

Unifying Approaches in Data Subset Selection via Fisher Information and Information-Theoretic Quantities

Anonymous authors

Paper under double-blind review

Abstract

The mutual information between predictions and model parameters—also referred to as *expected information gain* or BALD in machine learning—measures informativeness. It is a popular acquisition function in Bayesian active learning and Bayesian optimal experiment design. In data subset selection, i.e. active learning and active sampling, several recent works use Fisher information, Hessians, similarity matrices based on the gradients, or simply the gradient lengths to compute the acquisition scores that guide sample selection. Are these different approaches connected, and if so how? In this paper, we revisit the Fisher information and use it to show how several otherwise disparate methods are connected as approximations of information-theoretic quantities.

1 Introduction

Label and training efficiency are the key to a wider deployment of deep learning. Deep learning generally requires a lot of data, most of which has to be annotated. This is expensive and time-consuming. Together with semi-supervised and unsupervised approaches, active learning helps increase label efficiency. Given access to unlabeled data, *active learning* selects the most informative samples to label for a given model, thus decreasing the number of required annotations. Apart from label efficiency, training deep learning models is expensive and time-consuming, too, and *active sampling* improves training efficiency by filtering the training set to focus on the samples that will be the most informative for the model.

The acquisition functions used to select informative samples can often be traced back to information-theoretic objectives that are known from Bayesian optimal experiment design (Lindley, 1956). The connection between contemporary acquisition functions and basic information-theoretic quantities (short: *information quantities*) is the topic of this work. Amongst them we will examine: for active learning, the *expected information gain* (EIG) (Lindley, 1956; Hounsby et al., 2011; Kirsch et al., 2019)), the *(joint) expected predictive information gain* (JEPIG or EPIG, respectively) (Kirsch et al., 2021b; MacKay, 1992), and, for active sampling, the *information gain* (IG) (Sun et al., 2022) and *(joint) predictive information gain* (JPIG or PIG, respectively) (Mindermann et al., 2022). We will show that many contemporary methods can be seen as approximating information quantities using different approximations, which have different trade-offs.

We develop a unifying perspective on various existing approaches: we look at second-order posterior approximations (Laplace approximations) in §3, which we use to revisit the Fisher information and its approximations in §4. This leads to approximations of the information quantities mentioned above in §5. We pay special attention to the limitations: for example, we will see that approximations that use the trace of the Fisher information do not take redundancies between samples into account. They exhibit the same pathologies as other methods that in essence score points individually, also known as top-K batch acquisition (Kirsch et al., 2021a). In §6, we connect these approximations to approximations that use similarity matrices of the log loss gradients. In §7, we show that (Batch-)BALD and EPIG on the one hand; and BADGE, BAIT, SIMILAR¹, and PRISM¹ on the other hand can be seen as optimizing the same objectives. The difference is that (Batch-)BALD (Hounsby et al., 2011; Kirsch et al., 2019) and EPIG (Kirsch et al., 2021b) operate in prediction space, while Fisher information-based methods operate in weight space: we show that

¹using log det objectives

an approximation of EPIG, a transductive active learning objective, using Fisher information, matches the BAIT objective (Ash et al., 2021). We show how BADGE (Ash et al., 2019) can be seen as approximating the EIG, using the connection between similarity matrices and EIG approximations; similarly, we find that submodularity-based approaches (Iyer et al., 2021) such as SIMILAR (Kothawade et al., 2021) and PRISM (Kothawade et al., 2022), which report their best results using the log det of similarity matrices, approximate information quantities. We also show that gradient-length-based methods like *expected gradient length* (Settles et al., 2007) and *gradient norm score* (Paul et al., 2021) can be connected to information quantities.

Although our results employ a hierarchy of approximations, we do not examine the error terms in detail. This is in line with how these approximations are used in deep learning, where they often only provide motivation for useful mechanisms. However, we try to identify where these approximations might break, enumerate the limitations they introduce, and raise research questions that can be verified empirically as future work.

In general, we show that various disparate non-Bayesian approaches in data subset selection approximate information quantities from the Bayesian literature, going back to Lindley (1956) and MacKay (1992).

2 Background & Setting

This section introduces the relevant notation and the probabilistic model we use throughout this paper.

Information Theory. We follow the practical notation from Kirsch and Gal (2021). In particular, for entropy, we use an implicit or explicit notation, $H[X]$ or $H(p(X))$, while $H(p \parallel q)$ denotes the cross-entropy similar to the Kullback-Leibler divergence $D_{KL}(p \parallel q)$, and $H(p)$ denotes Shannon’s information content:

$$H(p) := -\log p, \quad (1)$$

$$H[x] := H(p(x)), \quad (2)$$

$$H(p(X) \parallel q(X)) := \mathbb{E}_{p(x)}[H(q(x))], \quad (3)$$

$$H[X] := H(p(X)) := H(p(X) \parallel p(X)), \quad (4)$$

where p and q are probabilities distributions, X is a random variable, and x is an outcome. Conditional and joint entropies are defined as usual (note that we take an expectation over y as well):

$$H[X \mid Y] = \mathbb{E}_{p(x,y)}[-\log p(x \mid y)]. \quad (5)$$

The mutual information $I[X; Y]$ for random variables X and Y is defined as:

$$I[X; Y] := H[X] - H[X \mid Y], \quad (6)$$

and is also referred to as expected uncertainty reduction or expected information gain (Lindley, 1956) because the entropy $H[X]$ quantifies the uncertainty about the random variable X , and $H[X \mid Y]$ about X after observing Y (in expectation).

Probabilistic Model. We assume a supervised setting: for inputs X , we have a Bayesian model with parameters Ω that makes predictions Y . What makes the model Bayesian is that the parameters follow a distribution $p(\omega)$, and we have a probabilistic model:

$$p(y, \omega \mid x) = p(y \mid x, \omega) p(\omega). \quad (7)$$

We extend this model to additional data $\mathcal{D} := \{(x_i, y_i)\}_i$ as follows: $p(\{y_i\}, \omega \mid \{x_i\}) = p(\{y_i\} \mid \{x_i\}, \omega) p(\omega)$. That is, we examine the common discriminative case where, unlike in the generative case, we do not model $p(x)$. The corresponding marginal prediction of the model is $p(y \mid x) = \mathbb{E}_{p(\omega)}[p(y \mid x, \omega)]$.

Transductive Objectives. When an objective uses additional data (unlabeled or labeled) to guide acquisitions, we refer to this as a *transductive* objective (Yu et al., 2006; Wang et al., 2020).

Active Learning. To increase label efficiency, instead of labeling data indiscriminately, active learning iteratively selects and acquires labels for the *most informative* unlabeled data from an unlabeled *pool set* according to some underlying *acquisition function*. An acquisition function scores the informativeness of an

unlabeled candidate batch $\{x_i^{\text{acq}}\}_{i=1}^N$, and the batch that maximizes the score is selected for labeling. After each such acquisition step, the model is re-trained to take the newly labeled data into account. Labels can be acquired individually or in batches (*batch active learning*, Gal et al. (2017)). The *expected information gain (EIG)*

$$I[\Omega; Y^{\text{acq}} | x^{\text{acq}}] \quad (\text{EIG/BALD})$$

and (*joint*) *expected predictive information gain (EPIG and JEPIG, respectively)*

$$I[Y^{\text{eval}}; Y^{\text{acq}} | X^{\text{eval}}, x^{\text{acq}}], \quad (\text{EPIG})$$

$$I[\{Y_i^{\text{eval}}\}; Y^{\text{acq}} | \{x_i^{\text{eval}}\}, x^{\text{acq}}] \quad (\text{JEPIG})$$

are examples of such acquisition functions.

Active Sampling. To increase training efficiency, instead of training with all samples, active sampling selects the informative samples $\mathcal{D}^{\text{acq}} = \{(x_i^{\text{acq}}, y_i^{\text{acq}})\}_{i=1}^N$ from the training set to train on. This can be done statically before training the model, in which case this is also referred to as core-set selection, or more dynamically, in which case it is also referred to as curriculum learning. The *information gain (IG)*

$$I[\Omega; y^{\text{acq}} | x^{\text{acq}}], \quad (\text{IG})$$

and (*joint*) *predictive information gain (PIG or JPIG, respectively)*

$$I[Y^{\text{eval}}; y^{\text{acq}} | X^{\text{eval}}, x^{\text{acq}}] \quad (= \mathbb{E}_{\hat{p}(x^{\text{eval}}, y^{\text{eval}})} I[y^{\text{eval}}; y^{\text{acq}} | x^{\text{eval}}, x^{\text{acq}}]), \quad (\text{PIG})$$

$$I[\{y_i^{\text{eval}}\}; y^{\text{acq}} | \{x_i^{\text{eval}}\}, x^{\text{acq}}] \quad (\text{JPIG})$$

are examples of such acquisition functions.

Log Loss. While many active learning and active sampling methods are motivated independently of the underlying loss, we will focus on log losses, such as the common cross-entropy loss or squared error loss, to be principled. These log-losses can be viewed through an information-theoretic lens.

3 Second-Order Posterior Approximation

To compute approximations of information quantities using Fisher information, we need to approximate the posterior $p(\omega | \mathcal{D}, \mathcal{D}^{\text{train}})$, where $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ are additional (new) samples, and we start with $p(\omega | \mathcal{D}^{\text{train}})$ as the “prior” distribution. We will drop $\mathcal{D}^{\text{train}}$ and use $p(\omega)$ when possible.

To begin, we can complete the square of a second-order Taylor approximation around the log parameter likelihood around a fixed ω^* :

$$\log p(\omega) \approx \log p(\omega^*) + \nabla_{\omega} \log p(\omega^*)(\omega - \omega^*) + \frac{1}{2}(\omega - \omega^*)^T \nabla_{\omega}^2 \log p(\omega^*)(\omega - \omega^*) \quad (8)$$

$$= \frac{1}{2}(\omega - (\omega^* - \nabla_{\omega}^2 \log p(\omega^*)^{-1} \nabla_{\omega} \log p(\omega^*))^T \nabla_{\omega}^2 \log p(\omega^*)(\omega - (\omega^* - \nabla_{\omega}^2 \log p(\omega^*)^{-1} \nabla_{\omega} \log p(\omega^*))) + \dots \quad (9)$$

We can express this in more concise terms by extending the notation of $H[\cdot]$ to its derivatives:

Notation 3.1. We write $H'[\cdot]$ for the Jacobian and $H''[\cdot]$ for the Hessian of $H[\cdot]$:

$$H'[\cdot] := -\nabla_{\omega} \log p(\cdot), \quad (10)$$

$$H''[\cdot] := -\nabla_{\omega}^2 \log p(\cdot). \quad (11)$$

Then, we have:

$$H[\omega] \approx H[\omega^*] + H'[\omega^*](\omega - \omega^*) + \frac{1}{2}(\omega - \omega^*)^T H''[\omega^*](\omega - \omega^*) \quad (12)$$

$$= \frac{1}{2}(\omega - (\omega^* - H''[\omega^*]^{-1} H'[\omega^*])^T H''[\omega^*] (\omega - (\omega^* - H''[\omega^*]^{-1} H'[\omega^*])) + \dots \quad (13)$$

Comparing this to the information content of a multivariate Gaussian distribution:

$$H[\mathcal{N}(w; \mu, \Sigma)] = \frac{1}{2}(\omega - \mu)^T \Sigma^{-1} (\omega - \mu) + \dots, \quad (14)$$

we obtain:

Proposition 3.2. *An approximation of the distribution $p(\omega)$ of Ω around some ω^* is given by:*

$$\Omega \approx \mathcal{N}(\omega^* - H''[\omega^*]^{-1} H'[\omega^*], H''[\omega^*]^{-1}), \quad (15)$$

where $H''[\omega^*]$ must be positive-definite. If ω^* is also a (global) minimizer of $H[\omega]$ (i.e. $H'[\omega^*] = 0$), we obtain the Laplace approximation:

$$\Omega \approx \mathcal{N}(\omega^*, H''[\omega^*]^{-1}). \quad (16)$$

Approximation Quality. Obviously, this approximation can be arbitrarily bad depending on $p(\omega)$ and ω^* —however, given enough data, we expect that $p(w)$ will concentrate around a maximum a posteriori (MAP) estimate, giving rise to the Laplace approximation. Obviously, insufficient data to reach concentration of the parameters and multimodality in over-parameterized models (Long, 2021) can be issues for active learning and active sampling.

