# p-Generalized Probit Regression and Scalable Maximum Likelihood Estimation via Sketching and Coresets 

Alexander Munteanu, Simon Omlor, Christian Peters<br>\{alexander.munteanu, simon.omlor, christian2.peters\}@tu-dortmund.de<br>Faculties of Statistics and Computer Science<br>TU Dortmund University


#### Abstract

We study the $p$-generalized probit regression model, which is a generalized linear model for binary responses. It extends the standard probit model by replacing its link function, the standard normal cdf, by a $p$-generalized normal distribution for $p \in[1, \infty)$. The $p$ generalized normal distributions (Subbotin, 1923) are of special interest in statistical modeling because they fit much more flexibly to data. Their tail behavior can be controlled by choice of the parameter $p$, which influences the model's sensitivity to outliers. Special cases include the Laplace, the Gaussian, and the uniform distributions. We further show how the maximum likelihood estimator for $p$-generalized probit regression can be approximated efficiently up to a factor of $(1+\varepsilon)$ on large data by combining sketching techniques with importance subsampling to obtain a small data summary called coreset.


## 1 INTRODUCTION

Probit regression is arguably one of the most successful models when considering predictive models with binary responses. Historically it preceded the logit model and those two are still the gold standards in many statistical domains (Cramer, 2002). Probit regression enjoys large popularity in toxicology (Lei and Sun, 2018), biostatistics (Varin and Czado, 2009), and also in econometrics (Gu et al., 2009; Moussa, 2019). Viewing it as a latent variable model yields an efficient Gibbs sampler for Bayesian analysis (Albert and Chib, 1993).

[^0]The probit model can be defined as a generalized linear model (McCullagh and Nelder, 1989) where the expected value of the response is connected to a linear predictor using the cumulative distribution function (cdf) of a standard normal distribution $\Phi(\cdot)$ as its link function, i.e.,

$$
\mathbb{E}[Y]=\Phi(Z \beta)
$$

In principle, the link function can be replaced by any cdf, e.g., by the cdf of a standard logistic distribution for logistic regression. In what follows, we extend the probit model to a $p$-generalized probit model by employing the cdf of a $p$-generalized normal distribution (Subbotin, 1923) as its link function. We work with a standardized form given in (Kalke and Richter, 2013):

$$
\Phi_{p}(x)=\frac{p^{1-1 / p}}{2 \Gamma(1 / p)} \int_{-\infty}^{x} \exp \left(-|t|^{p} / p\right) d t, x \in \mathbb{R}, p>0 .
$$

This family of distributions is of special interest in statistical modeling. It fits more flexibly to data because the tail behavior (kurtosis) can be controlled by choice of the parameter $p \geq 1^{1}$. Probability density functions (pdf) and cumulative distribution functions (cdf) for various values of $p$ are shown in Figure 1. The standard normal distribution and the standard probit model are obtained in the case $p=2$ with squared exponential tails. The case $p=1$ corresponds to the Laplace distribution with exponential tails, akin to the logistic distribution, and thus models more robustness to outliers. Finally, in the limiting case for $p \rightarrow \infty$ we obtain a uniform distribution over $[-1,1]$ whose tails are cut-off. Such a model is very sensitive to outliers and will fit the extreme data points, dominating other points with average behavior.
After introducing the generalized probit model, we set up the likelihood in order to learn the parameter $\beta \in \mathbb{R}^{d}$ via maximum likelihood estimation resp. minimization

[^1]

Figure 1: Probability density functions (pdf) and cumulative distribution functions (cdf) of $p$-generalized normal distributions for various values of $p$.
of the negative log-likelihood. The resulting $p$-probit loss function is strictly monotonic and convex for $p \geq 1$. We note that efficient routines are available for evaluating the probability density function, the cumulative distribution function and drawing random variables (Kalke and Richter, 2013) ${ }^{2}$. This allows to employ standard optimization methods like gradient descent or Newton's method (Bubeck, 2015). However, working with large scale data poses a severe limitation since each iteration scales at least linearly with the number of data points $n$. We note that stochastic gradient descent (SGD) and mini batching strategies provide scalable alternatives, and often work well in practice, but do not give the desired accuracy guarantees that we pursue in this paper. Examples and experiments where those methods fail were given in (Munteanu et al., 2018, 2021). Further, well-known random projections (e.g. Johnson-Lindenstrauss, Count Sketch, etc.) do not work for asymmetric functions since the sign of vectors are not preserved, and they have near linear lower bounds for $p>2$. Another disadvantage of random projections is that linear combinations of data decrease their interpretability. We thus focus on approximating the full gradient via importance sampling to cope with the limitations arising with large $n$.
Two methods called coresets and sketching (Phillips, 2017; Munteanu and Schwiegelshohn, 2018; Feldman, 2020) received a lot of attention recently and led to the so called sketch and solve paradigm (Woodruff, 2014). The first step of sketch and solve is a data reduction: we compute coresets and sketches which are concise data summaries that approximate the original full data with respect to a given loss function. In a second step, we pass the reduced data to the standard optimization algorithm that we would have used on the full data. But the algorithm now performs much more efficiently on the summary due to its reduced size. The theoretical analysis ensures that the result is close to the result of analyzing the original large data. In this paper we combine methods from the sketching and

[^2]coreset literature to facilitate an efficient estimation for $p$-generalized probit regression on large data.

### 1.1 Related Work

The $p$-generalized normal distribution was introduced by Subbotin (1923) and became widely popular in the late twentieth century (Goodman and Kotz, 1973; Osiewalski and Steel, 1993; Johnson et al., 1994). We refer to (Dytso et al., 2018) for an extensive survey on applications and analytical properties of the generalized normal distribution. Among other results, this reference provides an asymptotic characterization of the tails of generalized normal distributions, which we concretize in a non-asymptotic way, similar to the classic work of Gordon (1941) on the standard normal distribution. Another nice property is the decomposability into independent marginals, which characterizes the class of multivariate generalized normal distributions (Sinz et al., 2009; Dytso et al., 2018). In summary, the class of $p$-generalized normal distributions naturally extends the standard normal distribution and retains several of its useful and desirable analytical properties. Hereby, it offers more parametric flexibility allowing for tails that are either heavier $(p<2)$ or lighter $(p>2)$ than normal $(p=2)$ which makes it an excellent choice in many modeling scenarios (Dytso et al., 2018).

Most related to our work are coreset and sketching algorithms for linear $\ell_{p}$ regression (Clarkson, 2005; Dasgupta et al., 2009; Sohler and Woodruff, 2011; Meng and Mahoney, 2013; Woodruff and Zhang, 2013; Clarkson et al., 2016), which aims at minimizing $\|Z \beta-Y\|_{p}$, and can be seen as a standard linear model $Y=Z \beta+\eta$, where the error term $\eta$ follows a $p$-generalized normal distribution ${ }^{3}$. The earlier works relied on subsampling according to $\ell_{p}$ norms derived from a well-conditioned basis, whose approximation posed the computational bottleneck. Subsequent works improved the previous results significantly by approximating those bases via fast linear sketching techniques. To our knowledge, we are the first to study coresets and sketching for $\ell_{p}$ regression in the setting of generalized linear models.
Coresets for generalized linear models were studied by Molina et al. (2018), who gave an impossibility result for Poisson regression and a heuristic based on a latent variable log-normal count model. Starting with (Reddi et al., 2015; Huggins et al., 2016), a series of works focused on logistic regression. In (Munteanu et al., 2018) it was shown that there are no sublinear sketches or coresets for logistic regression, and to overcome this limitation the authors introduced a complexity parameter $\mu$ for compressing the data, which is related to the sta-

[^3]tistical modeling and their assumption that $\mu$ is small is naturally met on real world data and applications. The theoretical bounds have been recently improved (Mai et al., 2021) using $\ell_{1}$ Lewis weights instead of sampling according to the square root of $\ell_{2}$ leverage scores. However, the practical performance was only slightly improved at the cost of a $O(\log \log n)$ factor increase in the running time and number of passes over the data, required for approximating the Lewis weights. Recently, Munteanu et al. (2021) developed the first oblivious linear sketch for solving logistic regression up to $O(1)$ error in a single pass over a turnstile stream.

A different series of related works studied regularization as a means to overcome the lower bounds (Tolochinsky and Feldman, 2018; Tukan et al., 2020). Recently it was shown that if the regularization is strong enough then a uniform sample suffices (Samadian et al., 2020). In this paper we do not rely on regularization because a coreset for the regularized version of a problem does not yield a coreset for the unregularized version. On the other hand, if we have a coreset for the unregularized setting, it remains valid under any regularization term that is non-negative and does not depend on the data. This holds for the most common cases, e.g. for ridge, LASSO and elastic net regularization.

### 1.2 Our Contribution

A simple observation is that the tails of the logistic and the 1-generalized normal distributions are both exponential, and thus very similar up to a constant factor ${ }^{4}$. We thus expect similar results for the generalized linear models employing those cdfs as their link functions with respect to both, lower and upper bounds. Moreover, the intuition behind this paper is that we can generalize the aforementioned notion of $\mu$-complexity and hereby extend the class of generalized linear models that admit small coresets in a natural way, to obtain results for the broader class of $p$-generalized probit regression. Specifically, our contributions are:

- We introduce the $p$-generalized probit model as a flexible framework for modeling binary data, and for classification.
- To facilitate an efficient and scalable maximum likelihood estimation for the parameters of the $p$-generalized probit model, we develop coreset constructions for the associated $p$-probit loss function via the sensitivity framework (Langberg and Schulman, 2010; Feldman and Langberg, 2011) and combine with sketching techniques of Woodruff and Zhang (2013) to obtain wellconditioned $\ell_{p}$ bases and to approximate $\ell_{p}$ leverage scores efficiently.

[^4]- To this end we derive analytical properties of the $p$ generalized normal distribution that continue and concretize the asymptotic tail bounds of Dytso et al. (2018) using similar techniques as the classic non-asymptotic work of Gordon (1941) on the tails of the standard normal distribution. This result may be of broader independent interest.
- We provide a new VC dimension bound for our weighted loss functions by a novel fine-grained analysis, that also improves the $O(d \log n)$ bound of Munteanu et al. (2018) for logistic regression to $O(d \log (\mu / \varepsilon))^{5}$.
- We conduct an empirical evaluation on benchmark data. We compare the case $p=1$ to logistic regression, assessing their proximity. We assess the results obtained from $p$-probit regression for different values of $p$. We further evaluate our coreset constructions for different values of $p$ with respect to their approximation accuracy and algorithmic efficiency.

In summary, we advance the statistical modeling of binary data, make contributions in the field of analytical properties of the $p$-generalized normal distributions, generalize existing coreset constructions to a broader class of loss functions, and demonstrate the practical relevance of our methods.

## 2 TECHNICAL OVERVIEW

Several details are omitted due to page limitations. All missing pieces can be found in the appendix.

