample-efficiency for these approaches could be improved by performing active testing within the strata. Sawade et al. (2010) similarly explore active risk estimation through importance sampling, but rely on sampling with replacement which is suboptimal in pool-based settings (see Appendix D). Moreover, like the other aforementioned works, they do not consider the use of surrogates to allow for more effective acquisition strategies.

5. Empirical Investigation

We now assess the empirical performance of active testing and investigate the relative merits of different strategies. Similar to active learning, we assume a setting where sample acquisition is expensive, and therefore, per-sample efficiency is critical. Full details as well as additional results are provided in the appendix, and we release code for reproducing the results at github.com/jlko/active-testing.

We note a small but important practicality: we ensure all points have a minimum proposal probability regardless of the acquisition function value, to ensure that the weights are bounded even if q is badly calibrated (cf. Appendix B.1).

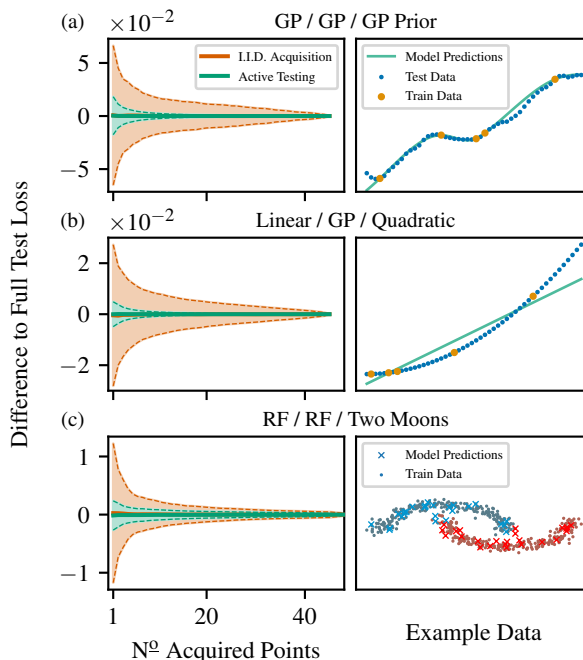


Figure 3. Active testing yields unbiased estimates of the test loss with significantly reduced variance. Each row shows a different combination of *model/surrogate/data*. GP is short for Gaussian process, RF for random forest. The first column displays the mean difference of the estimators to the true loss on the full test set (known only to an oracle). We retrain surrogates after each acquisition on all observed data. Shading indicates standard deviation over 5000 (a-b) / 2500 (c) runs; data is randomized between runs. The second column shows example data with *model predictions*, and the points used for *training* and *testing* (a-b).

5.1. Synthetic Data

We first show that active testing on synthetic datasets offers sample-efficient model evaluations. By way of example, we actively evaluate a Gaussian process (Rasmussen & Ghahramani, 2003) and a linear model for regression, and a random forest (Breiman, 2001) for classification. For regression, we estimate the squared error and acquire test labels via Eq. (8); for classification, we estimate the cross-entropy loss and acquire with Eq. (12). We use Gaussian process and random forest surrogates that are retrained on all observed data after each acquisition.

Figure 3 shows how the difference between our test loss estimation and the truth (known only to an oracle) is much smaller than the naive \hat{R}_{id} : active testing allows us to precisely estimate the empirical test loss using far fewer samples. For example, after acquiring labels for only 5 test points in (a), the standard deviation of active testing is already as low as it is for i.i.d. acquisition at step 40, nearly the entire test set. Further, we can see that the estimates of \hat{R}_{LURE} are indeed unbiased. Appendix A.1 gives experiments on additional synthetic datasets.