Flat Minimum Intuition. The information-geometric interpretation of the Hessian being positive definite is that the information content (pointwise entropy) is convex around ω^* and, equivalently, that the (log) posterior is concave around ω^* . The latter provides an intuition for the Gaussian approximation: the Hessian measures curvature, and the “flatter” the Hessian, e.g., the smaller the largest eigenvalue or the smaller the determinant, the less the loss changes when ω^* is perturbed. This leads to the search for flat minima as a way to improve generalization (Hinton and Van Camp, 1993; Hochreiter and Schmidhuber, 1994).

Notation 3.3. *To further shorten the notation, we will write $H''[\mathcal{D} | \omega^*]$ instead of $H''[\{y_i\} | \{x_i\}, \omega^*]$.*

Posterior Approximation of $\Omega | \mathcal{D}$. While the Laplace approximation is centered on a (global) minimizer, the approach above can be used for a (potentially low-quality) posterior approximation in general. We can expand $H[\omega^* | \mathcal{D}]$ using Bayes’ law and the additivity of the logarithm. That is, we have:

$$H[\omega^* | \mathcal{D}] = H[\mathcal{D} | \omega^*] + H[\omega^*] - H[\mathcal{D}], \quad (17)$$

and then, as $H[\mathcal{D}]$ is independent of ω :

$$H'[\omega^* | \mathcal{D}] = H'[\mathcal{D} | \omega^*] + H'[\omega^*] + 0 = H'[\mathcal{D} | \omega^*] + H'[\omega^*], \quad (18)$$

$$H''[\omega^* | \mathcal{D}] = H''[\mathcal{D} | \omega^*] + H''[\omega^*]. \quad (19)$$

Proposition 3.4. *The observed information $H''[\{y_i\} | \{x_i\}, \omega^*]$ is additive:*

$$H''[\{y_i\} | \{x_i\}, \omega^*] = \sum_i H''[y_i | x_i, \omega^*] = \sum_i -\nabla_{\omega}^2 \log p(y_i | x_i, \omega^*). \quad (20)$$

The observed information is defined with opposite sign compared to other works as it simplifies the exposition.

Uninformative Prior. For a typical Gaussian prior $p(\omega) \sim \mathcal{N}(\mu, \Sigma)$, we have $H''[\omega^*] = \Sigma^{-1}$ and $H''[\omega^* | \mathcal{D}] = H''[\mathcal{D} | \omega^*] + \Sigma^{-1}$. For an uninformative prior with “infinite prior variance” $\Sigma^{-1} \rightarrow 0$, we have $H''[\omega^*] = 0$ and $H''[\omega^* | \mathcal{D}] = H''[\mathcal{D} | \omega^*]$.

Proposition 3.5. *The entropy of the second-order approximation of $p(\omega)$ around ω^* is*

$$H[\Omega] \approx -\frac{1}{2} \log \det H''[\omega^*] + C_k, \quad (21)$$

where $C_k = \frac{k}{2} \log 2\pi e$ is a constant (independent of \mathcal{D} and ω^*) and k is the number of dimensions of ω .

4 Fisher Information

The following section revisits the Fisher information and its properties. All proofs can be found in §A. In particular, we look at two special cases with more favorable properties: following [Kunstner et al. \(2019\)](#), when we can write our model as $p(y | \hat{z} = \hat{f}(x; \omega))$, where $\hat{f}(x; \omega)$ are the logits, and $p(y | \hat{z})$ is a distribution from an exponential family, the Fisher information is independent of y , which has useful consequences as we will see; and following [Chaudhuri et al. \(2015\)](#), when we have a Generalized Linear Model (GLM), the Hessian is independent of y . The results for the GLM are also used as an approximation known as *Generalized Gauss-Newton approximation*. Together with other numerical approximations, such as using a diagonal matrix or low-rank factorizations, the Hessian and Fisher information can be efficiently approximated for large deep neural networks ([Daxberger et al., 2021](#)).

Definition 4.1. *The Fisher information $F(\{x_i\}, \omega^*)$ is the expectation over the observed information using the model’s own predictions for given $\{x_i\}$ at ω^* :*

$$F(\{x_i\}, \omega^*) = \mathbb{E}_{p(\{y_i\} | \{x_i\}, \omega^*)} [H''[\{y_i\} | \{x_i\}, \omega^*]]. \quad (22)$$

Note that this definition of the Fisher information is also referred to as “empirical” Fisher information in statistics because we do not take an expectation over x but use empirical samples for x ([Kunstner et al., 2019](#)).

Proposition 4.2. *Like the observed information, the Fisher information is additive:*

$$F(\{x_i\}, \omega^*) = \sum_i F(x_i, \omega^*). \quad (23)$$

There are two other equivalent definitions of the Fisher information:

Proposition 4.3. *The Fisher information is equivalent to:*

$$F(x, \omega^*) = \mathbb{E}_{p(y | x, \omega^*)} [H'[y | x, \omega^*]^T H'[y | x, \omega^*]] = \text{Cov}[H'[Y | x, \omega^*]]. \quad (24)$$

Note that the expectation is over $p(y | x, \omega^*)$ and not $p(y | x)$. The proof in §A applies two generally useful lemmas:

Lemma 4.4. *For the Jacobian $H'[y | x, \omega^*]$, we have:*

$$H'[y | x, \omega^*] = -\frac{\nabla_{\omega} p(y | x, \omega^*)}{p(y | x, \omega^*)}, \quad (25)$$

and for the Hessian $H''[y | x, \omega^*]$, we have:

$$H''[y | x, \omega^*] = H'[y | x, \omega^*]^T H'[y | x, \omega^*] - \frac{\nabla_{\omega}^2 p(y | x, \omega^*)}{p(y | x, \omega^*)}. \quad (26)$$

Lemma 4.5. *The following expectations over the model’s own predictions vanish:*

$$\mathbb{E}_{p(y | x, \omega^*)} [H'[y | x, \omega^*]] = 0, \quad (27)$$

$$\mathbb{E}_{p(y | x, \omega^*)} \left[\frac{\nabla_{\omega}^2 p(y | x, \omega^*)}{p(y | x, \omega^*)} \right] = 0. \quad (28)$$

Special Case: Exponential Family. [Kunstner et al. \(2019\)](#) show in their appendix that if we split a discriminative model into prelogits $\hat{f}(x; \omega)$ and a predictor $p(y | \hat{z} = \hat{f}(x; \omega))$, the Hessian does not depend on y when $p(y | \hat{z})$ is a distribution from an exponential family (independent of ω). Examples include using a normal distribution for regression parameterized by mean and variance predictions or a categorical distribution via the softmax function. The following statements and proofs follow [Kunstner et al. \(2019\)](#):

Proposition 4.6. *The Fisher information $F(\{x_i\}, \omega^*)$ for a model $p(y | \hat{z} = \hat{f}(x; \omega^*))$ is equivalent to:*

$$F(x, \omega^*) = \nabla_{\omega} \hat{f}(x; \omega^*)^T \mathbb{E}_{p(y|x, \omega^*)} [\nabla_{\hat{z}}^2 H[y | \hat{z} = \hat{f}(x; \omega^*)]] \nabla_{\omega} \hat{f}(x; \omega^*), \quad (29)$$

where $\nabla_{\hat{z}}^2 H[y | \hat{z} = \hat{f}(x; \omega^*)]$ is short for $\nabla_{\hat{z}}^2 H[y | \hat{z}]|_{\hat{z}=\hat{f}(x; \omega^*)}$.

Proposition 4.7. *The Fisher information $F(x, \omega^*)$ of a model of the form $p(y | \hat{z} = \hat{f}(x; \omega^*))$ is independent of y , where $p(y | \hat{z})$ is an exponential distribution, i.e., $\log p(y | \hat{z}) = \hat{z}^T T(y) - A(\hat{z}) + \log h(y)$:*

$$F(x, \omega^*) = \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 A(\hat{z} = \hat{f}(x; \omega^*)) \nabla_{\omega} \hat{f}(x; \omega^*). \quad (30)$$

It is crucial that the exponential distribution not depend on ω . This simplifies computing the Fisher information: no expectation over y s is needed anymore. The full outer product may not be needed explicitly either.

To make this more concrete, we will consider two common exponential distributions:

Gaussian Distribution. When $p(y | \hat{z}) = \mathcal{N}(y; \hat{z}, 1)$, we have $H''[y | \hat{z}] = 1$, and thus

$$F(x, \omega^*) = \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\omega} \hat{f}(x; \omega^*). \quad (31)$$

Categorical Distribution. When $p(y | \hat{z}) = \text{softmax}(\hat{z})_y$, we have $H''[y | \hat{z}] = \text{diag}(\pi) - \pi \pi^T$, with $\pi_y = p(y | \hat{z})$, and thus:

$$F(x, \omega^*) = \nabla_{\omega} \hat{f}(x; \omega^*)^T (\text{diag}(\pi) - \pi \pi^T) \nabla_{\omega} \hat{f}(x; \omega^*). \quad (32)$$

Special Case: Generalized Linear Models. Chaudhuri et al. (2015) require that the Hessian is independent of y , which we will make use of later as well. This holds for Generalized Linear Models:

Definition 4.8. *A generalized linear model (GLM) is a model $p(y | \hat{z} = \hat{f}(x; \omega))$ such that $\log p(y | \hat{z}) = \hat{z}^T T(y) - A(\hat{z}) + \log h(y)$ is a distribution of the exponential family, independent of ω , and $\hat{f}(x; \omega) = \omega^T x$ is linear in the parameters ω .*

Proposition 4.9. *The Hessian $H''[y | x, \omega^*]$ of a GLM is independent of y .*

$$H''[y | x, \omega^*] = \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 H[y | \hat{z} = \hat{f}(x; \omega^*)] \nabla_{\omega} \hat{f}(x; \omega^*) \quad (33)$$

$$= \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 A(w^T x) \nabla_{\omega} \hat{f}(x; \omega^*). \quad (34)$$

Proposition 4.10. *For a model such that the Hessian $H''[y | x, \omega^*]$ is independent of y , we have:*

$$F(x, \omega^*) = H''[y^* | x, \omega^*] \quad (35)$$

for any y^* , and:

$$\mathbb{E}_{p(y|x)} [H''[y | x, \omega^*]] = F(x, \omega^*). \quad (36)$$

Note that the expectation is over $p(y|x)$ and not $p(y|x, \omega^*)$, and $\mathbb{E}_{p(\{y_i\}|\{x_i\})} [H''[\{y_i\}|\{x_i\}, \omega^*]] = F(\{x_i\}, \omega^*)$ is additive.

Proposition 4.11. *When $\hat{f}(x; \omega) : \mathbb{R}^D \rightarrow \mathbb{R}^C$, where C is the number of classes (outputs) and D is the number of input dimensions, and $\omega \in \mathbb{R}^{D \times C}$, and assuming the parameters are flattened into a single vector for the Jacobian, $\nabla_{\omega} \hat{f}(x; \omega^*) = \text{Id}_C \otimes x^T \in \mathbb{R}^{C \times (C \cdot D)}$, where \otimes denotes the Kronecker product, we have:*

$$\nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 A(w^T x) \nabla_{\omega} \hat{f}(x; \omega^*) = \nabla_{\hat{z}}^2 A(w^T x) \otimes x x^T. \quad (37)$$

$p(y | x, \omega^*)$ vs $p(y | x)$. Having a GLM solves an important issue we will encounter in §5: approximating the EIG requires taking an expectation over $p(y | x)$ and not $p(y | x, \omega^*)$. One can approximate $p(y | x) \approx p(y | x, \omega^*)$, which can be justified in the limit, but this is likely not a good approximation in the cases interesting for active learning and active sampling. With a GLM, this is not a problem.