### 2.1 Preliminaries

For a fixed constant $p \in[1, \infty)$, and an unknown paramter $\beta \in \mathbb{R}^{d}$ we define the $p$-generalized probit model as a generalized linear model (McCullagh and Nelder, 1989):

$$
\mathbb{E}[Y]=\Phi_{p}(Z \beta)
$$

whose link function is given by $\Phi_{p}(r)=$ $\frac{p^{1-1 / p}}{2 \Gamma(1 / p)} \int_{-\infty}^{r} \exp \left(-|t|^{p} / p\right) d t$, i.e., the cumulative distribution function of the $p$-generalized normal distribution. Suppose we observe data $\left\{\left(z_{i}, y_{i}\right)\right\}_{i=1}^{n}$, where we have row vectors $z_{i} \in \mathbb{R}^{d}$ and $y_{i} \in\{0,1\}$ for each $i \in[n]$. We have $\operatorname{Pr}\left[y_{i}=1\right]=\mathbb{E}\left[y_{i}\right]=\Phi_{p}\left(z_{i} \beta\right)$ and similarly $\operatorname{Pr}\left[y_{i}=0\right]=1-\Phi_{p}\left(z_{i} \beta\right)=\Phi_{p}\left(-z_{i} \beta\right)$. The likelihood of the model is given by

$$
\begin{aligned}
\mathcal{L}(\beta \mid Z, Y) & =\prod_{i=1}^{n} \Phi_{p}\left(z_{i} \beta\right)^{y_{i}} \Phi_{p}\left(-z_{i} \beta\right)^{\left(1-y_{i}\right)} \\
& =\prod_{i=1}^{n} \Phi_{p}\left(\left(2 y_{i}-1\right) \cdot z_{i} \beta\right)
\end{aligned}
$$

[^5]We set $x_{i}=-\left(2 y_{i}-1\right) z_{i}$ for all $i \in[n]$ for convenience of presentation ${ }^{6}$. Consequently the log-likelihood simplifies to $\ell(\beta \mid X, Y)=\sum_{i=1}^{n} \ln \left(\Phi_{p}\left(-x_{i} \beta\right)\right)$. Maximizing the likelihood is thus equivalent to minimizing the following $p$-probit loss function, where we add a weight vector $w \in \mathbb{R}_{>0}^{n}$ for technical reasons detailed below:

$$
f_{w}(X \beta)=\sum_{i=1}^{n}-\ln \left(\Phi_{p}\left(-x_{i} \beta\right)\right) \cdot w_{i}
$$

We omit the subscript whenever the weights are uniform, i.e., $w_{i}=1$ for all $i \in[n]$. Moreover, to simplify notations we define the individual loss function

$$
\begin{equation*}
g(r)=-\ln \left(\Phi_{p}(-r)\right) \tag{1}
\end{equation*}
$$

The above loss function $f_{w}$ can be minimized within the framework of convex optimization (Bubeck, 2015) and can be regarded as solved for small to medium sized data. We provide a derivation of the gradient and the Hessian in Appendix A for completeness.

### 2.2 Coresets for $p$-Probit Regression

We focus on the situation where the number of data points $n \gg d$ is very large. In this case, applying a standard algorithm directly is not a viable option due to being either too slow or even impossible requiring too much working memory. Following the sketch and solve paradigm, our goal is to reduce the data without losing much information, before we approximate the problem efficiently on the reduced data using a standard solver. More formally our goal is now to develop an $\varepsilon$-coreset:
Definition 2.1. A weighted $\varepsilon$-coreset $C=\left(X^{\prime}, w\right)$ for $f$ is a matrix $X^{\prime} \in \mathbb{R}^{k \times d}$ together with a weight vector $w \in \mathbb{R}_{>0}^{k}$ such that for all $\beta \in \mathbb{R}^{d}$ it holds that

$$
\left|f_{w}\left(X^{\prime} \beta\right)-f(X \beta)\right| \leq \varepsilon \cdot f(X \beta) .
$$

Small $\varepsilon$-coresets with $k \ll n$ cannot be obtained in general. There are examples where no $\varepsilon$-coreset of size $k=o(n / \log n)$ exist, even if $\varepsilon$ attains an arbitrarily large value (Munteanu et al., 2018). When $X^{\prime}$ is constrained to be a subset of the input, then the bound can be strengthened to $\Omega(n)$ (Tolochinsky and Feldman, 2018). Those impossibility results rely on the monotonicity of the loss function and thus extend to all $p$-generalized probit losses. To get around those strong limitations, Munteanu et al. (2018) introduced a parameter $\mu$ as a natural notion for the complexity of compressing the input matrix $X$ for logistic regression, which we adapt to $p$-generalized probit models to parameterize our results.

[^6]Definition 2.2. Let $X \in \mathbb{R}^{n \times d}$ be any matrix. For $a$ fixed $p \geq 1$ we define

$$
\mu_{p}(X)=\sup _{\beta \in \mathbb{R}^{d} \backslash\{0\}} \frac{\sum_{x_{i} \beta>0}\left|x_{i} \beta\right|^{p}}{\sum_{x_{i} \beta<0}\left|x_{i} \beta\right|^{p}}
$$

We say that $X$ is $\mu$-complex if $\mu_{p}(X) \leq \mu<\infty$.
Computing $\mu$ is hard in general and is mainly for parameterizing the size of coresets in the theoretical analysis. In practice we spend as much memory as we can afford and do not precompute $\mu$. Fortunately, for many real world data sets $\mu$ turns out to be sufficiently small (Munteanu et al., 2018).

Our main result shows that if $\mu$ is small, then there exists a small coreset $C$. In fact, the size of $C$ does not depend on $n$ at all and it can be computed efficiently in two passes over the data:
Theorem 1. If $X \in \mathbb{R}^{n \times d}$ is $\mu$-complex for any fixed $p \in[1, \infty)$ then with constant probability we can compute an $\varepsilon$-coreset $C=\left(X^{\prime}, w\right)$ for $p$-probit regression of size $k=O\left(\frac{S}{\varepsilon^{2}}\left(d \ln \left(\varepsilon^{-1} \mu\right) \ln S\right)\right)$ in two passes over the data, where

$$
S= \begin{cases}O(\mu d), & \text { for } p=2 \\ O\left(\mu d^{p}(d \log d)^{2}\right), & \text { for } p \in[1,2) \\ O\left(\mu d^{2 p}(d \log d)^{2}\right), & \text { for } p \in(2, \infty) .\end{cases}
$$

Algorithm 1 runs in $O(\operatorname{nnz}(X) d+\operatorname{poly}(d))^{7}$ time for $p \in[1,2]$ and in $O\left(\mathrm{nnz}(X) d+\operatorname{poly}(d) n^{1-\frac{2}{p}} \log n\right)$ time for $p>2$, were nnz denotes the number of non-zeros.

In the case $p=2$, which is of special importance since it corresponds to the standard probit regression model, we have the following improvements:
Corollary 2.3. Consider the setting of Theorem 1, for $p=2$. The running time can be reduced to $O(\operatorname{nnz}(X) \log n+\operatorname{poly}(d))$. Moreover there exists a single pass online algorithm (Algorithm 2) that runs in time $O\left(n d^{2}+\operatorname{poly}(d)\right)$ and computes a coreset of size

$$
O\left(\frac{\mu d^{2} \ln \left(\|X\|_{2}\right)}{\varepsilon^{2}} \ln \left(\varepsilon^{-1} \mu\right) \ln \left(\mu d \ln \left(\|X\|_{2}\right)\right)\right)
$$

where $\|X\|_{2}$ denotes the largest singular value of $X$.
The coreset of Theorem 1 or Corollary 2.3 can then be used to compute a $(1+\varepsilon)$-approximation for the optimal maximum likelihood estimator for $\beta$ :
Corollary 2.4. Let $\left(X^{\prime}, w\right)$ be a weighted $\varepsilon$-coreset for $f$. Let $\tilde{\beta} \in \operatorname{argmin}_{\beta \in \mathbb{R}^{d}} f_{w}\left(X^{\prime} \beta\right)$. Then it holds that $f(X \tilde{\beta}) \leq(1+3 \varepsilon) \min _{\beta \in \mathbb{R}^{d}} f(X \beta)$.

[^7]High Level Description of the Algorithm Before getting into the details we outline Algorithm 1:

1. We make a first pass to sketch the data for the purpose of estimating their individual importance.
2. We make another pass to subsample the data proportional to their importance to obtain a coreset.
3. We solve the reduced problem on the coreset using a standard algorithm for convex optimization.

This approach implements the sensitivity sampling framework. The importance measure that it builds upon is called sensitivity, which measures the worst case contribution of each input point to the objective function. For efficiency reasons we first compute a sketch of the data in one pass. In the second pass, the sketch is used to approximate the $\ell_{p}$ leverage scores, which upper bound the sensitivities of the input points. Hereby, we pass them one-by-one to a reservoir sampler to obtain the coreset. In the following paragraphs we will prove that the output has the desired coreset property. To this end we will outline the sensitivity framework, then establish the connection of our loss function to $\ell_{p}$ spaces. We will further bound the VC dimension for our loss function and bound the sensitivities in terms of $\ell_{p}$ leverage scores. We will then show how we can approximate the $\ell_{p}$ leverage scores by means of well-conditioned bases via sketching, to use them as importance measure in the actual algorithm. Putting all those pieces together will prove Theorem 1. Finally, we can solve the original problem approximately using gradient descent or other standard methods for convex optimization (see Bubeck, 2015) on the resulting coreset. We continue with a more detailed description of all pieces and then put them together.

The Sensitivity Framework To prove Theorem 1, we start with the so called sensitivity framework (Langberg and Schulman, 2010; Feldman and Langberg, 2011; Braverman et al., 2016; Feldman et al., 2020). This framework provides a meta theorem, that can be used to construct coresets via importance subsampling. The main parameters that we need to bound is the VC dimension $\Delta$ of a weighted set of functions that can roughly be thought of as the dimension of the parameter space, and the sensitivities of the input points. The latter quantify the worst-case contributions of single input points to the loss function. Recall Eqn. (1).
Definition 2.5 (Langberg and Schulman (2010)). For all $i \in[n]$, we define the sensitivity $\zeta_{i}$ of point $x_{i}$ by

$$
\zeta_{i}=\sup _{\beta \in \mathbb{R}^{d}} \frac{g\left(x_{i} \beta\right)}{f(X \beta)}
$$

The total sensitivity is given by $Z=\sum_{i=1}^{n} \zeta_{i}$

Calculating the sensitivities is often difficult in the sense that it requires to solve the problem to optimality first, before it allows us to obtain an approximation. Fortunately it suffices to calculate upper bounds $S=\sum_{i=1}^{n} s_{i} \geq \sum_{i=1}^{n} \zeta_{i}=Z$, which is often tractable, though it is crucial to control the increased total sensitivity bound $S$. Our bound $S$ will depend on $p, d$, and $\mu$ as detailed in Theorem 1. The meta theorem (Proposition B.3) then ensures that if we sample $k=O\left(\frac{S}{\varepsilon^{2}}\left(\Delta \ln S+\ln \left(\frac{1}{\delta}\right)\right)\right)$ points each with probability $p_{i}=\frac{s_{i}}{S}$, then the resulting subsample, reweighted by $w_{i}=\frac{1}{k p_{i}}$, is an $\varepsilon$-coreset with probability at least $1-\delta$. More formal details are deferred to Appendix B.