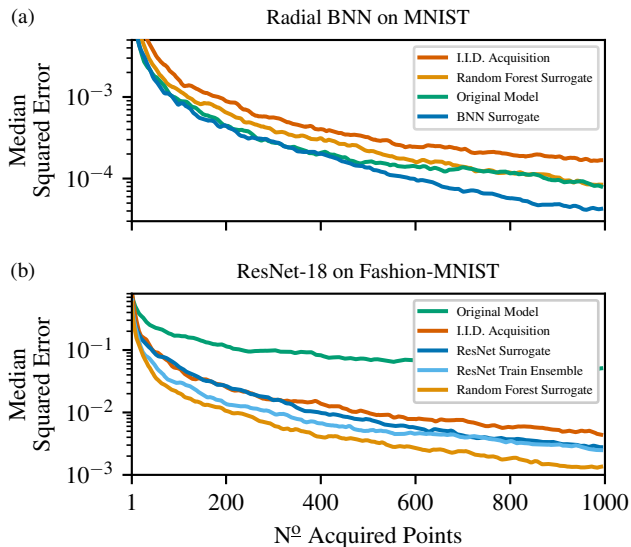


Figure 4. Median squared errors for (a) Radial BNN on **MNIST** and (b) ResNet-18 on **Fashion-MNIST** in a small-data setting. **Original Model** samples proportional to predictive entropy, **X Surrogate** iteratively retrains a surrogate on all observed data, and **ResNet Train Ensemble** is a deep ensemble trained on $\mathcal{D}_{\text{train}}$ once. Lower is better; medians are over 1085 runs for (a), 872 for (b).

Here we have actually acquired the full test set. This lets us show that both \hat{R}_{LURE} and \hat{R}_{iid} converge to the empirical test loss on the entire test set. However, typically we cannot do this which makes the difference in variance between \hat{R}_{iid} and \hat{R}_{LURE} at lower acquisition numbers crucial.

5.2. Surrogate Choice Case Study: Image Classification

We now investigate the impact of the different surrogate choices. For this, we move to more complex image classification tasks and additionally restrict the number of training points to only 250. This makes it harder to predict the true loss. Therefore, the strategies discussed in §3.4 are especially important to maximize sample-efficiency.

We evaluate two model types for this examination. First, a Radial Bayesian Neural Network (Radial BNN) (Farquhar et al., 2020) on the MNIST dataset (LeCun et al., 1998) in Fig. 4 (a). Radial BNNs are a recent approach for variational inference in BNNs (Blundell et al., 2015) and we use them because of their well-calibrated uncertainties. We also evaluate a ResNet-18 (He et al., 2016) trained on Fashion-MNIST (Xiao et al., 2017) in Fig. 4 (b) to investigate active testing with conventional neural network architectures. In these figures, we show the median squared error of the different surrogate strategies on a logarithmic scale to highlight differences between the approaches. While not shown, note that all approaches do still obtain unbiased estimates. We again use Eq. (12) to estimate the cross-entropy loss of the models.

Predictive Entropy. We first consider the most naive of the approaches mentioned in §3.4: using the unchanged **original model** as the surrogate, which leads to acquisitions based on model predictive entropy. For the Radial BNN, this approach already yields improvements over **i.i.d. acquisition** in Fig. 4 (a). The same can not be said for the ResNet in Fig. 4 (b), for which predictive entropy actually performs worse than i.i.d. acquisition. Presumably, this is the case because the standard neural network struggles to model epistemic uncertainty. Now, we progress to more complex surrogates, improving performance over the naive approach.

Retraining. **BNN surrogate** is a surrogate with identical setup as the original model that is retrained on the total observed data, $\mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{test}}^{\text{observed}}$, 12 times with increasingly large gaps. This leads to improved performance over the naive model predictive entropy, especially as more data is acquired. Similarly, the **ResNet surrogate** shows much-improved performance over predictive entropy when regularly retrained, now outperforming i.i.d. acquisition.

Different Model. As discussed in §3.4, it may be beneficial to choose the surrogate from a different model family to promote diversity in its predictions. We use a **random forest** as a surrogate for both Fig. 4 (a) and (b). For the Radial BNN on MNIST, the random forest, while better than i.i.d. acquisition, does not improve over the model predictive entropy. However, for the ResNet on Fashion-MNIST, we find that the random forest surrogate outperforms everything, despite being a cheaper surrogate. This demonstrates that for a surrogate to be successful, it does not necessarily need to be more accurate—although the difference in accuracy is small with so few data. Instead, the surrogate can be also be successful by being *different* from the original model, i.e. having structural biases that lead to it making different predictions and therefore discovering mistakes of the original model, with any new mistakes made less important. Further, if compute is limited, the random forest is attractive because retraining it is much faster.