Generalized Gauss-Newton Approximation. When we have an exponential family but not a GLM, the equality in Proposition 4.9 is often used as an approximation for the Hessian, i.e., we simply use the Fisher information as an approximation of the Hessian (via Proposition 4.7):

$$H''[y | x, \omega^*] \approx \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_z^2 A(w^T x) \nabla_{\omega} \hat{f}(x; \omega^*) = F(x, \omega^*). \quad (38)$$

This is known as *Generalized Gauss-Newton (GGN) approximation* (Kunstner et al., 2019; Immer et al., 2020). This approximation has the advantage that it is always positive semidefinite unlike the true Hessian.

Last-Layer Approaches. GLMs are often used in deep active learning (Gal et al., 2017; Ash et al., 2019; 2021). If we split the model into $p(y | x, \omega) = p(y | z = \omega^T f(x))$, where $z = f(x)$ are the embeddings, and treat the encoder $f(x)$ as fixed, we obtain a GLM that only uses the embeddings and last-layer weights.

5 Approximating Information Quantities

We use the results so far to derive various approximations for information quantities using the Hessian and Fisher information. This will help us connect information quantities to existing literature in §7.

5.1 Approximate Expected Information Gain

The expected information gain is a popular acquisition function in active learning, where it is also known as BALD, and in Bayesian optimal experimental design (Gal et al., 2017; Houlby et al., 2011; Lindley, 1956).

We can approximate the EIG $I[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}]$ of acquisition candidates $\{x_i^{\text{acq}}\}$ using the Fisher information:

$$I[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}] = H[\Omega] - H[\Omega | \{Y_i^{\text{acq}}\}, \{x_i^{\text{acq}}\}] \quad (39)$$

$$= H[\Omega] - \mathbb{E}_{p(\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\})} [H[\Omega | \{y_i^{\text{acq}}\}, \{x_i^{\text{acq}}\}]] \quad (40)$$

$$\approx -\frac{1}{2} \log \det H''[\omega^*] - \mathbb{E}_{p(\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\})} [-\frac{1}{2} \log \det H''[\omega | \{y_i^{\text{acq}}\}, \{x_i^{\text{acq}}\}]] \quad (41)$$

$$= \frac{1}{2} \mathbb{E}_{p(\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\})} [\log \det ((H''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*] + H''[\omega^*]) H''[\omega^*]^{-1})] \quad (42)$$

$$= \frac{1}{2} \mathbb{E}_{p(\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\})} [\log \det (H''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*] H''[\omega^*]^{-1} + Id)]. \quad (43)$$

The constant C_k from Proposition 3.5 cancels out as we subtract two entropy terms.

Generalized Linear Model. When we have a GLM, we can use Proposition 4.10 to obtain:

$$I[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}] \approx \dots = \frac{1}{2} \mathbb{E}_{p(\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\})} [\log \det (H''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*] H''[\omega^*]^{-1} + Id)] \quad (44)$$

$$= \frac{1}{2} \log \det (F(\{x_i^{\text{acq}}\}, \omega^*) H''[\omega^*]^{-1} + Id). \quad (45)$$

We can upper-bound the log determinant and obtain:

$$I[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}] \approx \frac{1}{2} \log \det (F(\{x_i^{\text{acq}}\}, \omega^*) H''[\omega^*]^{-1} + Id) \quad (46)$$

$$\leq \frac{1}{2} \text{tr} (F(\{x_i^{\text{acq}}\}, \omega^*) H''[\omega^*]^{-1}) \quad (47)$$

$$= \frac{1}{2} \sum_i \text{tr} (F(x_i^{\text{acq}}, \omega^*) H''[\omega^*]^{-1}). \quad (48)$$

where we have used the following inequality (proof in §B.1):

Lemma 5.1. *For symmetric, positive-semidefinite matrices A , we have (with equality iff $A = 0$):*

$$\log \det(A + Id) \leq \text{tr}(A). \quad (49)$$

General Case & Exponential Family. For the general case, we need to make a strong approximation:

$$p(\{y_i^{acq}\} | \{x_i^{acq}\}) \approx p(\{y_i^{acq}\} | \{x_i^{acq}\}, \omega^*), \quad (50)$$

which likely holds for a mostly converged posterior, but probably not for cases with little data. Alternatively, we could use the GGN approximation when we have an exponential family for the same result (but not an upper bound). See §B.1 for the derivation.

Proposition 5.2 (EIG). *The expected information gain can be approximately upper bounded via:*

$$I[\Omega; \{Y_i^{acq}\} | \{x_i^{acq}\}, \mathcal{D}^{train}] \lesssim \frac{1}{2} \log \det \left(\sum_i F(x_i^{acq}, \omega^*) H''[\omega^* | \mathcal{D}^{train}]^{-1} + Id \right) \quad (51)$$

$$\leq \frac{1}{2} \sum_i \text{tr} (F(x_i^{acq}, \omega^*) H''[\omega^* | \mathcal{D}^{train}]^{-1}). \quad (52)$$

Furthermore, we have the following:

$$\arg \max_{\{x_i^{acq}\}} \{I[\Omega; \{Y_i^{acq}\} | \{x_i^{acq}\}, \mathcal{D}^{train}]\} = \arg \max_{\{x_i^{acq}\}} \{-H[\Omega | \{Y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}]\}, \text{ and} \quad (53)$$

$$-H[\Omega | \{Y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] \lesssim \frac{1}{2} \log \det \left(\sum_i F(x_i^{acq}, \omega^*) + H''[\omega^* | \mathcal{D}^{train}] \right) - C_k. \quad (54)$$

The second statement follows from Equation (39), since $H[\Omega | \mathcal{D}^{train}]$ is constant, and provides an alternative objective when we are only interested in optimization of the EIG. We use it together to show a connection to the expected gradient length approach in active learning in §7.

Batch Acquisition Pathologies. Importantly, this approximation of the EIG using the trace is additive, whereas the one using the log determinant is not. This means that the trace approximation ignores dependencies between samples and only leads to naive top-k batch acquisition; see Kirsch et al. (2019; 2021a) for details. On the other hand, the log determinant of the Fisher information version might well capture these dependencies.

5.2 Approximate Information Gain

Following the same steps, we can also approximate the information gain, which is useful for active sampling:

Proposition 5.3 (IG). *The information gain $I[\Omega; \{y_i^{acq}\} | \{x_i^{acq}\}, \mathcal{D}^{train}] = H[\Omega | \mathcal{D}^{train}] - H[\Omega | \{y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}]$ can be approximately upper bounded via:*

$$I[\Omega; \{y_i^{acq}\} | \{x_i^{acq}\}, \mathcal{D}^{train}] \lesssim \frac{1}{2} \log \det (H''[\{y_i^{acq}\} | \{x_i^{acq}\}, \omega^*] H''[\omega^* | \mathcal{D}^{train}]^{-1} + Id) \quad (55)$$

$$\leq \frac{1}{2} \sum_i \text{tr} (H''[\{y_i^{acq}\} | \{x_i^{acq}\}, \omega^*] H''[\omega^* | \mathcal{D}^{train}]^{-1}). \quad (56)$$

Furthermore, we have the following:

$$\arg \max_{\{x_i^{acq}\}} \{I[\Omega; \{y_i^{acq}\} | \{x_i^{acq}\}, \mathcal{D}^{train}]\} = \arg \max_{\{x_i^{acq}\}} \{-H[\Omega | \{y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}]\}, \text{ and} \quad (57)$$

$$-H[\Omega | \{y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] \lesssim \frac{1}{2} \log \det (H''[\{y_i^{acq}\} | \{x_i^{acq}\}, \omega^*] + H''[\omega^* | \mathcal{D}^{train}]) - C_k. \quad (58)$$

Comparison to EIG. Importantly, when we have a GLM or use the GGN approximation, this approximation of the IG is equal to the one of the EIG. This tells us that active learning on a GLM with the EIG approximation will work as well as if we had access to the labels. On the other hand, active sampling via IG with the GGN approximation will not work better than the equivalent active learning approach without labels.

5.3 Approximate Expected Predictive Information Gain

In transductive active learning, we have access to an (empirical) distribution $\hat{p}(x^{\text{eval}})$ (e.g., the pool set) and want to find x^{acq} that maximizes the *expected predictive information gain (EPIG)* (Kirsch et al., 2021b):

$$\arg \max_{x^{\text{acq}}} \mathbb{I}[Y^{\text{eval}}; Y^{\text{acq}} | X^{\text{eval}}, x^{\text{acq}}, \mathcal{D}^{\text{train}}] = \arg \max_{x^{\text{acq}}} \mathbb{E}_{\hat{p}(x^{\text{eval}})} \mathbb{I}[Y^{\text{eval}}; Y^{\text{acq}} | x^{\text{eval}}, x^{\text{acq}}, \mathcal{D}^{\text{train}}], \quad (59)$$

We expand the objective as follows:

$$\mathbb{I}[Y^{\text{eval}}; Y^{\text{acq}} | X^{\text{eval}}, x^{\text{acq}}] = \mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}] - \mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}, Y^{\text{acq}}, x^{\text{acq}}], \quad (60)$$

where $\mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}]$ can be removed from the objective because it is independent of x^{acq} , thus, optimizing EPIG is equivalent to *minimizing* $\mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}, Y^{\text{acq}}, x^{\text{acq}}]$:

$$\arg \max_{x^{\text{acq}}} \mathbb{I}[Y^{\text{eval}}; Y^{\text{acq}} | X^{\text{eval}}, x^{\text{acq}}] = \arg \min_{x^{\text{acq}}} \mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}, Y^{\text{acq}}, x^{\text{acq}}]. \quad (61)$$

Following Proposition 5.2, this can be approximated by:

$$\begin{aligned} & \mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}, Y^{\text{acq}}, x^{\text{acq}}] \\ & \approx \frac{1}{2} \mathbb{E}_{p(y^{\text{eval}}, y^{\text{acq}} | x^{\text{eval}}, x^{\text{acq}})} \hat{p}(x^{\text{eval}}) [\log \det (H''[y^{\text{eval}} | x^{\text{eval}}, \omega^*] (H''[y^{\text{acq}} | x^{\text{acq}}, \omega^*] + H''[\omega^*]^{-1} + Id)]. \end{aligned} \quad (62)$$

Generalized Linear Model. For a generalized linear model, we can drop the expectation and obtain:

$$\begin{aligned} & \mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}, Y^{\text{acq}}, x^{\text{acq}}] \\ & \approx \frac{1}{2} \mathbb{E}_{\hat{p}(x^{\text{eval}})} [\log \det (F(x^{\text{eval}}, \omega^*) (F(x^{\text{acq}}, \omega^*) + H''[\omega^*]^{-1} + Id)] \end{aligned} \quad (63)$$

$$\leq \frac{1}{2} \log \det (\mathbb{E}_{\hat{p}(x^{\text{eval}})} [F(x^{\text{eval}}, \omega^*)] (F(x^{\text{acq}}, \omega^*) + H''[\omega^*]^{-1} + Id) \quad (64)$$

$$\leq \frac{1}{2} \text{tr} (\mathbb{E}_{\hat{p}(x^{\text{eval}})} [F(x^{\text{eval}}, \omega^*)] (F(x^{\text{acq}}, \omega^*) + H''[\omega^*]^{-1}), \quad (65)$$

where we have again used the concavity of the log determinant.