Properties of the Loss Function The first step for bounding the parameters of the sensitivity framework is to analyze the tails of the generalized normal distribution and relate its negative logarithm to $g(r) \approx r^{p} / p$ for $r \geq 0$, and to $\exp \left(-|r|^{p} / p\right)$ for $r<0$. This generalizes classic results on the tails of the standard normal distribution (Gordon, 1941). The following lemma makes this more precise and follows by combining the technical derivations of Lemma C. 1 in the appendix.
Lemma 2.6. The function $g$ (see Eqn. (1)) is convex and strictly increasing. Further for any $r \geq 0$ we have

$$
g^{\prime}(r) \geq r^{p-1}
$$

for any $r \geq 1$ we have

$$
g^{\prime}(r) \leq r^{p-1}+\frac{p-1}{r}
$$

and there exists a constant $c_{1}>0$ such that

$$
g(r) \geq c_{1} e^{-|r|^{p} / p}
$$

for any $r<0$.
Bounding the VC Dimension It is well-known (Huggins et al., 2016; Munteanu et al., 2018) that for uniform weights, the VC dimension of any invertible function of $X \beta$ can be bounded by relating to the class of affine separators, whose VC dimension equals $d+1$, see (Kearns and Vazirani, 1994). Munteanu et al. (2018) showed that the number of distinct weights for logistic regression can be bounded by $|W|=O(\log n)$ by rounding the sensitivities to the next powers of 2 , implying the VC dimension is bounded by $O(d \log n)$. The same arguments would apply to $p$-probit loss as well. In our fine-grained analysis we start by additionally rounding all sensitivities below $\frac{S}{n}$ to $\frac{S}{n}$ while increasing the total sensitivity to at most $2 S+n \frac{S}{n}=3 S$. Then we inspect large and small sensitivities separately, depending on their relation to a threshold $s_{0}$. For the large sensitivities $s_{i}>s_{0}$ we can show using Lemmas 2.6 and C.3
that their contribution can be approximated well when we replace $g(r)$ by a function $G^{+}(r)$ that equals zero for $r<0$, and $\frac{r^{p}}{p}$ for $r \geq 0$, and whose VC dimension is bounded by $O(d)$.
Lemma 2.7. Let $I_{1}$ be the index set of all data points with $s_{i}>s_{0}:=\frac{\mu S c \ln \left(p \varepsilon^{-1}\right)}{\varepsilon n}$ for some constant $c \in \mathbb{R}_{>0}$. Then for all $\beta \in \mathbb{R}^{d}$ it holds that

$$
\begin{aligned}
\sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right) & \leq \sum_{i \in I_{1}} g\left(x_{i} \beta\right) \\
& \leq(1+\varepsilon) \sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right)+\varepsilon \cdot \frac{n}{\mu}
\end{aligned}
$$

The additive error can be charged using a lower bound on the $p$-probit loss function.
Lemma 2.8. Assume $X \in \mathbb{R}^{n \times d}$ is $\mu$-complex. Then we have for any $\beta \in \mathbb{R}^{d}$ that

$$
f(X \beta)=\Omega\left(\frac{n}{\mu}(1+\ln (\mu))\right)
$$

For the remaining small sensitivities, they are bounded between narrow thresholds $\frac{S}{n} \leq s_{i} \leq s_{0}$ and by the previous rounding to the next power of 2 we can argue that there are only $|W|=O\left(\log \frac{\mu}{\varepsilon}\right)$ different weights. Overall we can bound the VC dimension by $O\left(d \log \frac{\mu}{\varepsilon}\right)$ using the following argument.
By the technical Corollary C. 6 in the appendix, our goal of obtaining a coreset for $f$ reduces to obtaining a coreset for the substitute function

$$
\tilde{f}(X \beta)=\sum_{i \in[n] \backslash I_{1}} g\left(x_{i} \beta\right)+\sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right) .
$$

To this end we set $\mathcal{F}_{1}=\left\{w_{i} G_{i}^{+} \mid i \in I_{1}\right\}$ where $G_{i}^{+}(\beta)=G^{+}\left(x_{i} \beta\right)$ and $\mathcal{F}_{2}=\left\{w_{i} g_{i} \mid i \in I_{2}:=[n] \backslash I_{1}\right\}$ where $g_{i}(\beta)=g\left(x_{i} \beta\right)$. Further we set $\mathcal{F}=\mathcal{F}_{1} \cup \mathcal{F}_{2}$ and show that the VC dimension of $\mathcal{F}$ can be bounded as desired:

Lemma 2.9. For the $V C$ dimension $\Delta$ of $\mathfrak{R}_{\mathcal{F}}$ we have

$$
\begin{aligned}
\Delta & \leq(d+1)\left(\log _{2}\left(\mu c \varepsilon^{-2}\right)+2\right) \\
& =O(d \log (\mu / \varepsilon))
\end{aligned}
$$

Bounding the Sensitivities Using the analytic bounds and the assumption of $\mu$-complex data we can relate the sensitivities for the negative domain to roughly $\frac{\mu}{n}$, and for the positive part to the $\ell_{p}$ leverage scores, defined by $u_{i}=\sup _{\beta \neq 0} \frac{\left|x_{i} \beta\right|^{p}}{\|X \beta\|_{p}^{p}}$, cf. (Dasgupta et al., 2009). This gives us the following lemma.
Lemma 2.10. There is a constant $c_{s}$ such that the sensitivity $\zeta_{i}$ of $x_{i}, i \in[n]$ for $\tilde{f}$ is bounded by

$$
\zeta_{i} \leq c_{s} \mu\left(\frac{1}{n}+u_{i}\right)
$$

Summing over the combined upper bound yields a total sensitivity of $S \leq \sum_{i \in[n]} c_{s} \mu\left(1 / n+u_{i}\right)=$ $O\left(\mu \sum_{i \in[n]} u_{i}\right)$, which dominates the size of the coreset; see Theorem 1. The sum of leverage scores can further be bounded by roughly $d^{O(p)}$ by means of $\ell_{p^{-}}$ well-conditioned bases for the column space of $X$.

Well-conditioned Bases via Sketching Techniques One can approximate the $\ell_{p}$ leverage scores using an orthonormal basis for the column space of $X$. Unfortunately this gives only an $n^{c}$-approximation for $p \neq 2$. We thus work with a generalization to so called well-conditioned bases. An $(\alpha, \beta, p)$-well-conditioned basis $V$ is a basis that preserves the norm of each vector well, as detailed in the following definition.
Definition 2.11 (Dasgupta et al. (2009)). Let $X$ be an $n \times m$ matrix of rank $d$, let $p \in[1, \infty)$, and let $q$ be its dual norm, i.e., $q \in(1, \infty]$ satisfying $\frac{1}{p}+\frac{1}{q}=1$. Then an $n \times d$ matrix $V$ is an $(\alpha, \beta, p)$-well-conditioned basis for the column space of $X$ if
(1) $\|V\|_{p}:=\left(\sum_{i \leq n, j \leq d}\left|V_{i j}\right|^{p}\right)^{1 / p} \leq \alpha$, and
(2) for all $z \in \mathbb{R}^{d},\|z\|_{q} \leq \beta\|V z\|_{p}$.

We say that $V$ is a p-well-conditioned basis for the column space of $X$ if $\alpha$ and $\beta$ are $d^{O(1)}$, independent of $m$ and $n$.

We can bound the leverage scores in terms of the rowwise $p$-norms of such a basis.
Lemma 2.12. Let $V$ be an ( $\alpha, \beta, p$ )-well-conditioned basis for the column space of $X$. Then it holds for all $i \in[n]$ that $u_{i} \leq \beta^{p}\left\|v_{i}\right\|_{p}^{p}$. As a direct consequence we have $\sum_{i=1}^{n} u_{i} \leq \beta^{p}\|V\|_{p}^{p} \leq(\alpha \beta)^{p}=d^{O(p)}$.

A prominent example of a well-conditioned basis is the aforementioned orthonormal basis for $\ell_{2}$, which can be obtained by QR-decomposition (or SVD) in $O\left(n d^{2}\right)$. Such a basis $Q$ is $(\sqrt{d}, 1,2)$-well-conditioned, since $\|Q\|_{F}=\sqrt{d}$ and $\|Q z\|_{2}=\|z\|_{2}$ due to rotational invariance of the $\ell_{2}$-norm. For general $p$ there exist so called Auerbach bases (Auerbach, 1930) with $\alpha=d$ and $\beta=1$, and approximations thereof can be computed in time $O\left(n d^{5} \log n\right)$ via Löwner-John ellipsoids (Clarkson, 2005; Dasgupta et al., 2009).

To get around this computational bottleneck, we first apply a sketching matrix to our data that yields a so called subspace embedding for approximating the $\ell_{p}$ norms of all vectors in the column space of $X$ with a distortion of roughly $O\left((d \log d)^{1 / p}\right)$. To this end we apply an embedding into $\ell_{\infty}$ using $1 / p$-powers of inverse exponential random variables (Andoni, 2017; Woodruff and Zhang, 2013) followed by a dimensionality reduction via sparse random embeddings (Clarkson and Woodruff, 2017; Cohen, 2016) to obtain a small sketch $\Pi X$.

Lemma 2.13 (Woodruff and Zhang (2013); Clarkson and Woodruff (2017)). There exists a random embedding matrix $\Pi \in \mathbb{R}^{n^{\prime} \times n}$ and $\gamma=O(d \log (d))$ such that

$$
\forall \beta \in \mathbb{R}^{d}: \frac{1}{\gamma^{1 / p}}\|X \beta\|_{p} \leq\|\Pi X \beta\|_{q} \leq \gamma^{1 / p}\|X \beta\|_{p}
$$

holds with constant probability, where $q=2, n^{\prime}=O\left(d^{2}\right)$ if $p \in[1,2]$ and $q=\infty, n^{\prime}=O\left(n^{1-\frac{2}{p}} \log n(d \log d)^{1+\frac{2}{p}}+\right.$ $\left.d^{5+4 p}\right)$ if $p \in(2, \infty)$. For $p=2$ we have $\gamma=2$. Further $\Pi X$ can be computed in $O(\mathrm{nnz}(X))$ time.

For $p \in[1,2]$, we can obtain $n^{\prime}=O(d \log d)$ in exchange for an increased running time $O(\operatorname{nnz}(X) \log d)$ by using the $\ell_{2}$ subspace embedding of Cohen (2016) as a replacement for (Clarkson and Woodruff, 2017). We further note that $\Omega\left(n^{1-\frac{2}{p}} \log n\right)$ is necessary for $p>2$ due to tight lower bounds for sketching the $p$ th frequency moments (Andoni et al., 2013).

Now we decompose the sketch $\Pi X=Q R$, which is fast due to its reduced size and argue that $V=X R^{-1}$ is $p$-well-conditioned, leading to the desired $d^{O(p)}$ bound.
Lemma 2.14. If $\Pi$ satisfies Lemma 2.13 and $\Pi X=$ $Q R$ is the $Q R$-decomposition of $\Pi X$ then $V=X R^{-1}$ is an $(\alpha, \beta, p)$-well-conditioned basis for the columnspace of $X$, where for $\gamma=O(d \log (d))$ we have

$$
(\alpha, \beta)= \begin{cases}(\sqrt{2 d}, \sqrt{2}), & \text { for } p=2 \\ \left(d \gamma^{1 / p}, \gamma^{1 / p}\right), & \text { for } p \in[1,2) \\ \left(d \gamma^{1 / p}, d \gamma^{1 / p}\right), & \text { for } p \in(2, \infty)\end{cases}
$$

We note that similar methods have been used before in the context of linear $\ell_{p}$ regression (Sohler and Woodruff, 2011; Woodruff and Zhang, 2013) and for approximating the $\ell_{2}$ leverage scores (Drineas et al., 2012).

### 2.3 Proof of Theorem 1

The previous argumentation yields a coreset whose size does not depend on $n$ and that can be calculated efficiently on large data. We prove our main results.