Ensembling Diversity. §3.4 discussed two ways retraining may help: new data improves the surrogate’s predictive model and repeated training promotes diversity through an implicit ensemble. In Fig. 4 (b), we introduce the **ResNet train ensemble**—a deep ensemble of ResNets trained once on $\mathcal{D}_{\text{train}}$. This surrogate allows us to isolate the effect of predictive diversity since it is not exposed to any test data through retraining. We output mean predictions of the ensemble, and find that the deep ensemble can, a little unexpectedly, outperform the **ResNet surrogate** without accessing the extra data. This is likely because of better calibrated uncertainties and the increased model capacity.

In summary, we have shown that active testing reduces the number of labeled examples required to get a reliable estimate of the test loss for a Radial BNN on MNIST and




Figure 5. Predictive entropy underestimates the true loss of some points by orders of magnitude. Diverse predictions from the ensemble of surrogates help for these crucial high-confidence mistakes, even though they are noisier for low-loss points, improving sample-efficiency overall. (a) We sort values of the true losses and use the index order to plot the approximate losses for predictive entropy (b) and an ensemble of surrogates (c), ideally seeing few small approximated losses on the right. Shown is a ResNet-18 on CIFAR-100; note the log-scale on y and the use of clipping to avoid overly small acquisition probabilities.

ResNet-18 on FashionMNIST in a challenging setting, if appropriate surrogates are chosen.

5.3. Large-Scale Image Classification

We now additionally apply active testing to a Resnet-18 trained on CIFAR-10 (Krizhevsky et al., 2009) and a WideResNet (Zagoruyko & Komodakis, 2016) trained on CIFAR-100. As the complexity of the datasets increases, it becomes harder to estimate the loss, and hence, it is crucial to show that active testing scales to these scenarios. We use conventional training set sizes of 50 000 data points.

In the previous section, we have seen surrogates based on deep ensembles perform well, even if they are only trained once and not exposed to any acquired test labels. For the following experiments, we therefore use these ensembles as surrogates. This is even more justified in the common case where there is much more training data than test data; the extra information in the test labels will typically help less.

In Fig. 5, we further visualize how ensembles increase the quality of the approximated loss in this setting. The original model (b) makes overconfident false predictions with high losses which are rarely detected (box). But the ensemble avoids the majority of these mistakes (c, box) which contribute most to the weighted loss estimator Eq. (3).

In all cases, the active testing estimator has lower median squared error than the baseline, see Fig. 6 (a)—again note the log-scale. We further show in Fig. 6 (b) that using active testing is much more sample-efficient than i.i.d. sampling by calculating the ‘relative labeling cost’: the proportion of actively chosen points needed to get the same performance as naive uniform sampling. E.g., a cost of 0.25 means we need only $1/4$ of actively chosen labels to get equivalently precise test loss. Thus, for the less complex datasets, we see

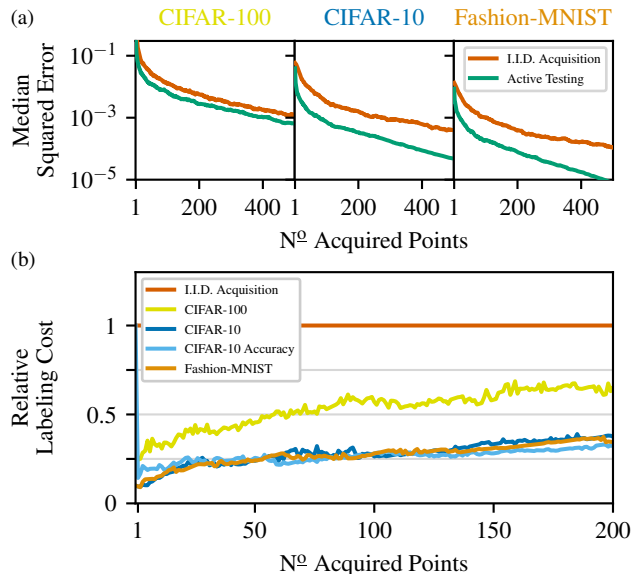


Figure 6. Active testing of a WideResNet on CIFAR-100 and ResNet-18 on CIFAR-10 and Fashion-MNIST. (a) Convergences of errors for active testing/i.i.d. acquisition. (b) Relative effective labeling cost. Active testing consistently improves the sample-efficiency. Lower is better; medians over 1000 random test sets.

efficiency gains are in the region of a factor of four, while for CIFAR-100 they are closer to a factor of two. We also show that there are similar gains in sample-efficiency when estimating accuracy—‘CIFAR-10 Accuracy’ in Fig. 6 (b).