General Case & Exponential Family. To our knowledge, there is no rigorous way to obtain a similar result in the general case as the Fisher information for an acquisition candidate now lies within an inverted term. Obviously, the GGN approximation can be applied when we have an exponential family, which leads back to the above GLM result as an approximation. See §B.2 for more details.

Proposition 5.4 (EPIG). *For a generalized linear model (or with the GGN approximation), we have:*

$$\arg \max_{\{x_i^{\text{acq}}\}} \mathbb{I}[Y^{\text{eval}}, \{Y_i^{\text{acq}}\} | X^{\text{eval}}, \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}] = \arg \min_{\{x_i^{\text{acq}}\}} \mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}, \{Y_i^{\text{acq}}\}, \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}], \quad (66)$$

with

$$\begin{aligned} & \mathbb{I}[\Omega; Y^{\text{eval}} | X^{\text{eval}}, \{Y_i^{\text{acq}}\}, \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}] \\ & \approx \frac{1}{2} \log \det (\mathbb{E}_{p(x^{\text{eval}})} [F(x^{\text{eval}}, \omega^*)] (F(\{x_i^{\text{acq}}\}, \omega^*) + H''[\omega^* | \mathcal{D}^{\text{train}}]^{-1} + Id) \end{aligned} \quad (67)$$

$$\leq \frac{1}{2} \text{tr} (\mathbb{E}_{p(x^{\text{eval}})} [F(x^{\text{eval}}, \omega^*)] (F(\{x_i^{\text{acq}}\}, \omega^*) + H''[\omega^* | \mathcal{D}^{\text{train}}]^{-1}). \quad (68)$$

Approximations for JEPIG, PIG and JPIG. The results can be found in §B.

6 Similarity Matrices and One-Sample Approximations of the Fisher Information

Many active learning methods do not use Fisher information, but use a kernel-based approach using similarity matrices of the loss gradients $H'[y^* | x]$, where y^* is either the true label or a hypothesized label y^* when no label is available (usually using the arg max prediction of the model).

Connection to Fisher Information. Crucially, given $\mathcal{D} = \{(y_i, x_i)_i\}$, if we let

$$\hat{H}'[\mathcal{D} | \omega^*] := \begin{pmatrix} \vdots \\ H'[y_i | x_i, \omega^*] \\ \vdots \end{pmatrix} \quad (69)$$

be a “data matrix” of the Jacobians, then $\hat{H}'[\mathcal{D} | \omega^*] \hat{H}'[\mathcal{D} | \omega^*]^T$ yields the similarity matrix $S[\mathcal{D} | \omega^*]$ using the Euclidean inner product:

$$S[\mathcal{D} | \omega^*]_{ij} := \langle H'[y_i | x_i, \omega^*], H'[y_j | x_j, \omega^*] \rangle = \hat{H}'[\mathcal{D} | \omega^*] \hat{H}'[\mathcal{D} | \omega^*]^T. \quad (70)$$

If we sample the $\{y_i\} \sim p(\{y_i\} | \{x_i\}, \omega^*)$, the “flipped” product $\hat{H}'[\mathcal{D} | \omega^*]^T \hat{H}'[\mathcal{D} | \omega^*]$ yields a one-sample estimate of the Fisher information $F(\{x_i\}, \omega^*)$:

$$F(\{x_i\}, \omega^*) = \sum_i F(x_i, \omega^*) = \mathbb{E}_{p(\{y_i\} | \{x_i\}, \omega^*)} \sum_i H'[y_i | x_i, \omega^*]^T H'[y_i | x_i, \omega^*] \quad (71)$$

$$= \mathbb{E}_{p(\{y_i\} | \{x_i\}, \omega^*)} \hat{H}'[\mathcal{D} | \omega^*]^T \hat{H}'[\mathcal{D} | \omega^*]. \quad (72)$$

If we do not sample but instead use the arg max class, we only obtain a biased estimate (Kunstner et al., 2019, §B).

Connection to the Expected Information Gain. More importantly, when we define an inner product $\langle \cdot, \cdot \rangle_{H''[\omega^* | \mathcal{D}^{\text{train}}]}$ using the Hessian $H''[\omega^* | \mathcal{D}^{\text{train}}]$, we can connect the similarity matrix, which uses this inner product, to our information gain approximations:

$$S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D} | \omega^*] := \hat{H}'[\mathcal{D} | \omega^*] H''[\omega^* | \mathcal{D}^{\text{train}}]^{-1} \hat{H}'[\mathcal{D} | \omega^*]^T \quad (73)$$

Specifically, we apply the matrix determinant lemma $\det(AB + M) = \det M \det(Id + BM^{-1}B)$ to obtain:

Proposition 6.1. *Given $\mathcal{D}^{\text{train}}$, $\{x_i^{\text{acq}}\}$ and (sampled) $\{y_i^{\text{acq}}\}$, we have for the EIG:*

$$I[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}] \lesssim \frac{1}{2} \log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{acq}} | \omega^*] + Id) \quad (74)$$

Proposition 6.2. *Assuming an uninformative posterior $H''[\omega^* | \mathcal{D}^{\text{train}}] = \lambda Id$ for $\lambda \rightarrow 0$, and given $\mathcal{D}^{\text{train}}$, $\{x_i^{\text{acq}}\}$, and (sampled) $\{y_i^{\text{acq}}\}$, we have for the EIG (before taking $\lambda \rightarrow 0$):*

$$I[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}] \lesssim \frac{1}{2} \log \det (S[\mathcal{D}^{\text{acq}} | \omega^*] + \lambda Id) - \frac{|\mathcal{D}^{\text{acq}}|}{2} \log \lambda. \quad (75)$$

The second term is constant independently of $\{x_i^{\text{acq}}\}$, and we can maximize $\{x_i^{\text{acq}}\}$ in $\log \det (S[\mathcal{D}^{\text{acq}} | \omega^] + \lambda Id)$ and then take the limit $\lambda \rightarrow 0$. Thus, we can use the following proxy objective:*

$$\log \det (S[\mathcal{D}^{\text{acq}} | \omega^*]). \quad (76)$$

Connection to Other Approximate Information Quantities. Interestingly, we can use the above approach to obtain valid approximations for the predictive information gains (EPIG and JEPIG) because the terms that would go towards $-\infty$ cancel out; see §C for details.

7 Connection to Other Acquisition Functions in the Literature

Here, we use the results so far to connect approaches in the literature to information quantities explicitly.

7.1 BAIT in “Gone Fishing” (Ash et al., 2021) and ActiveSetSelect in “Convergence Rates of Active Learning for Maximum Likelihood Estimation” (Chaudhuri et al., 2015)

While Chaudhuri et al. (2015) and Ash et al. (2021) use a similar objective, Chaudhuri et al. (2015) apply it only to GLMs, whereas Ash et al. (2021) use it for deep learning and introduce it as the *BAIT* objective:

$$\arg \min_{\{x_i^{\text{acq}}\}} \text{tr} \left(\left(F(\{x_i^{\text{acq}}\}, \omega^*) + F(x^{\text{train}}, \omega^*) + \lambda I \right)^{-1} F(\{x_i^{\text{eval}}\}, \omega^*) \right), \quad (\text{BAIT})$$

where λ is a hyperparameter. While Ash et al. (2021) use DNNs, they only use the last layer to compute the Fisher information. The last layer of the DNNs together with appropriate activation functions and losses constitute a generalized linear model. Following Proposition 5.4, Ash et al. (2021) thus perform transductive active learning (using the pool set as an evaluation set) and approximate EPIG as information quantity:

Proposition 7.1. *Both Chaudhuri et al. (2015) and Ash et al. (2021) perform transductive active learning, approximating EPIG:*

$$\arg \max_{x^{\text{acq}}} \mathbb{I}[Y^{\text{eval}}; Y^{\text{acq}} | X^{\text{eval}}, x^{\text{acq}}] \approx \arg \min_{x^{\text{acq}}} \text{tr} (F(x^{\text{eval}}, \omega^*) (\sum_i F(x^{\text{acq}}_i, \omega^*) + H''[\omega^* | \mathcal{D}^{\text{train}}])^{-1}), \quad (77)$$

with $H''[\omega^* | \mathcal{D}^{\text{train}}] = H''[\mathcal{D}^{\text{train}} | \omega^*] + H''[\omega^*]$ and $H''[\omega^*] = \lambda Id$.

Proof. This follows immediately for GLM (last-layer approaches) when we expand $H''[\omega^* | \mathcal{D}^{\text{train}}]$. \square

However, following §4, it does not seem that this approach will translate beyond a last-layer approach for DNNs. It thus remains an open question to find a principled approach for the general case and to go beyond last-layer active learning when using Fisher information without using the GGN approximation.

7.2 BADGE (Ash et al., 2019)

BADGE maximizes an approximation of the EIG.

Following §6, BADGE uses one-sample estimates with hard pseudo-labels $y^{\text{acq}} = \arg \max_y p(y | x^{\text{acq}}, \omega^*)$, and its sample selection is motivated by a k-DPP (Kulesza and Taskar, 2011) on the similarity matrix $S[\mathcal{D}^{\text{acq}} | \omega^*]$ using the pseudo-labels. This can be seen as approximating $\mathbb{I}[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}]$ with an uninformative posterior distribution following Proposition 6.1. However, BADGE actually uses k-MEANS++ (Arthur and Vassilvitskii, 2006; Ostrovsky et al., 2013) instead of a k-DPP to select samples to further speed up acquisition: it uses the Jacobians from the data matrix directly and samples a diverse batch based on the Euclidean distance between the Jacobians. We leave a comparison between k-DDP and k-means++ to future work and refer to the ablation in Ash et al. (2019) to see that the performance matches.

7.3 SIMILAR (Kothawade et al., 2021) and PRISM (Kothawade et al., 2022)

The LogDet objective maximizes an approximation of the EIG, and the LogDetMI objective of EPIG.

The results based on the log determinant of similarity matrices are reported as among the best in the experiments of Kothawade et al. (2021) and Kothawade et al. (2022). The LogDet objective $\log \det S[\mathcal{D}^{\text{acq}} | \omega^*]$ exactly matches the EIG approximation in §6.2. Furthermore, in §D.1, we show that the LogDetMI objective matches an approximation of JEPIG and similarly, re-derive the LogDetCMI objective.

7.4 Expected Gradient Length

Expected Gradient Length (EGL) (Settles et al., 2007; Settles, 2009) is usually defined for non-Bayesian models. Originally, it was an expectation over the gradient norm. In more recent literature (Huang et al.,

2016), it is introduced using the squared gradient norm:

$$\mathbb{E}_{p(y^{\text{acq}} | x^{\text{acq}}, \omega^*)} \left\| H'[y^{\text{acq}} | x^{\text{acq}}, \omega^*] \right\|^2. \quad (\text{EGL})$$

Using a diagonal approximation of the Fisher information, we obtain:

Proposition 7.2. *The EIG for a candidate sample x^{acq} approximately lower-bounds the EGL:*

$$2 \mathbb{I}[\Omega; Y^{\text{acq}} | x^{\text{acq}}] \lesssim \mathbb{E}_{p(y^{\text{acq}} | x^{\text{acq}}, \omega^*)} \left\| H'[y^{\text{acq}} | x^{\text{acq}}, \omega^*] \right\|^2 + \text{const.} \quad (78)$$

7.5 Deep Learning on a Data Diet

Paul et al. (2021) use the gradient length of given labeled samples x, y (averaged over multiple training runs) to select a subset of informative samples for training:

$$\mathbb{E}_{p(\omega | \mathcal{D}^{\text{train}})} \left\| H'[y^{\text{acq}} | x^{\text{acq}}, \omega] \right\|^2, \quad (\text{GraNd})$$

which they call the *gradient norm score* (GraNd).