Proof of Theorem 1. We first describe the algorithm (Algorithm 1) and its running time: First we apply our sketching matrix $\Pi$ from Lemma 2.13 to compute $\Pi X$ in time $O(\mathrm{nnz}(X))$ in one pass over the data. The number of rows is $n^{\prime}=O\left(d^{2}\right)$ for $p \in[1,2]$ and $n^{\prime}=$ $O\left(n^{1-\frac{2}{p}} \log n \operatorname{poly}(d)\right)$ for $p>2$. Then we calculate the $Q R$-decomposition of $\Pi X=Q R$ in time $O\left(d^{4}\right)$ respectively in time $O\left(n^{1-\frac{2}{p}} \log n\right.$ poly $\left.(d)\right)$ depending on the value of $p$, which is faster than $O\left(n d^{2}\right)$ without sketching. In a second pass over the data we compute the row norms $\left\|v_{i}\right\|_{p}^{p}=\left\|x_{i} R^{-1}\right\|_{p}^{p}$ used in our sampling probabilities. We set $s_{i}=\frac{1}{n}+\beta^{p}\left\|v_{i}\right\|_{p}^{p}$. By Lemma 2.12,
$S_{0}=1+(\alpha \beta)^{p}$ is an upper bound for $\sum_{i=1}^{n} s_{i}$. Next we set $s_{i}^{\prime}=\max \left\{2^{\left\lceil\log _{2}\left(s_{i}\right)\right\rceil}, \frac{S_{0}}{n}\right\}$, i.e. we round $s_{i}$ to the next power of 2 such that $S^{\prime}=\sum_{i=1}^{n} s_{i}^{\prime} \leq 2 S_{0}+n \cdot \frac{S_{0}}{n}=$ $3 S_{0}$. As we calculate those values, we can feed the point $x_{i}$ augmented with the corresponding sampling weight $s_{i}^{\prime}$ directly to $k$ independent copies of a weighted reservoir sampler (Chao, 1982). The latter is an online algorithm and updates its sample in constant time. The second pass takes $O(\operatorname{nnz}(X) d+\operatorname{poly}(d))$ time for $p \in[1,2]$, respectively $O\left(\operatorname{nnz}(X) d+\operatorname{poly}(d) n^{1-\frac{2}{p}} \log n\right)$ for $p>2$. Lemma 2.10 yields $S=\sum_{i \in[n]} c_{s} \mu(1 / n+$ $\left.u_{i}\right)=O\left(\mu \sum_{i \in[n]} u_{i}\right)$ and by Lemma 2.12 we have that $\sum_{i \in[n]} u_{i} \leq(\alpha \beta)^{p}=d^{O(p)}$ where the values of $\alpha, \beta$ are detailed in Lemma 2.14. Using Lemmas 2.9, 2.10, and 2.12 to bound the parameters of Proposition B. 3 we get for the substitute function $\tilde{f}$ that $\forall \beta: \mid \tilde{f}_{w}\left(X^{\prime} \beta\right)-$ $\tilde{f}(X \beta) \mid \leq \varepsilon \tilde{f}(X \beta)$ with probability at least $1-\delta$. By Corollary C. 6 and Lemma C. 7 this implies with high probability that $\left(X^{\prime}, w\right)$ is a $7 \varepsilon$-coreset for $f$. Folding the constant into $\varepsilon$ completes the proof.

Proof of Corollary 2.3. If $p=2$ and $d=\omega(\ln n)$, we can use a Johnson-Lindenstrauss transform, i.e., a matrix $G \in \mathbb{R}^{d \times m}$ where $m=O(\ln (n))$ and whose entries are i.i.d. $G_{i j} \sim N\left(0, \frac{1}{m}\right)$ (Johnson and Lindenstrauss, 1984) to compute a $\frac{1}{2}$-approximation to the row norms: We have $\left\|v_{i}^{\prime}\right\|_{2}^{2}:=\left\|x_{i}\left(R^{-1} G\right)\right\|_{2}^{2} \geq\left\|x_{i} R^{-1}\right\|_{2}^{2} / 2$ for all $i \in[n]$ simultanously with constant probability. The running time reduces to $O(\mathrm{nnz}(X) \ln (n)+\operatorname{poly}(d))$. The online algorithm (Algorithm 2) is obtained by running the online $\ell_{2}$ leverage score algorithm of Chhaya et al. (2020) that recently extended the previous work of Cohen et al. (2020). Each row update takes $O\left(d^{2}\right)$ time except for at most $O(d)$ updates that take $O\left(d^{3}\right)$ time, implying $O\left(n d^{2}+\operatorname{poly}(d)\right)$ total running time. The slightly increased coreset size results from an increase of the total sensitivity by at most $\log \left(\|X\|_{2}\right)$ due to the online procedure (Chhaya et al., 2020).

## 3 EMPIRICAL EVALUATION

Our intention is to corroborate our theoretical results by investigating the following questions empirically:
(Q1) How does the 1-probit model compare to logistic regression?
(Q2) How do the $p$-probit models for different values of $p$ compare to one another?
(Q3) How accurate are the maximum likelihood (ML) estimators obtained from $p$-probit coresets?
(Q4) How fast can we obtain an accurate ML estimator from $p$-probit coresets?

All experiments were conducted on an AMD Ryzen 7 2700x processor ( 8 cores, $3.7 \mathrm{GHz}, 16 \mathrm{~GB}$ RAM).

Our results can be reproduced with our open Python implementation available at https://github.com/ cxan96/efficient-probit-regression. More information on data sets, pseudo code and several plots are in Appendix D.
I. Statistical Modeling Aspects: For (Q1) \& (Q2) no data reduction is applied. Specifically, only statistical modeling aspects are studied here.
(Q1) We investigate how closely the 1-probit model estimate equals logistic regression. This seems to be a natural question since both functions have similar tail behaviors as we have argued before in the introduction regarding their link functions. The similarities naturally extend to their loss functions. We observe that the logistic and 1-probit loss functions both converge to a linear function for positive arguments and to the exponential function for negative arguments. Near zero, however, the two loss functions differ more significantly. We compare the coefficients obtained for the two models on full data sets from public repositories, see Figure 2. The coefficients $\beta_{i}$ are very close to each other for the Covertype and Kddcup data. This is what we expected since the tails of the 1-generalized normal (Laplace) distribution and the logistic distribution have a very similar exponential decay. For Webspam, however, there are few single values of $\beta_{i}$ that differ more significantly for the two models. Still, most coefficients are very similar and all less similar values have the same sign, so they point in the same direction regarding the influence of single variables. The deviating cases indicate that many points are located close to the separating hyperplane, where the two distributions deviate most from each other. This impression is affirmed in Figure 2, where we see the resulting linear separators and (misclassified) residuals for a 2D data example. The logit regression line passes through a whole bunch of points, and as a result the logit model is more strongly attracted to the blue outliers, being closer to the cases $p=1.5$ or $p=2$ than to $p=1$. The logistic model thus seems to adapt to different situations, while the 1-probit model always yields a robust estimator.
(Q2) We assess the difference between $p$-probit models for different values of $p \in\{1,1.5,2,3,5\}$ on a 2 D data set (plus intercept). In Figure 2 we see that the case $p=1$ is most robust, fitting the majority of the points, and ignoring the outliers. With an increasing value of $p$, the model becomes more sensitive and attracted to the outliers. The boxplots show the residuals of the misclassified points. For small values of $p \in\{1,1.5,2\}$ the model decreasingly ignores the blue outliers, which lead to large positive residuals. There are no orange outliers, so the negative residuals are relatively small. With increasing $p \in\{3,5\}$ the model gradually tends


Figure 2: (Q1) \& (Q2): (top) Comparison of normed coefficients for the logit vs. the 1-probit model on different data sets. (bottom) A 2D example data set that demonstrates how different values of $p$ affect the linear separator, as well as the distribution of the residuals $X \beta /\|\beta\|_{2}$ of misclassifications.
to minimize the distance to the most outlying misclassifications, independent of what happens between those extremes. We finally note that we have tried larger values but the effect does not change significantly beyond $p=5$, which indicates that relatively small values are already close to the limiting case $p \rightarrow \infty$.

## II. Accuracy and Algorithmic Efficiency As-

 pects: We compare our coresets to uniform sampling and for $p=1$ additionally to the square root of $\ell_{2}$ leverage scores and to $\ell_{1}$-Lewis weights, that were developed for logistic regression (Munteanu et al., 2018; Mai et al., 2021). For $p=2$ we additionally compare to the online algorithm for online $\ell_{2}$ leverage scores. See Figure 3.(Q3) For $p=1$ the Lewis weights have the best accuracy, closely followed by the 1-probit coreset. Square root $\ell_{2}$ sampling is third on Webspam and Covertype but seems superior and more stable on Kddcup. Uniform sampling is good only on Covertype (which seems to be a very uniform data set) but otherwise uniform sampling is considerably worse than all competitors. For $p=2$ we see that the online leverage score approximation performs worse than uniform sampling on some data sets but is close to the probit model on others. This might depend on the condition of the data as indicated by the theoretical analyses of online leverage scores (Chhaya et al., 2020; Cohen et al., 2020). Uniform sampling generally detoriorates with increasing $p$ since it misses the small number of outliers which become increasingly important for obtaining a good approximation. Only for Covertype it seems to be a good sampling probability independent of $p$ due to the


Figure 3: (Q3) \& (Q4): (top) Accuracy of coreset constructions for $p=1$ on three data sets. (middle) Accuracy on the Webspam data for increasing values of $p$. (bottom left) Accuracy of coreset constructions for $p=2$ on two data sets. (bottom right) Time vs. accuracy plot for different coreset constructions for $p=1$. The median approximation ratio denotes for each sample size the median of the approximation ratios over all repetitions.
uniformity of the data set. For more complicated Webspam and especially for Kddcup, uniform sampling fails. Our coresets perform well for all values of $p$, although for Covertype we note that for $p>2$ the convergence seems to start very late. We believe this is due to the considerably higher complexity of sketching and consequently weaker estimates of the norms for $p>2$.
(Q4) We fist note that the online algorithm is extremely slow, running in $\Theta\left(n d^{2}+\operatorname{poly}(d)\right)$; its advantage is that it needs only one single row-order pass as opposed to the other two-pass algorithms. Although Lewis weights won the competition with respect to accuracy, we note that they require repeated calculations of leverage scores and are thus much slower, which is also reflected when considering the time vs. accuracy trade-off. Uniform sampling seems to be the fastest choice when we only need a weak approximation, but our coresets are the best choice to obtain both, fast and accurate estimates at the same time.

In summary, the answers to questions (Q1)-(Q4) show the benefits and limitations of our methods.

## 4 CONCLUSION

We introduce the $p$-generalized probit model as a flexible framework for modeling binary data, generalizing the standard probit model to control the tail behavior (kurtosis) of the link, and to model different levels of sensitivity to outliers. To facilitate an efficient and scalable maximum likelihood estimation for the parameters of the $p$-generalized probit model, we develop coreset constructions and combine with sketching techniques for the associated $p$-probit loss function. Hereby, we advance the analytical analysis on the tails of the $p$-generalized normal distribution, which may be of independent interest. Our algorithms run in nearly input sparsity time in two passes over the data which can be improved for $p=2$, in which case we also have a single pass online algorithm. Our experiments demonstrate the usefulness of our $p$-generalization to model the sensitivity to outliers, as well as the efficiency and accuracy of our coreset and sketching algorithms. It would be interesting to continue this work by extending modeling flexibility to allow for skewed distributions.

## Acknowledgements

We thank the anonymous reviewers for their helpful comments. We thank Prof. Dr. Katja Ickstadt for pointing us to the probit model and for valuable discussions on that topic. The authors were supported by the German Science Foundation (DFG), Collaborative Research Center SFB 876, project C4 and by the Dortmund Data Science Center (DoDSc).

## References

James H. Albert and Siddhartha Chib. Bayesian analysis of binary and polychotomous response data. Journal of the American Statistical Association, 88(422): 669-679, 1993.

Alexandr Andoni. High frequency moments via maxstability. In IEEE International Conference on Acoustics, Speech and Signal Processing, (ICASSP), pages 6364-6368, 2017.
Alexandr Andoni, Huy L. Nguyên, Yury Polyanskiy, and Yihong Wu. Tight lower bound for linear sketches of moments. In Proceedings of the 40 th International Colloquium on Automata, Languages, and Programming (ICALP), pages 25-32, 2013.
Herman Auerbach. On the area of convex curves with conjugate diameters. PhD thesis, University of Lwów, 1930.