5.4. Diversity and Fidelity in Ensemble Surrogates

We now perform an ablation to study the relative effects of surrogate fidelity and diversity on active testing performance. For this, we evaluate a ResNet-18 trained on CIFAR-10 using different ResNet ensembles as surrogates. Starting with a base surrogate of a single ResNet-18, we increase the size of the ensemble (mainly increasing diversity) as well as the capacity of the layers (increasing fidelity). Given the success of the ‘Train Ensemble’ in §5.2, we train surrogates only once on $\mathcal{D}_{\text{train}}$, rather than retraining as data is acquired.

As Fig. 7 shows, both fidelity and diversity contribute to active testing performance: the best performance is obtained for the most diverse and complex surrogate, justifying our claims in §3.4. We see that increasing fidelity and diversity both individually help performance, with the effect of the latter seeming to be most pronounced (e.g. an ensemble of 5 Resnet-18s outperforms a single ResNet-50).

5.5. Optimal Proposals and Unbiasedness

Fig. 8 (a) confirms our theoretical assumptions by showing that sampling proportional to the true loss, i.e. cheating by asking an oracle for the true outcomes beforehand, does indeed yield exact, single-sample estimates of the loss if

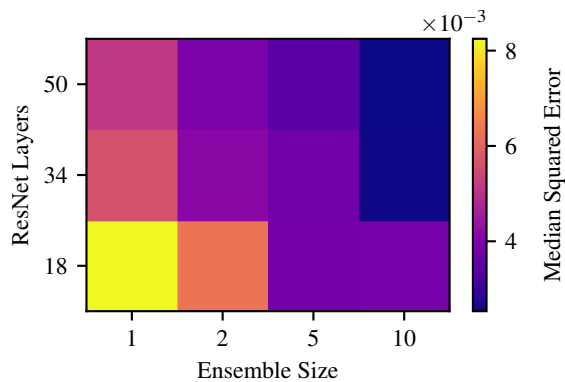


Figure 7. Both diversity and fidelity of the surrogate contribute to sample-efficient active testing. However, the effect of increasing diversity seems larger than that of increased fidelity. We vary the layers (fidelity) and ensemble size (diversity) of the surrogate for active evaluation of a ResNet-18 trained on CIFAR-10. Experiments are repeated for 1000 randomly drawn test sets and we report average values over acquisition steps 100–200.

combined with \hat{R}_{LURE} . Further, it confirms the need for a bias-correcting estimator such as \hat{R}_{LURE} : without it, the risk estimates are biased and clearly overestimate model error.

5.6. Active Testing vs. Active Learning

As mentioned in §3.4, we expect there to be differences in acquisition function requirements for active learning and active testing. For example, mutual information is a popular acquisition function in active learning (Houlsby et al., 2011), but our derivations for classification lead to acquisition strategies based on predictive entropy. Can mutual information also be used for active testing? In Fig. 8 (b) we see that even the simple approach of using the original model as a surrogate and a predictive entropy acquisition outperforms mutual information. Acquiring with mutual information helps active learning because it focuses on uncertainty that can be reduced by more information rather than irreducible noise. While this focus helps learning, it is unhelpful for evaluation where all uncertainty is relevant. This is just one way active testing needs special examination and cannot just re-use results from active learning.

5.7. Practical Advice

Empirically, we find that deep learning ensemble surrogates appear to robustly achieve sample-efficient active testing when using our acquisition strategies. Increases in surrogate fidelity further seem to benefit sample-efficiency.