Proposition 7.3. *The IG for a candidate sample x^{acq} approximately lower-bounds the gradient norm (GraNd) score at ω^* up to a second-order term:*

$$2 \mathbb{I}[\Omega; y^{\text{acq}} | x^{\text{acq}}] \lesssim \mathbb{E}_{p(\omega | \mathcal{D}^{\text{train}})} \left[\left\| H'[y^{\text{acq}} | x^{\text{acq}}, \omega] \right\|^2 \right] - \mathbb{E}_{p(\omega | \mathcal{D}^{\text{train}})} \left[\text{tr} \left(\frac{\nabla_{\omega}^2 p(y | x, \omega)}{p(y | x, \omega)} \right) \right] + \text{const.} \quad (79)$$

The second term might not be negligible. Hence, GraNd (the first term on the left) might deviate from the information gain. This raises the question how the information gain compares to GraNd in practice.

8 Conclusion

We have looked at Fisher information and Laplace approximations and have derived weight-space approximations for the expected information gain and expected predictive information gain. This has allowed us to connect various information quantities to objectives already used in the literature.

Last-Layer Active Learning. Methods that only use last-layer Fisher information or similar perform active learning on the embeddings only, despite that feature learning is arguably the most important feature of deep neural networks. However, these approaches can find great use with large pre-trained models, which are only fine-tuned on new data domains anyway (Tran et al., 2022).

Batch Acquisition Pathologies. Approaches that use the matrix trace instead of the log determinant are additive in the batch candidates $\{x_i^{\text{acq}}\}$ and can thus by definition not take redundancies between batch candidates into account, leading to failures detailed in Kirsch et al. (2019; 2021a).

Weight vs. Prediction Space. Ash et al. (2019; 2021); Kothawade et al. (2021; 2022) approximate the relevant information quantity in weight space, while Kirsch et al. (2021b); Houlsby et al. (2011); Kirsch et al. (2019); Mindermann et al. (2022) approximate the relevant information quantities in prediction space.

Informativeness. Taking a step back, we have seen that a Bayesian perspective using information quantities connects seemingly disparate literature. Although Bayesian methods are often seen as separate from regular active learning and active sampling, the sometimes fuzzy notion of “informativeness” expressed through various different objectives in non-Bayesian settings collapses to the same couple of information quantities, which were, in principle, already well known by Lindley (1956) and MacKay (1992).

References

- David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. Technical report, 2006.
- Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds, 2019.
- Jordan T. Ash, Surbhi Goel, Akshay Krishnamurthy, and Sham Kakade. Gone fishing: Neural active learning with fisher embeddings, 2021.
- Kamalika Chaudhuri, Sham Kakade, Praneeth Netrapalli, and Sujay Sanghavi. Convergence rates of active learning for maximum likelihood estimation, 2015.
- Thomas M Cover and A Thomas. Determinant inequalities via information theory. *SIAM journal on Matrix Analysis and Applications*, 9(3):384–392, 1988.
- Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace redux-effortless bayesian deep learning. *Advances in Neural Information Processing Systems*, 34:20089–20103, 2021.
- Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192. PMLR, 2017.
- Geoffrey E Hinton and Drew Van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13, 1993.
- Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. *Advances in neural information processing systems*, 7, 1994.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- Jiaji Huang, Rewon Child, Vinay Rao, Hairong Liu, Sanjeev Satheesh, and Adam Coates. Active learning for speech recognition: the power of gradients, 2016.
- Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of bayesian neural nets via local linearization, 2020.
- Rishabh Iyer, Ninad Khargoankar, Jeff Bilmes, and Himanshu Asanani. Submodular combinatorial information measures with applications in machine learning. In *Algorithmic Learning Theory*, pages 722–754. PMLR, 2021.
- Andreas Kirsch and Yarin Gal. A practical & unified notation for information-theoretic quantities in ml. *arXiv preprint arXiv:2106.12062*, 2021.
- Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. In *Advances in Neural Information Processing Systems*, pages 7024–7035, 2019.
- Andreas Kirsch, Sebastian Farquhar, Parmida Atighehchian, Andrew Jesson, Frederic Branchaud-Charron, and Yarin Gal. Stochastic batch acquisition for deep active learning, 2021a.
- Andreas Kirsch, Tom Rainforth, and Yarin Gal. Test distribution-aware active learning: A principled approach against distribution shift and outliers, 2021b.
- Suraj Kothawade, Nathan Beck, Krishnateja Killamsetty, and Rishabh Iyer. Similar: Submodular information measures based active learning in realistic scenarios. *Advances in Neural Information Processing Systems*, 34:18685–18697, 2021.

- Suraj Kothawade, Vishal Kaushal, Ganesh Ramakrishnan, Jeff Bilmes, and Rishabh Iyer. Prism: A rich class of parameterized submodular information measures for guided data subset selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10238–10246, 2022.
- Alex Kulesza and Ben Taskar. k-dpps: Fixed-size determinantal point processes. In *ICML*, 2011.
- Frederik Kunstner, Lukas Balles, and Philipp Hennig. Limitations of the empirical fisher approximation for natural gradient descent, 2019.
- Dennis V Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005, 1956.
- Quan Long. Multimodal information gain in bayesian design of experiments, 2021.
- David JC MacKay. Information-based objective functions for active data selection. *Neural computation*, 4(4): 590–604, 1992.
- Sören Mindermann, Jan Brauner, Muhammed Razzak, Mrinank Sharma, Andreas Kirsch, Winnie Xu, Benedikt Höltingen, Aidan N. Gomez, Adrien Morisot, Sebastian Farquhar, and Yarin Gal. Prioritized training on points that are learnable, worth learning, and not yet learnt, 2022.
- Rafail Ostrovsky, Yuval Rabani, Leonard J Schulman, and Chaitanya Swamy. The effectiveness of lloyd-type methods for the k-means problem. *Journal of the ACM (JACM)*, 59(6):1–22, 2013.
- Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training, 2021.
- Burr Settles. Active learning literature survey. 2009.
- Burr Settles, Mark Craven, and Soumya Ray. Multiple-instance active learning. *Advances in neural information processing systems*, 20, 2007.
- Shengyang Sun, Daniele Calandriello, Huiyi Hu, Ang Li, and Michalis Titsias. Information-theoretic online memory selection for continual learning, 2022.
- Dustin Tran, Jeremiah Liu, Michael W Dusenberry, Du Phan, Mark Collier, Jie Ren, Kehang Han, Zi Wang, Zelda Mariet, Huiyi Hu, et al. Plex: Towards reliability using pretrained large model extensions. *arXiv preprint arXiv:2207.07411*, 2022.
- Chaoqi Wang, Shengyang Sun, and Roger Grosse. Beyond marginal uncertainty: How accurately can bayesian regression models estimate posterior predictive correlations?, 2020.
- Kai Yu, Jinbo Bi, and Volker Tresp. Active learning via transductive experimental design. In *Proceedings of the 23rd international conference on Machine learning*, pages 1081–1088, 2006.

A Fisher Information: Additional Derivations & Proofs

Proposition 4.2. *Like the observed information, the Fisher information is additive:*

$$F(\{x_i\}, \omega^*) = \sum_i F(x_i, \omega^*). \quad (23)$$

Proof. This follows immediately from $Y_i \perp\!\!\!\perp Y_j \mid x_i, x_j, \omega^*$ for $i \neq j$ and the additivity of the observed information:

$$F(\{x_i\}, \omega^*) = \mathbb{E}_{P(\{y_i\}|\{x_i\}, \omega^*)} [H''[\{y_i\} \mid \{x_i\}, \omega^*]] = \mathbb{E}_{P(\{y_i\}|\{x_i\}, \omega^*)} \left[\sum_i H''[y_i \mid x_i, \omega^*] \right] \quad (80)$$

$$= \sum_i \mathbb{E}_{P(y_i|x_i, \omega^*)} [H''[y_i \mid x_i, \omega^*]] = \sum_i F(x_i, \omega^*). \quad (81)$$

□

Proposition 4.3. *The Fisher information is equivalent to:*

$$F(x, \omega^*) = \mathbb{E}_{P(y|x, \omega^*)} [H'[y \mid x, \omega^*]^T H'[y \mid x, \omega^*]] = \text{Cov}[H'[Y \mid x, \omega^*]]. \quad (24)$$

To prove Proposition 4.3, we use the two lemmas below:

Lemma 4.4. *For the Jacobian $H'[y \mid x, \omega^*]$, we have:*

$$H'[y \mid x, \omega^*] = -\frac{\nabla_\omega p(y \mid x, \omega^*)}{p(y \mid x, \omega^*)}, \quad (25)$$

and for the Hessian $H''[y \mid x, \omega^*]$, we have:

$$H''[y \mid x, \omega^*] = H'[y \mid x, \omega^*]^T H'[y \mid x, \omega^*] - \frac{\nabla_\omega^2 p(y \mid x, \omega^*)}{p(y \mid x, \omega^*)}. \quad (26)$$

Proof. The result follows immediately from the application of the rules of multivariate calculus. □

Lemma 4.5. *The following expectations over the model's own predictions vanish:*

$$\mathbb{E}_{P(y|x, \omega^*)} [H'[y \mid x, \omega^*]] = 0, \quad (27)$$

$$\mathbb{E}_{P(y|x, \omega^*)} \left[\frac{\nabla_\omega^2 p(y \mid x, \omega^*)}{p(y \mid x, \omega^*)} \right] = 0. \quad (28)$$

Proof. We use the previous equivalences and rewrite the expectations as integral; the results follows:

$$\mathbb{E}_{P(y|x, \omega^*)} [H'[y \mid x, \omega^*]] = \mathbb{E}_{P(y|x, \omega^*)} [-\nabla_\omega \log p(y \mid x, \omega^*)] = -\mathbb{E}_{P(y|x, \omega^*)} \left[\frac{\nabla_\omega p(y \mid x, \omega^*)}{p(y \mid x, \omega^*)} \right] \quad (82)$$

$$= -\int \nabla_\omega p(y \mid x, \omega^*) dy = -\nabla_\omega \int p(y \mid x, \omega^*) dy = -\nabla_\omega 1 = 0, \quad (83)$$

$$\mathbb{E}_{P(y|x, \omega^*)} \left[\frac{\nabla_\omega^2 p(y \mid x, \omega^*)}{p(y \mid x, \omega^*)} \right] = \int \nabla_\omega^2 p(y \mid x, \omega^*) dy = \nabla_\omega^2 \int p(y \mid x, \omega^*) dy = \nabla_\omega^2 1 = 0. \quad (84)$$

□

Proof of Proposition 4.3. With the previous lemma, we have the following.

$$\text{Cov}[H'[Y \mid x, \omega^*]] = \mathbb{E}[H'[Y \mid x, \omega^*]^T H'[Y \mid x, \omega^*]] - \underbrace{\mathbb{E}[H'[Y \mid x, \omega^*]^T]}_{=0} \underbrace{\mathbb{E}[H'[Y \mid x, \omega^*]]}_{=0} \quad (85)$$

$$= \mathbb{E}[\mathbf{H}'[Y | x, \omega^*]^T \mathbf{H}'[Y | x, \omega^*]]. \quad (86)$$

For the expectation over the Hessian, we plug Lemma 4.4 into Lemma 4.5 and obtain:

$$\mathbf{F}(x, \omega^*) = \mathbb{E}_{\mathbf{p}(y|x, \omega^*)}[\mathbf{H}''[y | x, \omega^*]] = \mathbb{E}_{\mathbf{p}(y|x, \omega^*)} \left[\mathbf{H}'[y | x, \omega^*]^T \mathbf{H}'[y | x, \omega^*] - \frac{\nabla_{\omega}^2 \mathbf{p}(y | x, \omega^*)}{\mathbf{p}(y | x, \omega^*)} \right] \quad (87)$$

$$= \mathbb{E}_{\mathbf{p}(y|x, \omega^*)}[\mathbf{H}'[y | x, \omega^*]^T \mathbf{H}'[y | x, \omega^*]] - 0 = \text{Cov}[\mathbf{H}'[Y | x, \omega^*]]. \quad (88)$$

□

A.1 Special Case: Exponential Family

Proposition 4.6. *The Fisher information $\mathbf{F}(\{x_i\}, \omega^*)$ for a model $\mathbf{p}(y | \hat{z} = \hat{f}(x; \omega^*))$ is equivalent to:*

$$\mathbf{F}(x, \omega^*) = \nabla_{\omega} \hat{f}(x; \omega^*)^T \mathbb{E}_{\mathbf{p}(y|x, \omega^*)}[\nabla_{\hat{z}}^2 \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)]] \nabla_{\omega} \hat{f}(x; \omega^*), \quad (29)$$

where $\nabla_{\hat{z}}^2 \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)]$ is short for $\nabla_{\hat{z}}^2 \mathbf{H}[y | \hat{z}]|_{\hat{z}=\hat{f}(x; \omega^*)}$.