Vladimir Braverman, Dan Feldman, and Harry Lang. New frameworks for offline and streaming coreset constructions. CoRR, abs/1612.00889, 2016.
Sébastien Bubeck. Convex optimization: Algorithms and complexity. Found. Trends Mach. Learn., 8(3-4): 231-357, 2015.
Min-Te Chao. A general purpose unequal probability sampling plan. Biometrika, 69(3):653-656, 1982.
Rachit Chhaya, Jayesh Choudhari, Anirban Dasgupta, and Supratim Shit. Streaming coresets for symmetric tensor factorization. In Proceedings of the 37th International Conference on Machine Learning (ICML), volume 119, pages 1855-1865, 2020.
Kenneth L. Clarkson. Subgradient and sampling algorithms for $\ell_{1}$ regression. In Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 257-266, 2005.
Kenneth L Clarkson and David P Woodruff. Low-rank approximation and regression in input sparsity time. Journal of the ACM, 63(6):1-45, 2017.
Kenneth L. Clarkson, Petros Drineas, Malik MagdonIsmail, Michael W. Mahoney, Xiangrui Meng, and David P. Woodruff. The fast Cauchy transform and faster robust linear regression. SIAM Journal on Computing, 45(3):763-810, 2016.

Michael B. Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 278-287, 2016.
Michael B. Cohen, Cameron Musco, and Jakub Pachocki. Online row sampling. Theory Comput., 16: 1-25, 2020.
J.S. Cramer. The origins of logistic regression. Tinbergen Institute, Tinbergen Institute Discussion Papers, TI 2002(119/4):1-15, 112002.
Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W. Mahoney. Sampling algorithms and coresets for $\ell_{p}$ regression. SIAM J. Comput., 38(5):2060-2078, 2009.
Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, and David P. Woodruff. Fast approximation of matrix coherence and statistical leverage. J. Mach. Learn. Res., 13:3475-3506, 2012.
Alex Dytso, Ronit Bustin, H. Vincent Poor, and Shlomo Shamai. Analytical properties of generalized Gaussian distributions. Journal of Statistical Distributions and Applications, 5:1-40, 122018.
Dan Feldman. Core-sets: An updated survey. Wiley Interdiscip. Rev. Data Min. Knowl. Discov., 10(1), 2020.

Dan Feldman and Michael Langberg. A unified framework for approximating and clustering data. In Proceedings of the 43 rd ACM Symposium on Theory of Computing (STOC), pages 569-578, 2011.
Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning Big Data into tiny data: Constant-size coresets for k-means, PCA, and projective clustering. SIAM J. Comput., 49(3):601-657, 2020.
Irwin R. Goodman and Samuel Kotz. Multivariate $\theta$-generalized normal distributions. Journal of Multivariate Analysis, 3(2):204-219, 1973.
Robert D. Gordon. Values of Mills' ratio of area to bounding ordinate and of the normal probability integral for large values of the argument. The Annals of Mathematical Statistics, 12(3):364-366, 1941.

Yuanyuan Gu, Denzil G Fiebig, Edward Cripps, and Robert Kohn. Bayesian estimation of a random effects heteroscedastic probit model. The Econometrics Journal, 12(2):324-339, 2009.
Jonathan H. Huggins, Trevor Campbell, and Tamara Broderick. Coresets for scalable Bayesian logistic regression. In Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS), pages 4080-4088, 2016.
Norman L. Johnson, Samuel Kotz, and Narayanaswamy Balakrishnan. Continuous univariate distributions, Volume 1. Wiley \& Sons, 2nd edition, 1994.

William B Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. Contemporary mathematics, 26(1):189-206, 1984.
S. Kalke and W.-D. Richter. Simulation of the $p$ generalized Gaussian distribution. Journal of Statistical Computation and Simulation, 83(4):641-667, 2013.

Michael J. Kearns and Umesh V. Vazirani. An Introduction to Computational Learning Theory. MIT Press, Cambridge, 1994.

Michael Langberg and Leonard J. Schulman. Universal $\varepsilon$-approximators for integrals. In Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 598-607, 2010.
Chengfeng Lei and Xiulian Sun. Comparing lethal dose ratios using probit regression with arbitrary slopes. BMC Pharmacology and Toxicology, 19(1): 1-10, 2018.

Tung Mai, Anup B. Rao, and Cameron Musco. Coresets for classification - simplified and strengthened. CoRR, abs/2106.04254, 2021.
P. McCullagh and J. A. Nelder. Generalized Linear Models. Chapman \& Hall, London, 1989.

Xiangrui Meng and Michael W. Mahoney. Lowdistortion subspace embeddings in input-sparsity time and applications to robust linear regression. In Proceedings of the Twenty-First Annual ACMSIAM Symposium on Discrete Algorithms (SODA), pages 91-100, 2013.

Alejandro Molina, Alexander Munteanu, and Kristian Kersting. Core dependency networks. In Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), pages 3820-3827, 2018.

Richard Kouamé Moussa. Heteroskedasticity in oneway error component probit models. Econometrics, $7(3): 35,2019$.

Alexander Munteanu and Chris Schwiegelshohn. Coresets-methods and history: A theoreticians design pattern for approximation and streaming algorithms. Künstliche Intell., 32(1):37-53, 2018.

Alexander Munteanu, Chris Schwiegelshohn, Christian Sohler, and David P. Woodruff. On coresets for logistic regression. In Advances in Neural Information Processing Systems 31, (NeurIPS), pages 6562-6571, 2018.

Alexander Munteanu, Simon Omlor, and David P. Woodruff. Oblivious sketching for logistic regression. In Proceedings of the 38th International Conference on Machine Learning (ICML), pages 7861-7871, 2021.

Jacek Osiewalski and Mark F. J. Steel. Robust Bayesian inference in $\ell_{q}$-spherical models. Biometrika, 80(2): 456-460, 1993.
Jeff M Phillips. Coresets and sketches. In Handbook of Discrete and Computational Geometry, pages 12691288. Chapman and Hall/CRC, 3rd edition, 2017.

Sashank J. Reddi, Barnabás Póczos, and Alexander J. Smola. Communication efficient coresets for empirical loss minimization. In Proceedings of the ThirtyFirst Conference on Uncertainty in Artificial Intelligence (UAI), pages 752-761, 2015.

Alireza Samadian, Kirk Pruhs, Benjamin Moseley, Sungjin Im, and Ryan R. Curtin. Unconditional coresets for regularized loss minimization. In The 23rd International Conference on Artificial Intelligence and Statistics, (AISTATS), pages 482-492, 2020.

Fabian Sinz, Sebastian Gerwinn, and Matthias Bethge. Characterization of the $p$-generalized normal distribution. Journal of Multivariate Analysis, 100(5): 817-820, 2009.
Christian Sohler and David P. Woodruff. Subspace embeddings for the $\ell_{1}$-norm with applications. In Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC), pages 755-764, 2011.
M. T. Subbotin. On the law of frequency of error. Matematicheskǐ Sbornik, 31(2):296-301, 1923.

Elad Tolochinsky and Dan Feldman. Coresets for monotonic functions with applications to deep learning. CoRR, abs/1802.07382, 2018.
Murad Tukan, Alaa Maalouf, and Dan Feldman. Coresets for near-convex functions. In Advances in Neural Information Processing Systems 33, (NeurIPS), 2020.

Cristiano Varin and Claudia Czado. A mixed autoregressive probit model for ordinal longitudinal data. Biostatistics, 11(1):127-138, 112009.

David P. Woodruff. Sketching as a tool for numerical linear algebra. Found. Trends Theor. Comput. Sci., 10(1-2):1-157, 2014.
David P. Woodruff and Qin Zhang. Subspace embeddings and $\ell_{p}$-regression using exponential random variables. In Proceedings of the 26th Annual Conference on Learning Theory (COLT), pages 546-567, 2013.

## A GRADIENT AND HESSIAN MATRIX

The optimization of $f$ can be done by gradient descent or by applying the Newton-Raphson method (Bubeck, 2015), an iterative procedure that starts at an initial guess $\beta^{(0)}$ and successively applies the following update rule:

$$
\beta^{(t)}=\beta^{(t-1)}-\left(\frac{\partial^{2} f\left(\beta^{(t-1)}\right)}{\partial \beta \partial \beta^{T}}\right)^{-1} \cdot \frac{\partial f\left(\beta^{(t-1)}\right)}{\partial \beta}
$$

where $\left(\frac{\partial^{2} f\left(\beta^{(t-1)}\right)}{\partial \beta \partial \beta^{T}}\right)^{-1}$ refers to the inverse of the Hessian matrix of $f$, evaluated at $\beta^{(t-1)}$, and $\frac{\partial f\left(\beta^{(t-1)}\right)}{\partial \beta}$ refers to the gradient of $f$, evaluated at $\beta^{(t-1)}$. The idea behind this procedure is, broadly speaking, to approximate $f$ locally around $\beta^{(t)}$ by its second degree Taylor-polynomial and then analytically find the minimum of this polynomial. The minimum of this local polynomial approximation of $f$ is then used iteratively as a basis for the next step of the Newton-Raphson algorithm.
It remains to derive the gradient and the Hessian matrix of $f$. Since $f$ is a sum of the function $g$ evaluated at different points, it makes sense to first determine the derivative of $g$. To this end $\varphi_{p}(r)$ is the density function of the (standardized) p-generalized normal distribution function, cf. (Dytso et al., 2018; Kalke and Richter, 2013):

$$
\varphi_{p}(r)=\frac{p^{1-1 / p}}{2 \Gamma(1 / p)} \exp \left(-|r|^{p} / p\right)
$$

We proceed by using the chain rule as follows:

$$
\begin{aligned}
\frac{d}{d r} g(r)=\frac{d}{d r}-\ln \left(\Phi_{p}(-r)\right) & =\frac{d}{d r} \ln \left(\frac{1}{1-\Phi_{p}(r)}\right) \\
& =\left(1-\Phi_{p}(r)\right) \cdot \frac{d}{d r}\left(\frac{1}{1-\Phi_{p}(r)}\right) \\
& =\left(1-\Phi_{p}(r)\right) \cdot \frac{(-1)}{\left(1-\Phi_{p}(r)\right)^{2}} \cdot \frac{d}{d r}\left(1-\Phi_{p}(r)\right) \\
& =\frac{(-1)}{1-\Phi_{p}(r)} \cdot(-1) \cdot \varphi_{p}(r) \\
& =\frac{\varphi_{p}(r)}{1-\Phi_{p}(r)}
\end{aligned}
$$

We can use this result to calculate the gradient of $f$ :

$$
\begin{aligned}
\frac{\partial}{\partial \beta} f_{w}(X \beta) & =\frac{\partial}{\partial \beta} \sum_{i=1}^{n} w_{i} g\left(x_{i} \beta\right) \\
& =\sum_{i=1}^{n} w_{i} x_{i} g^{\prime}\left(x_{i} \beta\right) \\
& =\sum_{i=1}^{n} w_{i} x_{i} \frac{\varphi_{p}\left(x_{i} \beta\right)}{1-\Phi_{p}\left(x_{i} \beta\right)}
\end{aligned}
$$

Next, we need to determine the Hessian matrix of $f$. To this end, we again start by finding the second derivative of $g$, this time using the quotient rule.