Active testing generally assumes that acquisitions of labels for samples are expensive, hence we recommend retraining the surrogate whenever new data becomes available. However, if the cost of this is noticeable relative to that of labeling, our results indicate that not retraining the




Figure 8. (a) Naively acquiring proportional to the predictive entropy and using the unweighted estimator \hat{R}_{id} leads to biased estimates with high variance compared to active testing with \hat{R}_{LURE} . Sampling from the unknown true loss distribution would yield unbiased, zero-variance estimates. While this is in practice impossible, the result validates a main theoretical assumption. (b) Mutual information, popular in active learning, underperforms for active testing, even compared to the simple predictive entropy approach. This is because it does not target expected loss. Shown for 692 runs of a Radial BNN on Fashion-MNIST.

surrogates is an option, especially when the number of acquired test labels is small compared to the training data.

In general, we do not recommend the naive strategy that relies entirely on the original model and does not introduce a dedicated surrogate model. As §5.2 has shown, this method can fail to achieve sample-efficient active testing if the original model does not have trustworthy uncertainties. This strategy should remain a last resort and used only when there is significant reason to trust the original model’s uncertainties; we find the diversity provided by a surrogate is critical, even if that surrogate is itself simple.

6. Conclusions

We have introduced the concept of active testing and given principled derivations for acquisition functions suitable for model evaluation. Active testing allows much more precise estimates of test loss and accuracy using fewer data labels.

While our work provides an exciting starting point for active testing, we believe that the underlying idea of sample-efficient evaluation leaves significant scope for further development and alternative approaches. We therefore eagerly anticipate what might be achieved with future work.

Acknowledgements

We acknowledge funding from the New College Yeotown Scholarship (JK) and Oxford CDT in Cyber Security (SF).

References

- Atlas, L., Cohn, D., and Ladner, R. Training connectionist networks with queries and selective sampling. In *Advances in Neural Information Processing Systems*, volume 2, pp. 566–573, 1990.
- Bach, F. Active learning for misspecified generalized linear models. In *Advances in Neural Information Processing Systems*, volume 19, pp. 65–72, 2007.
- Bennett, P. N. and Carvalho, V. R. Online stratified sampling: evaluating classifiers at web-scale. In *International conference on Information and knowledge management*, volume 19, pp. 1581–1584, 2010.
- Beygelzimer, A., Dasgupta, S., and Langford, J. Importance weighted active learning. In *International Conference on Machine Learning*, volume 26, pp. 49–56, 2009.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *International Conference on Machine Learning*, volume 32, pp. 1613–1622, 2015.
- Breiman, L. Random forests. *Machine learning*, 45(1): 5–32, 2001.
- Chai, H., Ton, J.-F., Osborne, M. A., and Garnett, R. Automated model selection with Bayesian quadrature. In *International Conference on Machine Learning*, volume 36, pp. 931–940, 2019.
- Chaloner, K. and Verdinelli, I. Bayesian experimental design: A review. *Statistical Science*, pp. 273–304, 1995.
- Chapelle, O., Scholkopf, B., and Zien, A. Semi-supervised learning. *IEEE Transactions on Neural Networks*, 20(3): 542–542, 2009.
- Chen, Y., Welling, M., and Smola, A. Super-samples from kernel herding. *arXiv preprint arXiv:1203.3472*, 2012.
- Dasgupta, S. and Hsu, D. Hierarchical sampling for active learning. In *International Conference on Machine Learning*, pp. 208–215. ACM Press, 2008.
- DeVries, T. and Taylor, G. W. Improved regularization of convolutional neural networks with cutout. *arXiv:1708.04552*, 2017.
- Erhan, D., Courville, A., Bengio, Y., and Vincent, P. Why does unsupervised pre-training help deep learning? In *International Conference on Artificial Intelligence and Statistics*, volume 13, pp. 201–208, 2010.
- Farquhar, S., Osborne, M. A., and Gal, Y. Radial Bayesian neural networks: Beyond discrete support in large-scale Bayesian deep learning. In *International Conference on Artificial Intelligence and Statistics*, volume 23, pp. 1352–1362, 2020.
- Farquhar, S., Gal, Y., and Rainforth, T. On statistical bias in active learning: How and when to fix it. In *International Conference on Learning Representations*, 2021.
- Foster, A., Jankowiak, M., O’Meara, M., Teh, Y. W., and Rainforth, T. A unified stochastic gradient approach to designing bayesian-optimal experiments. In *International Conference on Artificial Intelligence and Statistics*, pp. 2959–2969. PMLR, 2020.
- Foster, A., Ivanova, D. R., Malik, I., and Rainforth, T. Deep adaptive design: Amortizing sequential bayesian experimental design. In *International Conference on Machine Learning*, 2021.
- Ganti, R. and Gray, A. Upal: Unbiased pool based active learning. *Artificial Intelligence and Statistics*, 15, 2012.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.
- Houlsby, N., Huszár, F., Ghahramani, Z., and Lengyel, M. Bayesian active learning for classification and preference learning. *arXiv:1112.5745*, 2011.
- Imberg, H., Jonasson, J., and Axelson-Fisk, M. Optimal sampling in unbiased active learning. *Artificial Intelligence and Statistics*, 23, 2020.
- Ji, D., Logan, R. L., Smyth, P., and Steyvers, M. Active bayesian assessment of black-box classifiers. In *AAAI Conference on Artificial Intelligence*, volume 35, pp. 7935–7944, 2021.
- Kahn, H. *Use of different Monte Carlo sampling techniques*. Rand Corporation, 1955.
- Kahn, H. and Marshall, A. W. Methods of reducing sample size in monte carlo computations. *Journal of the Operations Research Society of America*, 1:263–278, 1953.
- Katariya, N., Iyer, A., and Sarawagi, S. Active evaluation of classifiers on large datasets. In *International Conference on Data Mining*, volume 12, pp. 329–338, 2012.
- Kendall, A. and Gal, Y. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *Advances In Neural Information Processing Systems*, 30, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.