Proof. We apply the second equivalence in Proposition 4.3 twice:

$$\mathbf{F}(x, \omega^*) = \text{Cov}[\mathbf{H}'[Y | x, \omega^*]] = \text{Cov}[\nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}} \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)] \nabla_{\omega} \hat{f}(x; \omega^*)] \quad (89)$$

$$= \nabla_{\omega} \hat{f}(x; \omega^*)^T \text{Cov}[\nabla_{\hat{z}} \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)]] \nabla_{\omega} \hat{f}(x; \omega^*) \quad (90)$$

$$= \nabla_{\omega} \hat{f}(x; \omega^*)^T \mathbb{E}_{\mathbf{p}(y|x, \omega^*)}[\nabla_{\hat{z}}^2 \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)]] \nabla_{\omega} \hat{f}(x; \omega^*) \quad (91)$$

□

A.2 Special Case: Generalized Linear Models

Proposition 4.9. *The Hessian $\mathbf{H}''[y | x, \omega^*]$ of a GLM is independent of y .*

$$\mathbf{H}''[y | x, \omega^*] = \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)] \nabla_{\omega} \hat{f}(x; \omega^*) \quad (33)$$

$$= \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 A(w^T x) \nabla_{\omega} \hat{f}(x; \omega^*). \quad (34)$$

Proof.

$$\mathbf{H}''[y | x, \omega^*] \quad (92)$$

$$= \nabla_{\omega}[\mathbf{H}'[y | x, \omega^*]] \quad (93)$$

$$= \nabla_{\omega}[\nabla_{\hat{z}} \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)] \nabla_{\omega} \hat{f}(x; \omega^*)] \quad (94)$$

$$= \nabla_{\hat{z}} \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)] \underbrace{\nabla_{\omega}^2 \hat{f}(x; \omega^*)}_{=\nabla_{\omega}^2 [w^T x]=0} + \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 \mathbf{H}[y | \hat{z} = \hat{f}(x; \omega^*)] \nabla_{\omega} \hat{f}(x; \omega^*) \quad (95)$$

$$= \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 A(w^T x) \nabla_{\omega} \hat{f}(x; \omega^*). \quad (96)$$

□

Proposition 4.11. *When $\hat{f}(x; \omega) : \mathbb{R}^D \rightarrow \mathbb{R}^C$, where C is the number of classes (outputs) and D is the number of input dimensions, and $\omega \in \mathbb{R}^{D \times C}$, and assuming the parameters are flattened into a single vector for the Jacobian, $\nabla_{\omega} \hat{f}(x; \omega^*) = \text{Id}_C \otimes x^T \in \mathbb{R}^{C \times (C \cdot D)}$, where \otimes denotes the Kronecker product, we have:*

$$\nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\hat{z}}^2 A(w^T x) \nabla_{\omega} \hat{f}(x; \omega^*) = \nabla_{\hat{z}}^2 A(w^T x) \otimes x x^T. \quad (37)$$

Proof. We begin with a few statements that lead to the conclusion step by step, where $x \in \mathbb{R}^D$, $A \in \mathbb{R}^{C \times C}$, $G \in \mathbb{R}^{C \times (C \cdot D)}$:

$$(x x^T)_{ij} = x_i x_j \quad (97)$$

$$(Id_C \otimes x^T)_{c, d D+i} = x_i \cdot \mathbb{1}\{c = d\} \quad (98)$$

$$(G^T A G)_{ij} = \sum_{k,l} G_{ki} A_{kl} G_{lj} \quad (99)$$

$$(A \otimes x x^T)_{c D+i, d D+j} = A_{cd} x_i x_j, \quad (100)$$

$$((Id_C \otimes x^T)^T A (Id_C \otimes x^T))_{c D+i, d D+j} = \sum_{k,l} (Id_C \otimes x^T)_{k, c D+i} A_{kl} (Id_C \otimes x^T)_{l, d D+j} \quad (101)$$

$$= \sum_{k,l} x_i \cdot \mathbb{1}\{k = c\} A_{kl} x_j \cdot \mathbb{1}\{l = d\} \quad (102)$$

$$= x_i A_{cd} x_j \quad (103)$$

$$= (A \otimes x x^T)_{c D+i, d D+j}. \quad (104)$$

$$\implies \nabla_{\omega} \hat{f}(x; \omega^*)^T \nabla_{\tilde{z}}^2 A(\omega^T x) \nabla_{\omega} \hat{f}(x; \omega^*) = \nabla_{\tilde{z}}^2 A(\omega^T x) \otimes x x^T. \quad (105)$$

□

Proposition 4.10. *For a model such that the Hessian $H''[y | x, \omega^*]$ is independent of y , we have:*

$$F(x, \omega^*) = H''[y^* | x, \omega^*] \quad (35)$$

for any y^* , and:

$$\mathbb{E}_{P(y|x)}[H''[y | x, \omega^*]] = F(x, \omega^*). \quad (36)$$

Proof. This follows directly from Proposition 4.7. In particular, we have:

$$F(x, \omega^*) = \mathbb{E}_{P(y|x, \omega^*)}[H''[y | x, \omega^*]] = H''[y^* | x, \omega^*], \quad (106)$$

where we have fixed y^* to an arbitrary value. □

B Approximating Information Quantities

B.1 Approximate Expected Information Gain

Lemma 5.1. *For symmetric, positive-semidefinite matrices A , we have (with equality iff $A = 0$):*

$$\log \det(A + Id) \leq \text{tr}(A). \quad (49)$$

Proof. When A is positive semidefinite and symmetric, its eigenvalues $(\lambda_i)_i$ are real and non-negative. Moreover, $A + Id$ has eigenvalues $(\lambda_i + 1)_i$; $\det(A + Id) = \prod_i (\lambda_i + 1)$; and $\text{tr } A = \sum_i \lambda_i$. These properties easily follow from the respective eigenvalue decomposition. Thus, we have:

$$\log \det(A + Id) \leq \log \prod_i (\lambda_i + 1) = \sum_i \log(\lambda_i + 1) \leq \sum_i \lambda_i = \text{tr}(A), \quad (107)$$

where we have used $\log(x + 1) \leq x$ iff equality for $x = 0$. □

General Case. In the main text, we only skimmed the general case and mentioned the main assumption. Here, we look at the general case in detail.

For the general case, we need to make strong approximations to be able to pursue a similar derivation. First, we cannot drop the expectation; instead, we note that the log determinant is a concave function on the

positive-semidefinite symmetric cone (Cover and Thomas, 1988), and we can use Jensen’s inequality on the log determinant term from Equation (43) as follows:

$$\mathbb{E}_{\mathbf{p}(\{y_i^{\text{acq}}\}|\{x_i^{\text{acq}}\})}[\log \det (\mathbf{H}''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*] \mathbf{H}''[\omega^*]^{-1} + Id)] \quad (108)$$

$$\leq \log \det \left(\mathbb{E}_{\mathbf{p}(\{y_i^{\text{acq}}\}|\{x_i^{\text{acq}}\})}[\mathbf{H}''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*]] \mathbf{H}''[\omega^*]^{-1} + Id \right). \quad (109)$$

Second, we need use the following approximation:

$$\mathbf{p}(\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}) \approx \mathbf{p}(\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*) \quad (110)$$

to obtain a Fisher information and use its additivity. That is, we obtain:

$$\mathbb{E}_{\mathbf{p}(\{y_i^{\text{acq}}\}|\{x_i^{\text{acq}}\}, \omega^*)}[\mathbf{H}''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*]] = \mathbf{F}(\{x_i^{\text{acq}}\}, \omega^*) = \sum_i \mathbf{F}(x_i^{\text{acq}}, \omega^*). \quad (111)$$

Plugging all of this together and applying Lemma 5.1, we obtain the same final approximation:

$$\mathbf{I}[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}] = \dots \approx \frac{1}{2} \mathbb{E}_{\mathbf{p}(\{y_i^{\text{acq}}\}|\{x_i^{\text{acq}}\})}[\log \det (\mathbf{H}''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*] \mathbf{H}''[\omega^*]^{-1} + Id)] \quad (112)$$

$$\leq \frac{1}{2} \log \det \left(\mathbb{E}_{\mathbf{p}(\{y_i^{\text{acq}}\}|\{x_i^{\text{acq}}\})}[\mathbf{H}''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*]] \mathbf{H}''[\omega^*]^{-1} + Id \right) \quad (113)$$

$$\approx \frac{1}{2} \log \det \left(\mathbb{E}_{\mathbf{p}(\{y_i^{\text{acq}}\}|\{x_i^{\text{acq}}\}, \omega^*)}[\mathbf{H}''[\{y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \omega^*]] \mathbf{H}''[\omega^*]^{-1} + Id \right) \quad (114)$$

$$= \frac{1}{2} \log \det \left(\sum_i \mathbf{F}(x_i^{\text{acq}}, \omega^*) \mathbf{H}''[\omega^*]^{-1} + Id \right) \quad (115)$$

$$\leq \frac{1}{2} \sum_i \text{tr} (\mathbf{F}(x_i^{\text{acq}}, \omega^*) \mathbf{H}''[\omega^*]^{-1}). \quad (116)$$

Unlike in the case of generalized linear models, a stronger assumption was necessary to reach the same result. Alternatively, we could use the GGN approximation, which leads to the same result. It does not lead to an upper bound, however.

B.2 Approximate Expected Predicted Information Gain

In the main text, we only briefly referred to not knowing a principled way to arrive at the same result of Proposition 5.4 for the general case. This is because unlike for the expected information gain, the Fisher information for an acquisition candidate now lies within an matrix inversion. Even if we use the fact that $\log \det(Id + XY^{-1})$ is concave in X and convex in Y , we end up with:

$$\dots \quad (117)$$

$$\approx \frac{1}{2} \mathbb{E}_{\mathbf{p}(y^{\text{eval}}, y^{\text{acq}} | x^{\text{eval}}, x^{\text{acq}}) \mathbf{p}(x^{\text{eval}})}[\log \det (\mathbf{H}''[y^{\text{eval}} | x^{\text{eval}}, \omega^*] (\mathbf{H}''[y^{\text{acq}} | x^{\text{acq}}, \omega^*] + \mathbf{H}''[\omega^*]^{-1} + Id)] \quad (118)$$

$$\leq \frac{1}{2} \mathbb{E}_{\mathbf{p}(y^{\text{acq}} | x^{\text{acq}})}[\log \det (\mathbb{E}_{\mathbf{p}(y^{\text{eval}}, x^{\text{eval}})}[\mathbf{H}''[y^{\text{eval}} | x^{\text{eval}}, \omega^*]] (\mathbf{H}''[y^{\text{acq}} | x^{\text{acq}}, \omega^*] + \mathbf{H}''[\omega^*]^{-1} + Id)] \quad (119)$$

$$\geq \frac{1}{2} \log \det (\mathbb{E}_{\mathbf{p}(y^{\text{eval}}, x^{\text{eval}})}[\mathbf{H}''[y^{\text{eval}} | x^{\text{eval}}, \omega^*]] (\mathbb{E}_{\mathbf{p}(y^{\text{acq}} | x^{\text{acq}})}[\mathbf{H}''[y^{\text{acq}} | x^{\text{acq}}, \omega^*]] + \mathbf{H}''[\omega^*]^{-1} + Id) \quad (120)$$

$$= \frac{1}{2} \log \det (\mathbb{E}_{\mathbf{p}(x^{\text{eval}})}[\mathbf{F}(x^{\text{eval}}, \omega^*)] (\mathbf{F}(x^{\text{acq}}, \omega^*) + \mathbf{H}''[\omega^*]^{-1} + Id) \quad (121)$$

$$\geq \frac{1}{2} \text{tr} (\mathbb{E}_{\mathbf{p}(x^{\text{eval}})}[\mathbf{F}(x^{\text{eval}}, \omega^*)] (\mathbf{F}(x^{\text{acq}}, \omega^*) + \mathbf{H}''[\omega^*]^{-1})). \quad (122)$$

Note the $\leq \dots \geq$, which invalidates the chain. Obviously, the errors could cancel out, but a principled statement seems hardly possible using this deduction.