$$
\begin{aligned}
\frac{d^{2}}{d r^{2}} g(r) & =\frac{d}{d r} \frac{\varphi_{p}(r)}{1-\Phi_{p}(r)} \\
& =\frac{\varphi_{p}^{\prime}(r)\left(1-\Phi_{p}(r)\right)-\varphi_{p}(r) \cdot(-1) \cdot \varphi_{p}(r)}{\left(1-\Phi_{p}(r)\right)^{2}} \\
& =\frac{(-1) \cdot \operatorname{sgn}(r) \cdot|r|^{p-1} \cdot \varphi_{p}(r)\left(1-\Phi_{p}(r)\right)-\varphi_{p}(r) \cdot(-1) \cdot \varphi_{p}(r)}{\left(1-\Phi_{p}(r)\right)^{2}}
\end{aligned}
$$

$$
\begin{aligned}
& =\frac{\left[\varphi_{p}(r)\right]^{2}-\left.\operatorname{sgn}(r)|r|\right|^{p-1} \cdot \varphi_{p}(r) \cdot\left(1-\Phi_{p}(r)\right)}{\left(1-\Phi_{p}(r)\right)^{2}} \\
& =\left(\frac{\varphi_{p}(r)}{1-\Phi_{p}(r)}\right)^{2}-\operatorname{sgn}(r)|r|^{p-1} \cdot \frac{\varphi_{p}(r)}{1-\Phi_{p}(r)} \\
& =\frac{\varphi_{p}(r)}{1-\Phi_{p}(r)}\left(\frac{\varphi_{p}(r)}{1-\Phi_{p}(r)}-\operatorname{sgn}(r)|r|^{p-1}\right) \\
& =g^{\prime}(r) \cdot\left(g^{\prime}(r)-\operatorname{sgn}(r)|r|^{p-1}\right)
\end{aligned}
$$

We can now use this result to find the Hessian matrix of $f$. Recall that $x_{i}$ are row vectors in our paper and thus each $x_{i}^{T} x_{i}$ is a $d \times d$-matrix.

$$
\begin{aligned}
\frac{\partial^{2}}{\partial \beta \partial \beta^{T}} f_{w}(X \beta) & =\sum_{i=1}^{n} \frac{\partial^{2}}{\partial \beta \partial \beta^{T}} w_{i} g\left(x_{i} \beta\right) \\
& =\sum_{i=1}^{n} w_{i} x_{i}^{T} x_{i} g^{\prime}\left(x_{i} \beta\right)\left(g^{\prime}\left(x_{i} \beta\right)-\operatorname{sgn}\left(x_{i} \beta\right)\left|x_{i} \beta\right|^{p-1}\right) \\
& =\sum_{i=1}^{n} w_{i} x_{i}^{T} x_{i} \frac{\varphi_{p}\left(x_{i} \beta\right)}{1-\Phi_{p}\left(x_{i} \beta\right)}\left(\frac{\varphi_{p}\left(x_{i} \beta\right)}{1-\Phi_{p}\left(x_{i} \beta\right)}-\operatorname{sgn}\left(x_{i} \beta\right)\left|x_{i} \beta\right|^{p-1}\right) .
\end{aligned}
$$

It can be shown, that $f_{w}(X \beta)$ is a convex function of $\beta$, and that the Newton-Raphson algorithm converges to a global optimum when applied to a convex function (Bubeck, 2015). The optimization procedure thus converges to the maximum likelihood estimate $\hat{\beta} \in \operatorname{argmin}_{\beta \in \mathbb{R}^{d}} f_{w}(X \beta)$ provided it exists.

## B SENSITIVITY FRAMEWORK

Our approach is based on the so called sensitivity framework (Langberg and Schulman, 2010). The sensitivity of a point is its worst case contribution to the entire loss function $\zeta_{i}=\sup _{\beta \in \mathbb{R}^{d}} \frac{g\left(x_{i} \beta\right)}{f(X \beta)}$, cf. Definition 2.5. Computing the sensitivity of a point can be difficult. Fortunately it suffices to get a reasonably tight upper bound on the sensitivity of a point. In Section C. 2 we will show that such a bound can be derived in our setting.
Since sensitivities are not sufficient to get a good bound for all solutions $\beta$, we also need the terminology of the VC-dimension:
Definition B.1. The range space for a set $\mathcal{F}$ is a pair $\mathfrak{R}=(\mathcal{F}$, ranges) where ranges is a family of subsets of $\mathcal{F}$. The $V C$-dimension $\Delta(\mathfrak{R})$ of $\mathfrak{R}$ is the size $|G|$ of the largest subset $G \subseteq \mathcal{F}$ such that $G$ is shattered by ranges, i.e., $\mid\{G \cap R \mid R \in$ ranges $\} \mid=2^{|G|}$.
Definition B.2. Let $\mathcal{F}$ be a finite set of functions mapping from $\mathbb{R}^{d}$ to $\mathbb{R}_{\geq 0}$. For every $x \in \mathbb{R}^{d}$ and $r \in \mathbb{R}_{\geq 0}$, let $\operatorname{range}_{\mathcal{F}}(x, r)=\{f \in \mathcal{F} \mid f(x) \geq r\}$, and $\operatorname{ranges}(\mathcal{F})=\left\{\operatorname{range}_{\mathcal{F}}(x, r) \mid x \in \mathbb{R}^{\boldsymbol{d}}, r \in \mathbb{R}_{\geq 0}\right\}$, and $\mathfrak{R}_{\mathcal{F}}=(\mathcal{F}, \operatorname{ranges}(\mathcal{F}))$ be the range space induced by $\mathcal{F}$.

The VC-dimension can be thought of something similar to the dimension of our problem. For example the VC-dimension of the set of hyperplane classifiers in $\mathbb{R}^{d}$ is $d+1$ (Kearns and Vazirani, 1994). We analyze the VC-dimension of our problem in Section C.1. The sensitivity scores were combined with a theory on the VC-dimension of range spaces in (Feldman and Langberg, 2011; Braverman et al., 2016). We use a more recent version of Feldman et al. (2020).
Proposition B.3. (Feldman et al., 2020) Consider a family of functions $\mathcal{F}=\left\{f_{1}, \ldots, f_{n}\right\}$ mapping from $\mathbb{R}^{d}$ to $[0, \infty)$ and a vector of weights $u \in \mathbb{R}_{>0}^{n}$. Let $\varepsilon, \delta \in(0,1 / 2)$. Let $s_{i} \geq \zeta_{i}$. Let $S=\sum_{i=1}^{n} s_{i} \geq \sum_{i=1}^{n} \zeta_{i}=Z$. Given $s_{i}$ one can compute in time $O(|\mathcal{F}|)$ a set $R \subset \mathcal{F}$ of

$$
O\left(\frac{S}{\varepsilon^{2}}\left(\Delta \ln S+\ln \left(\frac{1}{\delta}\right)\right)\right)
$$

weighted functions such that with probability $1-\delta$, we have for all $x \in \mathbb{R}^{d}$ simultaneously

$$
\left|\sum_{f_{i} \in \mathcal{F}} u_{i} f_{i}(x)-\sum_{f_{i} \in R} w_{i} f_{i}(x)\right| \leq \varepsilon \sum_{f_{i} \in \mathcal{F}} u_{i} f_{i}(x),
$$

where each element of $R$ is sampled i.i.d. with probability $p_{j}=\frac{s_{j}}{S}$ from $\mathcal{F}, w_{i}=\frac{S u_{j}}{s_{j}|R|}$ denotes the weight of a function $f_{i} \in R$ that corresponds to $f_{j} \in \mathcal{F}$, and where $\Delta$ is an upper bound on the $V C$-dimension of the range space $\Re_{\mathcal{F}^{*}}$ induced by $\mathcal{F}^{*}$ obtained by defining $\mathcal{F}^{*}$ to be the set of functions $f_{j} \in \mathcal{F}$, where each function is scaled by $\frac{S u_{j}}{s_{j}|R|}$.

In the following section we will show how to compute upper bounds for the sensitivities and VC-dimension for our loss function. In order to derive those bounds we first need to analyze the loss function $g$ of an individual point. This will enable us to compute the desired upper bounds for both, the VC-dimension in Section C. 1 and the sensitivities in Section C.2. Finally we will be able to apply Proposition B. 3 to obtain a coreset, and to prove Theorem 1.

## C PROPERTIES OF $g$

In this section we will determine useful properties of $g$. Recall that for any function $f: \mathbb{R} \rightarrow \mathbb{R}$ the derivative $\frac{d}{d r} \int_{r}^{\infty} f(t) d t$ equals $\lim _{t \rightarrow \infty} f(t)-f(r)$ if the integral is finite. First we note that for $r \in \mathbb{R}$ with $r \geq 0$

$$
\begin{aligned}
g^{\prime}(r) & =\frac{\varphi_{p}(r)}{1-\Phi(r)}=\frac{\exp \left(-|r|^{p} / p\right)}{\int_{r}^{\infty} \exp \left(-|t|^{p} / p\right) d t} \\
& =\frac{1}{\exp \left(|r|^{p} / p\right) \int_{r}^{\infty} \exp \left(-|t|^{p} / p\right) d t}>0
\end{aligned}
$$

We set $h(r):=\frac{1}{g^{\prime}(r)}=\exp \left(|r|^{p} / p\right) \int_{r}^{\infty} \exp \left(-|t|^{p} / p\right) d t$. Our aim is to characterize the tail behavior of the $p$-generalized normal distribution. To this end we will first analyze $h$ using similar methods as Gordon (1941) who considered the case $p=2$, i.e., the standard normal distribution.
Lemma C.1. The following holds for any $r>0$ :

$$
\begin{align*}
h^{\prime}(r) & =r^{p-1} h(r)-1  \tag{2}\\
h^{\prime \prime}(r) & =(p-1) r^{p-2} h(r)+r^{p-1} h^{\prime}(r)  \tag{3}\\
h^{\prime \prime}(r) & =\frac{r^{p}+p-1}{r} h^{\prime}(r)+\frac{p-1}{r}  \tag{4}\\
h^{\prime \prime \prime}(r) & =\left(1+\frac{p}{r^{p}+p-1}+\frac{p-2}{r^{p}}\right) r^{p-1} h^{\prime \prime}(r)-\frac{(p-1) p r^{p-2}}{r^{p}+p-1}  \tag{5}\\
h(r) & >0  \tag{6}\\
h^{\prime}(r) & <0  \tag{7}\\
h(r) & <\frac{1}{r^{p-1}} \tag{8}
\end{align*}
$$

Further if $r \geq 1$ then it holds that

$$
\begin{align*}
h^{\prime \prime}(r) & \geq 0  \tag{9}\\
h(r) & \geq \frac{r}{r^{p}+p-1} \tag{10}
\end{align*}
$$

Proof of Lemma C.1. Equations (2) and (3) can be derived by a direct calculation of the derivatives. Note that (2) is equivalent to

$$
\begin{equation*}
h(r)=\frac{h^{\prime}(r)+1}{r^{p-1}} \tag{11}
\end{equation*}
$$

Equation (4) follows by substitution of (11) in (3). Equation (3) is equivalent to

$$
\begin{equation*}
h(r)=\frac{h^{\prime \prime}(r)}{(p-1) r^{p-2}}-\frac{r}{p-1} h^{\prime}(r) \tag{12}
\end{equation*}
$$

To get (5) we first note that by (2) and then (12) it holds

$$
\begin{align*}
\frac{p-1}{r^{2}} h^{\prime}(r) & =\frac{p-1}{r^{2}}\left(r^{p-1} h(r)-1\right) \\
& =\frac{p-1}{r^{2}} r^{p-1} h(r)-\frac{p-1}{r^{2}} \\
& =\frac{h^{\prime \prime}(r)}{r}-r^{p-2} h^{\prime}(r)-\frac{p-1}{r^{2}} . \tag{13}
\end{align*}
$$

Further note that (4) is equivalent to

$$
\begin{equation*}
h^{\prime}(r)=\frac{r h^{\prime \prime}(r)}{r^{p}+p-1}-\frac{p-1}{r^{p}+p-1} \tag{14}
\end{equation*}
$$