- Kingma, D. P., Rezende, D. J., Mohamed, S., and Welling, M. Semi-supervised learning with deep generative models. *arXiv:1406.5298*, 2014.
- Kirsch, A., van Amersfoort, J., and Gal, Y. Batchbald: Efficient and diverse batch acquisition for deep Bayesian active learning. In *Advances in Neural Information Processing Systems*, volume 32, pp. 7026–7037, 2019.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images, 2009.
- Kumar, A. and Raj, B. Classifier risk estimation under limited labeling resources. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pp. 3–15, 2018.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, volume 30, pp. 6402–6413, 2017.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lindley, D. V. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pp. 986–1005, 1956.
- Lowell, D., Lipton, Z. C., and Wallace, B. C. Practical Obstacles to Deploying Active Learning. *Empirical Methods in Natural Language Processing*, November 2019.
- MacKay, D. J. C. Information-Based Objective Functions for Active Data Selection. *Neural Computation*, 4(4): 590–604, 1992.
- Nguyen, P., Ramanan, D., and Fowlkes, C. Active testing: An efficient and robust framework for estimating accuracy. In *International Conference on Machine Learning*, volume 37, pp. 3759–3768, 2018.
- Osborne, M., Garnett, R., Ghahramani, Z., Duvenaud, D. K., Roberts, S. J., and Rasmussen, C. Active learning of model evidence using Bayesian quadrature. In *Advances in Neural Information Processing Systems*, volume 25, pp. 46–54, 2012.
- Osborne, M. A. *Bayesian Gaussian processes for sequential prediction, optimisation and quadrature*. PhD thesis, Oxford University, UK, 2010.
- Owen, A. B. *Monte Carlo theory, methods and examples*. 2013.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32, pp. 8024–8035, 2019.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Rasmussen, C. E. Gaussian processes in machine learning. In *Summer school on machine learning*, pp. 63–71. Springer, 2003.
- Rasmussen, C. E. and Ghahramani, Z. Bayesian Monte Carlo. *Advances in neural information processing systems*, pp. 505–512, 2003.
- Sawade, C., Landwehr, N., Bickel, S., and Scheffer, T. Active risk estimation. In *International Conference on Machine Learning*, volume 27, pp. 951–958, 2010.
- Sebastiani, P. and Wynn, H. P. Maximum entropy sampling and optimal Bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1):145–157, 2000.
- Sener, O. and Savarese, S. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018.
- Settles, B. Active Learning Literature Survey. *Machine Learning*, 2010.
- Sugiyama, M. Active learning for misspecified models. In *Advances in Neural Information Processing Systems*, volume 18, pp. 1305–1312, 2006.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. In *British Machine Vision Conference*, pp. 87.1–87.12, 2016.