B.3 Approximate Predictive Information Gain

Similar to Proposition 5.3, we can approximate the predictive information gain. We assume that we have access to an (empirical) distribution $\hat{\mathbf{p}}(x^{\text{eval}}, y^{\text{eval}})$:

Proposition B.1. We have (where we take the expectation over $\hat{p}(x^{eval}, y^{eval})$):

$$\arg \max_{x^{acq}} I[Y^{eval}, y^{acq} | X^{eval}, x^{acq}, \mathcal{D}^{train}] = \arg \min_{x^{acq}} I[\Omega; Y^{eval} | X^{eval}, y^{acq}, x^{acq}, \mathcal{D}^{train}] \quad (123)$$

with

$$I[\Omega; Y^{eval} | X^{eval}, y^{acq}, x^{acq}, \mathcal{D}^{train}] \quad (124)$$

$$\begin{aligned} &= \mathbb{E}_{\hat{p}(x^{eval}, y^{eval})} I[\Omega; y^{eval} | x^{eval}, y^{acq}, x^{acq}, \mathcal{D}^{train}] \\ &\approx \mathbb{E}_{\hat{p}(x^{eval}, y^{eval})} \left[\frac{1}{2} \log \det (H''[y^{eval} | x^{eval}, \omega^*] (H''[y^{acq} | x^{acq}, \omega^*] + H''[\omega^* | \mathcal{D}^{train}])^{-1} + Id) \right] \end{aligned} \quad (125)$$

$$\leq \frac{1}{2} \log \det (\mathbb{E}_{\hat{p}(x^{eval}, y^{eval})} [H''[y^{eval} | x^{eval}, \omega^*]] (H''[y^{acq} | x^{acq}, \omega^*] + H''[\omega^* | \mathcal{D}^{train}])^{-1} + Id) \quad (126)$$

$$\leq \frac{1}{2} \text{tr} (\mathbb{E}_{\hat{p}(x^{eval}, y^{eval})} [H''[y^{eval} | x^{eval}, \omega^*]] (H''[y^{acq} | x^{acq}, \omega^*] + H''[\omega^* | \mathcal{D}^{train}])^{-1}). \quad (127)$$

All of this follows immediately. Only for the second inequality, we need to use Jensen's inequality and that the log determinant is on the positive semidefinite symmetric cone (Cover and Thomas, 1988). Like for the information gain, there is no difference between having access to labels or not when we have a GLM or use the GGN approximation.

B.4 Approximate Joint (Expected) Predictive Information Gain

A comparison of EPIG and JEPIG shows that JEPIG does not require an expectation over $\hat{p}(x^{eval})$ but uses a set of *evaluation samples* $\{x_i^{eval}\}$. As such, we can easily adapt Proposition 5.4 to JEPIG and obtain:

Proposition B.2 (JEPIG). For a generalized linear model (or with the GGN approximation), we have:

$$\arg \max_{\{x_i^{acq}\}} I[\{Y_i^{eval}\}; \{Y_i^{acq}\} | \{x_i^{eval}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] = \arg \min_{\{x_i^{acq}\}} I[\Omega; \{Y_i^{eval}\} | \{x_i^{eval}\}, \{Y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] \quad (128)$$

with

$$\begin{aligned} &I[\Omega; \{Y_i^{eval}\} | \{x_i^{eval}\}, \{Y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] \\ &\approx \frac{1}{2} \log \det (F(\{x_i^{eval}\}, \omega^*) (F(\{x_i^{acq}\}, \omega^*) + H''[\omega^* | \mathcal{D}^{train}])^{-1} + Id) \end{aligned} \quad (129)$$

$$\leq \frac{1}{2} \text{tr} (F(\{x_i^{eval}\}, \omega^*) (F(\{x_i^{acq}\}, \omega^*) + H''[\omega^* | \mathcal{D}^{train}])^{-1}). \quad (130)$$

And similarly for JPIG, we obtain:

Proposition B.3 (JPIG). For a generalized linear model (or with the GGN approximation), we have:

$$\arg \max_{\{x_i^{acq}\}} I[\{y_i^{eval}\}; \{y_i^{acq}\} | \{x_i^{eval}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] = \arg \min_{\{x_i^{acq}\}} I[\Omega; \{y_i^{eval}\} | \{x_i^{eval}\}, \{y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] \quad (131)$$

with

$$\begin{aligned} &I[\Omega; \{y_i^{eval}\} | \{x_i^{eval}\}, \{Y_i^{acq}\}, \{x_i^{acq}\}, \mathcal{D}^{train}] \\ &\approx \frac{1}{2} \log \det (H''[\{y_i^{eval}\} | \{x_i^{eval}\}, \omega^*] (H''[\{y_i^{acq}\} | \{x_i^{acq}\}, \omega^*] + H''[\omega^* | \mathcal{D}^{train}])^{-1} + Id) \end{aligned} \quad (132)$$

$$\leq \frac{1}{2} \text{tr} (H''[\{y_i^{eval}\} | \{x_i^{eval}\}, \omega^*] (H''[\{y_i^{acq}\} | \{x_i^{acq}\}, \omega^*] + H''[\omega^* | \mathcal{D}^{train}])^{-1}). \quad (133)$$

C Similarity Matrices and One-Sample Approximations of the Fisher Information

Proposition 6.1. *Given $\mathcal{D}^{\text{train}}$, $\{x_i^{\text{acq}}\}$ and (sampled) $\{y_i^{\text{acq}}\}$, we have for the EIG:*

$$\mathbb{I}[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}] \lesssim \frac{1}{2} \log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{acq}} | \omega^*] + Id) \quad (74)$$

Proof.

$$\mathbb{I}[\Omega; \{Y_i^{\text{acq}}\} | \{x_i^{\text{acq}}\}, \mathcal{D}^{\text{train}}] \lesssim \frac{1}{2} \log \det (F(\{x_i^{\text{acq}}\}, \omega^*) H''[\omega^* | \mathcal{D}^{\text{train}}]^{-1} + Id) \quad (134)$$

$$= \frac{1}{2} \log \det ((F(\{x_i^{\text{acq}}\}, \omega^*) + H''[\omega^* | \mathcal{D}^{\text{train}}]) H''[\omega^* | \mathcal{D}^{\text{train}}]^{-1}) \quad (135)$$

$$\approx \frac{1}{2} \log \det \left((\hat{H}'[\mathcal{D}^{\text{acq}} | \omega^*]^T \hat{H}'[\mathcal{D}^{\text{acq}} | \omega^*] + H''[\omega^* | \mathcal{D}^{\text{train}}]) H''[\omega^* | \mathcal{D}^{\text{train}}]^{-1} \right) \quad (136)$$

$$= \frac{1}{2} \log \det \left(\hat{H}'[\mathcal{D}^{\text{acq}} | \omega^*] H''[\omega^* | \mathcal{D}^{\text{train}}]^{-1} \hat{H}'[\mathcal{D}^{\text{acq}} | \omega^*]^T + Id \right) \quad (137)$$

$$= \frac{1}{2} \log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{acq}} | \omega^*] + Id), \quad (138)$$

where we have used the matrix determinant lemma:

$$\det(AB + M) = \det(BM^{-1}A + Id) \det M. \quad (139)$$

□

Connection to the Joint (Expected) Predictive Information Gain. Following eq. (60), JEPIG can be decomposed as the difference between two EIG terms, which we can further split into three terms that are all only conditioned on $\mathcal{D}^{\text{train}}$:

$$\mathbb{I}[\{Y_i^{\text{eval}}\}; Y^{\text{acq}} | \{x_i^{\text{eval}}\}, x^{\text{acq}}, \mathcal{D}^{\text{train}}] \quad (140)$$

$$= \mathbb{I}[\Omega; \{Y_i^{\text{eval}}\} | \{x_i^{\text{eval}}\}, \mathcal{D}^{\text{train}}] - \mathbb{I}[\Omega; \{Y_i^{\text{eval}}\} | \{x_i^{\text{eval}}\}, Y^{\text{acq}}, x^{\text{acq}}, \mathcal{D}^{\text{train}}]$$

$$= \mathbb{I}[\Omega; \{Y_i^{\text{eval}}\} | \{x_i^{\text{eval}}\}, \mathcal{D}^{\text{train}}] - \mathbb{I}[\Omega; \{Y_i^{\text{eval}}\}, Y^{\text{acq}} | \{x_i^{\text{eval}}\}, x^{\text{acq}}, \mathcal{D}^{\text{train}}] + \mathbb{I}[\Omega; Y^{\text{acq}} | x^{\text{acq}}, \mathcal{D}^{\text{train}}] \quad (141)$$

Using Proposition 6.1, we can approximate this as:

$$\mathbb{I}[\{Y_i^{\text{eval}}\}; Y^{\text{acq}} | \{x_i^{\text{eval}}\}, x^{\text{acq}}, \mathcal{D}^{\text{train}}] \quad (142)$$

$$= \frac{1}{2} \log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{eval}} | \omega^*] + Id) - \frac{1}{2} \log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{acq}}, \mathcal{D}^{\text{eval}} | \omega^*] + Id) \quad (143)$$

$$+ \frac{1}{2} \log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{acq}} | \omega^*] + Id) \quad (144)$$

Further, we can use the approximation in Proposition 6.2 and find that the $\log \lambda$ terms cancel because we have $|\mathcal{D}^{\text{acq}}| + |\mathcal{D}^{\text{train}}| = |\mathcal{D}^{\text{acq}} \cup \mathcal{D}^{\text{train}}|$ obviously. Taking the limit, we obtain:

$$\mathbb{I}[\{Y_i^{\text{eval}}\}; Y^{\text{acq}} | \{x_i^{\text{eval}}\}, x^{\text{acq}}, \mathcal{D}^{\text{train}}] \quad (145)$$

$$\approx \frac{1}{2} \log \det (S[\mathcal{D}^{\text{eval}} | \omega^*] + \lambda Id) - \frac{1}{2} \log \det (S[\mathcal{D}^{\text{acq}}, \mathcal{D}^{\text{eval}} | \omega^*] + \lambda Id) + \frac{1}{2} \log \det (S[\mathcal{D}^{\text{acq}} | \omega^*] + \lambda Id) \quad (146)$$

$$\rightarrow \frac{1}{2} \log \det (S[\mathcal{D}^{\text{eval}} | \omega^*]) - \frac{1}{2} \log \det (S[\mathcal{D}^{\text{acq}}, \mathcal{D}^{\text{eval}} | \omega^*]) + \frac{1}{2} \log \det (S[\mathcal{D}^{\text{acq}} | \omega^*]). \quad (147)$$