Taking the derivative of (4) and using the equations (13) and (14) we get

$$
\begin{aligned}
& h^{\prime \prime \prime}(r)=(p-1) r^{p-2} h^{\prime}(r)-\frac{p-1}{r^{2}} h^{\prime}(r)+r^{p-1} h^{\prime \prime}(r) \\
&+\frac{p-1}{r} h^{\prime \prime}(r)-\frac{p-1}{r^{2}} \\
& \stackrel{(13)}{=} p r^{p-2} h^{\prime}(r)+r^{p-1} h^{\prime \prime}(r)+\frac{p-2}{r} h^{\prime \prime}(r) \\
& \stackrel{(14)}{=} p r^{p-2} \cdot\left(\frac{r}{r^{p}+p-1} h^{\prime \prime}(r)-\frac{p-1}{r^{p}+p-1}\right) \\
&+r^{p-1} h^{\prime \prime}(r)+\frac{p-2}{r} h^{\prime \prime}(r) \\
&=\left(1+\frac{p}{r^{p}+p-1}+\frac{p-2}{r^{p}}\right) r^{p-1} h^{\prime \prime}(r) \\
&-\frac{(p-1) p r^{p-2}}{r^{p}+p-1} .
\end{aligned}
$$

Equation (6) follows since all terms appearing in $h(r)$ are positive.
For (7) we note that

$$
\begin{align*}
r^{p-1} h(r) & =\exp \left(r^{p} / p\right) \int_{r}^{\infty} r^{p-1} \exp \left(-|t|^{p} / p\right) d t \\
& <\exp \left(r^{p} / p\right) \int_{r}^{\infty} \frac{p}{p} t^{p-1} \exp \left(-|t|^{p} / p\right) d t \\
& =\exp \left(r^{p} / p\right) \cdot \exp \left(-r^{p} / p\right)=1 \tag{15}
\end{align*}
$$

and thus (8) follows from dividing by $r^{p-1}$ and (7) also follows from (15) using Equation (2).
Next we prove (9): For $r \geq 1$ it holds that $\left(1+\frac{p}{r^{p}+p-1}+\frac{p-2}{r^{p}}\right) r^{p-1}>0$. Now let $r_{0} \geq 1$. Assume for the sake of contradiction that $h^{\prime \prime}\left(r_{0}\right)<0$. Then using (5) we also get $h^{\prime \prime \prime}\left(r_{0}\right) \leq\left(1+\frac{p}{r^{p}+p-1}+\frac{p-2}{r^{p}}\right) r^{p-1} h^{\prime \prime}\left(r_{0}\right)<0$. Thus we have $h^{\prime \prime}(r)<h^{\prime \prime}\left(r_{0}\right)$ for all $r>r_{0}$. Consequently $h^{\prime}$ is also strictly decreasing by a rate of at least $h^{\prime \prime}\left(r_{0}\right)$ starting at $r_{0}$. This implies that there exists $r^{\prime}>r_{0}$ with $h\left(r^{\prime}\right)<0$, which contradicts (6) and thus (9) follows. Lastly (10) follows by substitution of (2) in (3) and using (9).

Lemma C.2. [Copy of Lemma 2.6] The function $g$ is convex and strictly increasing. Further for any $r \geq 0$ we have

$$
g^{\prime}(r) \geq r^{p-1}
$$

for any $r \geq 1$ we have

$$
g^{\prime}(r) \leq r^{p-1}+\frac{p-1}{r}
$$

and there exists a constant $c_{1}>0$ such that

$$
g(r) \geq c_{1} e^{-|r|^{p} / p}
$$

for any $r<0$.
Proof of Lemma 2.6/C.2. First note that $g=-\ln \left(\Phi_{p}(-r)\right)$ is strictly increasing since $\Phi(-r) \in(0,1)$ is strictly decreasing for increasing $r$ and $-\ln (t)$ is strictly increasing for decreasing $t$. Next consider $r \geq 0$. Then $g^{\prime \prime}(r)=\left(\frac{1}{h(r)}\right)^{\prime}=-\frac{h^{\prime}(r)}{h(r)^{2}}>0$ by (7) thus $g$ is convex on $[0, \infty)$. For $r<0$ we have derived in Section A that

$$
g^{\prime \prime}(r)=g^{\prime}(r)\left(g^{\prime}(r)-\operatorname{sgn}(r)|r|^{p-1}\right)
$$

For $r<0$ all terms are positive. Thus $g(r)$ is convex for all $r \in \mathbb{R}$. The bounds for $g^{\prime}$ follow immediately by the bounds for $h$ from Lemma C. 1 (8) and (10).
Now, for $r<-1$ using the Taylor series of $-\ln (t)$ at $t=1$, the normalizing constant $C_{p}=\int_{-\infty}^{\infty} \exp \left(-|t|^{p} / p\right) d t=$ $\frac{2 \Gamma(1 / p)}{p^{1-1 / p}}$ and Equation (10) we have

$$
\begin{aligned}
g(r) & =-\ln \left(1-C_{p}^{-1} \int_{-r}^{\infty} \exp \left(-|t|^{p} / p\right) d t\right) \\
& \geq C_{p}^{-1} \int_{-r}^{\infty} \exp \left(-|t|^{p} / p\right) d t \\
& \geq C_{p}^{-1} \exp \left(-(-r)^{p} / p\right) \cdot \frac{r}{r^{p}+p-1} \geq \frac{\exp \left(-(-r)^{p} / p\right)}{p C_{p}}
\end{aligned}
$$

For any $r \in[-1,0]$ we have $g(r) \geq g(-1) \geq g(-1) \exp \left(-(-r)^{p} / p\right)$. Thus for $c_{1}=\min \left\{g(-1), 1 /\left(p C_{p}\right)\right\}$ we have $g(r) \geq c_{1} \exp \left(-(-r)^{p} / p\right)$.

These properties can be used to prove the following lemma:
Lemma C.3. Set $G^{+}(r)=\frac{r^{p}}{p}$ if $r \geq 0$ and $G^{+}(r)=0$ if $r<0$. There exists $c_{2}>0$ depending only on $p$ such that for any $\varepsilon \in\left(0, e^{-1}\right)$ and any $r \in \mathbb{R}$ it holds that

$$
\begin{equation*}
G^{+}(r) \leq g(r) \leq(1+\varepsilon) G^{+}(r)+c_{2} \ln \left(\frac{p}{\varepsilon}\right) \tag{16}
\end{equation*}
$$

Proof of Lemma C.3. For $r<0$ we have $g(r)>0=G^{+}(r)$. For $r \geq 0$ by using Lemma 2.6 we get

$$
\begin{aligned}
g(r) \geq g(0)+\int_{0}^{r} g^{\prime}(t) d t & \geq g(0)+\int_{0}^{r} t^{p-1} d t \\
& =g(0)+G^{+}(r) \geq G^{+}(r)
\end{aligned}
$$

For the second inequality we split the domain of $g$ into three parts: First since $g$ is monotonically increasing for any $r$, we have $g(r) \leq g(1)$ for $r \in(-\infty, 1]$. For $r \geq 1$, by using Lemma 2.6, it holds that

$$
\begin{align*}
g(r) & \leq g(1)+\int_{1}^{r} t^{p-1}+\frac{p-1}{t} d t \\
& =g(1)+G^{+}(r)-\frac{1}{p}+(p-1) \ln (r) . \tag{17}
\end{align*}
$$

Now consider $r \in\left(1, r_{0}\right]$ where $r_{0}=\frac{p^{3}}{\varepsilon^{3}}$. Then we have

$$
g(r) \leq g(1)+G^{+}(r)+(p-1) \ln \left(r_{0}\right)=g(1)+G^{+}(r)+3(p-1) \ln \left(\frac{p}{\varepsilon}\right)
$$

Our last step is to show that for $r>r_{0}$ it holds that $(p-1) \ln (r) \leq \varepsilon G^{+}(r)$. We assume without loss of generality that $\varepsilon^{-1} \geq 2$. Now the equation

$$
\varepsilon G^{+}(r)=\varepsilon \frac{r^{p}}{p} \geq(p-1) \ln (r)
$$

is equivalent to

$$
\exp \left(\frac{\varepsilon r^{p}}{p^{2}-p}\right) \geq r
$$

Note that $r^{p} \geq r$ holds since $r \geq r_{0}>1$ and thus we get for any $r=a r_{0}$ with $a \geq 1$ that

$$
\exp \left(\frac{\varepsilon r^{p}}{p^{2}-p}\right) \geq \exp \left(\frac{\varepsilon r}{p^{2}}\right) \geq \exp \left(\frac{\varepsilon a r_{0}}{p^{2}}\right) \geq \exp \left(\frac{a p}{\varepsilon^{2}}\right) \geq \exp \left(2 a \cdot \frac{p}{\varepsilon}\right) \geq a r_{0}=r
$$

The last inequality follows from the fact that $e^{2 a z} \geq a z^{3}$ always holds in our case where $z \geq 2$ and $a \geq 1$. Consequently it holds for any $r \in\left[r_{0}, \infty\right)$ that

$$
g(r) \leq g(1)+G^{+}(r)+(p-1) \ln (r) \leq g(1)+(1+\varepsilon) G^{+}(r)
$$

Combining all three inequalities we note that for any $r \in \mathbb{R}$ it holds that

$$
\begin{aligned}
g(r) & \leq g(1)+(1+\varepsilon) G^{+}(r)+(p-1) \ln \left(\frac{p^{3}}{\varepsilon^{3}}\right) \\
& =(1+\varepsilon) G^{+}(r)+\left(\frac{g(1)}{\ln (p / \varepsilon)}+3(p-1)\right) \ln \left(\frac{p}{\varepsilon}\right) \\
& \leq(1+\varepsilon) G^{+}(r)+c_{2} \ln \left(\frac{p}{\varepsilon}\right)
\end{aligned}
$$

where $c_{2}:=(g(1)+3(p-1)) \geq\left(\frac{g(1)}{\ln (p / \varepsilon)}+3(p-1)\right)$ holds, since $\varepsilon^{-1} \geq e$ and $p \geq 1$.
Lemma C.4. [Copy of Lemma 2.8] Assume $X \in \mathbb{R}^{n \times d}$ is $\mu$-complex. Then we have for any $\beta \in \mathbb{R}^{d}$ that

$$
f(X \beta)=\Omega\left(\frac{n}{\mu}(1+\ln (\mu))\right)
$$

Proof of Lemma 2.8/C.4. Let $z=X \beta$. For $r \leq 0$ we have $g(r) \geq c_{1} e^{-r^{p} / p}$ by Lemma 2.6. For $r \geq 0$ we have $g(r)=g(0)+\int_{0}^{r} g^{\prime}(t) d t$. Recall that

$$
g^{\prime}(t)=\frac{1}{h(t)} \geq t^{p-1}
$$

and thus

$$
\begin{equation*}
g(r) \geq g(0)+\int_{0}^{r} t^{p-1} d t=g(0)+\frac{r^{p}}{p} \tag{18}
\end{equation*}
$$

Set $z_{-}=\frac{1}{n} \sum_{z_{i} \leq 0}\left|z_{i}\right|^{p}$ and $z_{+}=\frac{1}{n} \sum_{z_{i} \geq 0}\left|z_{i}\right|^{p} \geq \frac{z_{-}}{\mu}$. We set $z^{-} \in \mathbb{R}^{n}$ to be the vector with $z_{i}^{-}=z_{i}$ if $z_{i}<0$ and $z_{i}^{-}=0$ else. Using convexity of $e^{-r}$ we can apply Jensens inequality to conclude that

$$
\begin{aligned}
f(X \beta) & =\sum_{i=1}^{n} g\left(z_{i}\right) \\
& =\sum_{i=1}^{n} \min \left\{g\left(z_{i}\right), g(0)\right\}+\sum_{z_{i} \geq 0} \int_{0}^{z_{i}} t^{p-1} d t \\
& \geq \sum_{i=1}^{n} c e^{-\left|z_{i}^{-}\right|^{p} / p}+\frac{1}{p} \sum_{z_{i} \geq 0} z_{i}^{p} \\
& \geq n c_{1} e^{-\left(z_{-}\right) / p}+\frac{n z_{+}}{p} \\
& \geq n c_{1} e^{-\left(z_{-}\right) / p}+\frac{n z_{-}}{\mu p} .
\end{aligned}
$$

Taking the derivative of $\ell(r)=n c_{1} e^{-(r) / p}+\frac{n r}{\mu p}$, i.e. $\ell^{\prime}(r)=\frac{n}{p}\left(-c_{1} e^{-(r) / p}+\frac{1}{\mu}\right)$ which is 0 if $\frac{r}{p}=\ln \left(c_{1} \mu\right)$. Thus it holds that

$$
f(X \beta) \geq \ell\left(z_{-}\right) \geq \frac{n}{\mu}\left(1+\ln \left(c_{1} \mu\right)\right)
$$

which is exactly what we needed to show.