Finally, the first term is independent of \mathcal{D}^{acq} , and if we are interested in approximately maximizing JEPIG, we can maximize as proxy:

$$\log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{acq}} | \omega^*] + Id) - \log \det (S_{H''[\omega^* | \mathcal{D}^{\text{train}}]}[\mathcal{D}^{\text{acq}}, \mathcal{D}^{\text{eval}} | \omega^*] + Id), \quad (148)$$

or

$$\log \det (S[\mathcal{D}^{\text{acq}} | \omega^*]) - \log \det (S[\mathcal{D}^{\text{acq}}, \mathcal{D}^{\text{eval}} | \omega^*]). \quad (149)$$

D Connection to Other Acquisition Functions in the Literature

D.1 SIMILAR (Kothawade et al., 2021) and PRISM (Kothawade et al., 2022)

Connection to LogDetMI. If we apply the Schur decomposition to $\log \det S[\mathcal{D}^{\text{acq}}, \mathcal{D}^{\text{eval}} | \omega^*]$ from eq. (149), we obtain the following:

$$\begin{aligned} \log \det S[\mathcal{D}^{\text{acq}}, \mathcal{D}^{\text{eval}} | \omega^*] \\ = \log \det S[\mathcal{D}^{\text{eval}} | \omega^*] + \log \det(S[\mathcal{D}^{\text{acq}} | \omega^*] - S[\mathcal{D}^{\text{acq}}; \mathcal{D}^{\text{eval}} | \omega^*] S[\mathcal{D}^{\text{eval}} | \omega^*]^{-1} S[\mathcal{D}^{\text{eval}}; \mathcal{D}^{\text{acq}} | \omega^*]), \end{aligned} \quad (150)$$

where $S[\mathcal{D}^{\text{acq}}; \mathcal{D}^{\text{eval}} | \omega^*]$ is the non-symmetric similarity matrix between \mathcal{D}^{acq} and $\mathcal{D}^{\text{eval}}$ etc.

Dropping $\log \det S[\mathcal{D}^{\text{eval}} | \omega^*]$ which is independent of \mathcal{D}^{acq} , we can also maximize:

$$\log \det S[\mathcal{D}^{\text{acq}} | \omega^*] - \log \det(S[\mathcal{D}^{\text{acq}} | \omega^*] - S[\mathcal{D}^{\text{acq}}; \mathcal{D}^{\text{eval}} | \omega^*] S[\mathcal{D}^{\text{eval}} | \omega^*]^{-1} S[\mathcal{D}^{\text{eval}}; \mathcal{D}^{\text{acq}} | \omega^*]), \quad (151)$$

which is exactly the LogDetMI objective from SIMILAR (Kothawade et al., 2021) and PRISM (Kothawade et al., 2022).

We can further rewrite this objective by extracting $S[\mathcal{D}^{\text{acq}} | \omega^*]$ from the second term, obtaining:

$$\log \det S[\mathcal{D}^{\text{acq}} | \omega^*] - \log \det(S[\mathcal{D}^{\text{acq}} | \omega^*] - S[\mathcal{D}^{\text{acq}}; \mathcal{D}^{\text{eval}} | \omega^*] S[\mathcal{D}^{\text{eval}} | \omega^*]^{-1} S[\mathcal{D}^{\text{eval}}; \mathcal{D}^{\text{acq}} | \omega^*]) \quad (152)$$

$$= -\log \det(Id - S[\mathcal{D}^{\text{acq}} | \omega^*]^{-1} S[\mathcal{D}^{\text{acq}}; \mathcal{D}^{\text{eval}} | \omega^*] S[\mathcal{D}^{\text{eval}} | \omega^*]^{-1} S[\mathcal{D}^{\text{eval}}; \mathcal{D}^{\text{acq}} | \omega^*]). \quad (153)$$

Connection to LogDetCMI. Using information-theoretic decompositions, it is easy to show that:

$$I[\{Y_i^{\text{eval}}\}; Y^{\text{acq}} | \{x_i^{\text{eval}}\}, x^{\text{acq}}, \{Y_i\}, \{x_i\}, \mathcal{D}^{\text{train}}] \quad (154)$$

$$= I[\{Y_i^{\text{eval}}\}; Y^{\text{acq}}, \{Y_i\} | \{x_i^{\text{eval}}\}, x^{\text{acq}}, \{x_i\}, \mathcal{D}^{\text{train}}] - I[\{Y_i^{\text{eval}}\}; \{Y_i\} | \{x_i\}, \mathcal{D}^{\text{train}}]. \quad (155)$$

These are two JEPIG terms, and using above approximations, including (153), leads to the LogDetCMI objective from Kothawade et al. (2021) and Kothawade et al. (2022):

$$\log \frac{\det(Id - S[\mathcal{D}^{\text{acq}} | \omega^*]^{-1} S[\mathcal{D}^{\text{acq}}; \mathcal{D}^{\text{eval}} | \omega^*] S[\mathcal{D}^{\text{eval}} | \omega^*]^{-1} S[\mathcal{D}^{\text{eval}}; \mathcal{D}^{\text{acq}} | \omega^*])}{\det(Id - S[\mathcal{D}^{\text{acq}}, \mathcal{D} | \omega^*]^{-1} S[\mathcal{D}^{\text{acq}}, \mathcal{D}; \mathcal{D}^{\text{eval}} | \omega^*] S[\mathcal{D}^{\text{eval}} | \omega^*]^{-1} S[\mathcal{D}^{\text{eval}}; \mathcal{D}^{\text{acq}}, \mathcal{D} | \omega^*])}. \quad (156)$$

D.2 Expected Gradient Length

Proposition 7.2. *The EIG for a candidate sample x^{acq} approximately lower-bounds the EGL:*

$$2 I[\Omega; Y^{\text{acq}} | x^{\text{acq}}] \lesssim \mathbb{E}_{\mathbf{p}(y^{\text{acq}} | x^{\text{acq}}, \omega^*)} \|H'[y^{\text{acq}} | x^{\text{acq}}, \omega^*]\|^2 + \text{const.} \quad (78)$$

Proof. The EIG is equal to the conditional entropy up to a constant term, via eq. (54) in Proposition 5.2:

$$I[\Omega; Y^{\text{acq}} | x^{\text{acq}}] \lesssim \frac{1}{2} \log \det (F(x^{\text{acq}}, \omega^*) + H''[\omega^* | \mathcal{D}^{\text{train}}]) + \text{const.} \quad (157)$$

We apply a diagonal approximation for the Fisher information and Hessian, noting that the determinant of the diagonal matrix upper-bounds the determinant of the full matrix:

$$\leq \frac{1}{2} \log \det (F_{\text{diag}}(x^{\text{acq}}, \omega^*) + H''_{\text{diag}}[\omega^* | \mathcal{D}^{\text{train}}]) + \text{const.} \quad (158)$$

$$= \frac{1}{2} \sum_k \log (F_{\text{diag},kk}(x^{\text{acq}}, \omega^*) + H''_{\text{diag},kk}[\omega^* | \mathcal{D}^{\text{train}}]) + \text{const.} \quad (159)$$

We use $\log x \leq x - 1$ and that $H''[\omega^* | \mathcal{D}^{\text{train}}]$ is constant:

$$\leq \frac{1}{2} \sum_k (F_{\text{diag},kk}(x^{\text{acq}}, \omega^*) + H''_{\text{diag},kk}[\omega^* | \mathcal{D}^{\text{train}}]) + \text{const.} \quad (160)$$

$$\leq \frac{1}{2} \sum_k F_{diag, kk}(x^{acq}, \omega^*) + \text{const.} \quad (161)$$

From Proposition 4.3, we know that the Fisher information is equivalent to the outer product of the Jacobians: $F(x^{acq}, \omega^*) = \mathbb{E}_{p(y^{acq}|x^{acq}, \omega^*)} [H'[y^{acq} | x^{acq}, \omega^*] H'[y^{acq} | x^{acq}, \omega^*]^T]$, and we finally obtain for the diagonal elements:

$$= \frac{1}{2} \sum_k \mathbb{E}_{p(y^{acq}|x^{acq}, \omega^*)} [H'_k[y^{acq} | x^{acq}, \omega^*]^2] + \text{const.} \quad (162)$$

$$= \frac{1}{2} \mathbb{E}_{p(y^{acq}|x^{acq}, \omega^*)} [\|H'[y^{acq} | x^{acq}, \omega^*]\|^2] + \text{const.} \quad (163)$$

□

D.3 Deep Learning on a Data Diet

Proposition 7.3. *The IG for a candidate sample x^{acq} approximately lower-bounds the gradient norm (GraNd) score at ω^* up to a second-order term:*

$$2 I[\Omega; y^{acq} | x^{acq}] \lesssim \mathbb{E}_{p(\omega | \mathcal{D}^{train})} [\|H'[y^{acq} | x^{acq}, \omega]\|^2] - \mathbb{E}_{p(\omega | \mathcal{D}^{train})} [\text{tr} \left(\frac{\nabla_\omega^2 p(y | x, \omega)}{p(y | x, \omega)} \right)] + \text{const.} \quad (79)$$

Proof. For any fixed ω^* , the IG is equal to the conditional entropy up to a constant term, via Proposition 5.3:

$$I[\Omega; Y^{acq} | x^{acq}] \lesssim \frac{1}{2} \log \det (H''[y^{acq} | x^{acq}, \omega^*] + H''[\omega^* | \mathcal{D}^{train}]) + \text{const.} \quad (164)$$

As in the previous proof, we apply a diagonal approximation for the Hessian, noting that the determinant of the diagonal matrix upper-bounds the determinant of the full matrix:

$$\leq \frac{1}{2} \log \det (H''_{diag}[y^{acq} | x^{acq}, \omega^*] + H''_{diag}[\omega^* | \mathcal{D}^{train}]) + \text{const.} \quad (165)$$

$$= \frac{1}{2} \sum_k \log (H''_{diag, kk}[y^{acq} | x^{acq}, \omega^*] + H''_{diag, kk}[\omega^* | \mathcal{D}^{train}]) + \text{const.} \quad (166)$$

Again, we use $\log x \leq x - 1$ and that $H''[\omega^* | \mathcal{D}^{train}]$ is constant:

$$\leq \frac{1}{2} \sum_k (H''_{diag, kk}[y^{acq} | x^{acq}, \omega^*] + H''_{diag, kk}[\omega^* | \mathcal{D}^{train}]) + \text{const.} \quad (167)$$

$$\leq \frac{1}{2} \sum_k H''_{diag, kk}[y^{acq} | x^{acq}, \omega^*] + \text{const.} \quad (168)$$

From Lemma 4.4, we know that the Hessian is equivalent to the outer product of the Jacobians plus a second-order term: $H''[y^{acq} | x^{acq}, \omega^*] = H'[y^{acq} | x^{acq}, \omega^*] H'[y^{acq} | x^{acq}, \omega^*]^T - \frac{\nabla_\omega^2 p(y^{acq} | x^{acq}, \omega^*)}{p(y^{acq} | x^{acq}, \omega^*)}$, and we finally obtain for the diagonal elements:

$$= \frac{1}{2} \sum_k H'_k[y^{acq} | x^{acq}, \omega^*]^2 - \frac{1}{2} \text{tr} \left(\frac{\nabla_\omega^2 p(y^{acq} | x^{acq}, \omega^*)}{p(y^{acq} | x^{acq}, \omega^*)} \right) + \text{const.} \quad (169)$$

$$= \frac{1}{2} \|H'[y^{acq} | x^{acq}, \omega^*]\|^2 - \frac{1}{2} \text{tr} \left(\frac{\nabla_\omega^2 p(y^{acq} | x^{acq}, \omega^*)}{p(y^{acq} | x^{acq}, \omega^*)} \right) + \text{const.} \quad (170)$$

Taking an expectation over $\omega^* \sim p(\omega^* | \mathcal{D}^{train})$ yields the statement. □