## C. 1 Bounding the VC-Dimension

In order to bound the VC-dimension of the range space induced by the weighted set of functions we need to reduce the number of distinct weights considered. We first round all sensitivities to their closest power of 2 . The new total sensitivity $S^{\prime}$ is at most twice the old sensitivity $S$. Next we increase all sensitivities smaller than $\frac{S}{n}$ to $\frac{S}{n}$. The new sensitivity is at most $S^{\prime}+n \cdot S / n=3 S$. The next step is to split the data into high sensitivity points and low sensitivity points.
Lemma C.5. [Copy of Lemma 2.7] Let $I_{1}$ be the index set of all data points with $s_{i}>s_{0}:=\frac{\mu S c \ln \left(p \varepsilon^{-1}\right)}{\varepsilon n}$ for some constant $c \in \mathbb{R}_{>0}$. Then for all $\beta \in \mathbb{R}^{d}$ it holds that

$$
\begin{aligned}
\sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right) & \leq \sum_{i \in I_{1}} g\left(x_{i} \beta\right) \\
& \leq(1+\varepsilon) \sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right)+\varepsilon \cdot \frac{n}{\mu}
\end{aligned}
$$

Proof of Lemma 2.7/C.5. We set $c=c_{2}$ as in Lemma C.3. Note that there are at most $\frac{S}{s_{0}}=\frac{\varepsilon n}{c \ln \left(p \varepsilon^{-1}\right) \mu}$ points in $I_{1}$. Thus the lemma follows by applying Lemma C. 3 to each point in $I_{1}$.

As a consequence we get the following corollary:
Corollary C.6. Let $I_{2}=[n] \backslash I_{1}$. Further let $\left(X^{\prime}, w\right) \in \mathbb{R}^{n^{\prime} \times d} \times \mathbb{R}^{n^{\prime}}$ with rows $x_{i}^{\prime}=x_{\pi(i)}$ for some mapping $\pi_{\tilde{\sim}}:\left[n^{\prime}\right] \rightarrow[n]$. We set $I_{1}^{\prime}=\left\{i \in\left[n^{\prime}\right] \mid \pi(i) \in I_{1}\right\}$ and similarly $I_{2}^{\prime}=\left\{i \in\left[n^{\prime}\right] \mid \pi(i) \in I_{2}\right\}$. Further define $\tilde{f}_{w}\left(X^{\prime} \beta\right)=\sum_{i \in I_{2}^{\prime}} w_{i} g\left(x_{i}^{\prime} \beta\right)+\sum_{i \in I_{1}^{\prime}} w_{i} G^{+}\left(x_{i}^{\prime} \beta\right)$ and by $\tilde{f}(X \beta)=\sum_{i \in I_{2}} g\left(x_{i} \beta\right)+\sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right)$. Assume that for all $\beta \in \mathbb{R}^{d}$ it holds

$$
\begin{equation*}
\left|\tilde{f}_{w}\left(X^{\prime} \beta\right)-\tilde{f}(X \beta)\right| \leq \varepsilon \tilde{f}(X \beta) \tag{19}
\end{equation*}
$$

and $\sum_{i \in I_{1}^{\prime}} w_{i} \leq \frac{2 S}{s_{0}}$. Further assume that $\varepsilon \leq \frac{1}{4}$. Then $\left(X^{\prime}, w\right)$ is a $7 \varepsilon$-coreset for the original $f$.
Proof of Corollary C.6. Observe that by triangle inequality

$$
\begin{equation*}
\left|f_{w}\left(X^{\prime} \beta\right)-f(X \beta)\right| \leq\left|f_{w}\left(X^{\prime} \beta\right)-\tilde{f}_{w}\left(X^{\prime} \beta\right)\right|+\left|\tilde{f}_{w}\left(X^{\prime} \beta\right)-\tilde{f}(X \beta)\right|+|\tilde{f}(X \beta)-f(X \beta)| \tag{20}
\end{equation*}
$$

By Lemma 2.7 it holds that

$$
\begin{aligned}
\tilde{f}(X \beta) \leq f(X \beta) & \leq \tilde{f}(X \beta)+\varepsilon \sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right)+\varepsilon \cdot \frac{n}{\mu} \\
& \leq \tilde{f}(X \beta)+2 \varepsilon f(X \beta)
\end{aligned}
$$

We thus have that

$$
|\tilde{f}(X \beta)-f(X \beta)| \leq 2 \varepsilon f(X \beta)
$$

Analogously to Lemma 2.7, using the bounded size of $\sum_{i \in I_{1}^{\prime}} w_{i}$ and the assumption (19) one can show that

$$
\begin{aligned}
\tilde{f}_{w}\left(X^{\prime} \beta\right) \leq f_{w}\left(X^{\prime} \beta\right) & \leq \tilde{f}_{w}\left(X^{\prime} \beta\right)+\varepsilon \sum_{i \in I_{1}^{\prime}} w_{i} G^{+}\left(x_{i}^{\prime} \beta\right)+\sum_{i \in I_{1}^{\prime}} w_{i} \cdot c_{2} \ln \left(\frac{p}{\varepsilon}\right) \\
& \leq \tilde{f}_{w}\left(X^{\prime} \beta\right)+\varepsilon \tilde{f}_{w}\left(X^{\prime} \beta\right)+\frac{2 S}{s_{0}} \cdot c_{2} \ln \left(\frac{p}{\varepsilon}\right) \\
& \stackrel{(19)}{\leq} \tilde{f}_{w}\left(X^{\prime} \beta\right)+\varepsilon(1+\varepsilon) \tilde{f}(X \beta)+2 \varepsilon \frac{n}{\mu} \\
& \leq \tilde{f}_{w}\left(X^{\prime} \beta\right)+2 \varepsilon \tilde{f}(X \beta)+2 \varepsilon f(X \beta) \\
& \leq \tilde{f}_{w}\left(X^{\prime} \beta\right)+4 \varepsilon f(X \beta)
\end{aligned}
$$

and thus we have

$$
\left|f_{w}\left(X^{\prime} \beta\right)-\tilde{f}_{w}\left(X^{\prime} \beta\right)\right| \leq 4 \varepsilon f(X \beta) .
$$

Now combining everything into Equation (20) yields

$$
\begin{aligned}
\left|f_{w}\left(X^{\prime} \beta\right)-f(X \beta)\right| & \leq\left|f_{w}\left(X^{\prime} \beta\right)-\tilde{f}_{w}\left(X^{\prime} \beta\right)\right|+\left|\tilde{f}_{w}\left(X^{\prime} \beta\right)-\tilde{f}(X \beta)\right|+|\tilde{f}(X \beta)-f(X \beta)| \\
& \leq 4 \varepsilon f(X \beta)+\varepsilon \tilde{f}(X \beta)+2 \varepsilon f(X \beta) \\
& \leq 7 \varepsilon f(X \beta)
\end{aligned}
$$

and thus $\left(X^{\prime}, w\right)$ is a $7 \varepsilon$-coreset.
Before we continue showing that for the set of functions that we consider, the VC-dimension is not too large, we show that the assumption made in Corollary C. 6 that $\sum_{i \in I_{1}^{\prime}} w_{i} \leq \frac{2 S}{s_{0}}$ is reasonable, i.e., that it holds with high probability in our context:
Lemma C.7. Assume, as in the context of Proposition B.3, that for $R$ with $|R|=k$ where each element of $R$ is sampled i.i.d. with probability $p_{j}=\frac{s_{j}}{S}$ from $\mathcal{F}$ and $w_{i}=\frac{S}{s_{j}|R|}=\frac{1}{k p_{j}}$ denotes the weight of a function $f_{i} \in R$ that corresponds to $f_{j} \in \mathcal{F}$. Then with probability at least $1-\frac{1}{k}$ it holds that $\sum_{i \in I_{1}^{\prime}} w_{i} \leq \frac{2 S}{s_{0}}$.

Proof of Lemma C.7. Let $x_{\pi(i)}$ be the $i$ th element of $R$. We set $Z_{i}=w_{\pi(i)}$ if $\pi(i) \in I_{1}$ and $Z_{i}=0$ otherwise. Then $\lambda=\mathbb{E}\left(Z_{i}\right)=\sum_{j \in I_{1}} p_{j} w_{j}=\sum_{j \in I_{1}} p_{j} \frac{1}{k p_{j}}=\frac{\left|I_{1}\right|}{k}$. Recall from Lemma 2.7 that $\left|I_{1}\right| \leq \frac{S}{s_{0}}$ and for $j \in I_{1}$ we have $s_{j}>s_{0}$. For the variance it follows

$$
\mathbb{E}\left[\left(Z_{i}-\lambda\right)^{2}\right]=\mathbb{E}\left[Z_{i}^{2}\right]-\mathbb{E}\left[Z_{i}\right]^{2} \leq \mathbb{E}\left[Z_{i}^{2}\right]=\sum_{j \in I_{1}} p_{j} w_{j}^{2}=\sum_{j \in I_{1}} \frac{1}{k^{2} p_{j}} \leq \sum_{j \in I_{1}} \frac{S}{k^{2} s_{0}}=\left|I_{1}\right| \cdot \frac{S}{k^{2} s_{0}} \leq \frac{S^{2}}{s_{0}^{2} k^{2}}
$$

Thus by independence of the $Z_{i}$ the variance of $Z=\sum_{i=1}^{k} Z_{i}$ is bounded by $\frac{S^{2}}{s_{0}^{2} k}$. Now applying Chebyshev's inequality yields

$$
P\left(Z \geq 2 \cdot \frac{S}{s_{0}}\right) \leq P\left(Z-\mathbb{E}(Z) \geq \frac{S}{s_{0}}\right) \leq \frac{\operatorname{Var}(Z)}{S / s_{0}} \leq \frac{S^{2} /\left(s_{0}^{2} k\right)}{S^{2} / s_{0}^{2}}=\frac{1}{k}
$$

By the technical Corollary C. 6 our goal of obtaining a coreset for $f$ reduces to obtaining a coreset for the substitute function

$$
\tilde{f}(X \beta)=\sum_{i \in[n] \backslash I_{1}} g\left(x_{i} \beta\right)+\sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right) .
$$

To this end we set $\mathcal{F}_{1}=\left\{w_{i} G_{i}^{+} \mid i \in I_{1}\right\}$ where $G_{i}^{+}(\beta)=G^{+}\left(x_{i} \beta\right)$ and $\mathcal{F}_{2}=\left\{w_{i} g_{i} \mid i \in I_{2}=[n] \backslash I_{1}\right\}$ where $g_{i}(\beta)=g\left(x_{i} \beta\right)$. Further we set $\mathcal{F}=\mathcal{F}_{1} \cup \mathcal{F}_{2}$ and show that the VC-dimension of $\mathcal{F}$ can be bounded as desired:
Lemma C.8. [Copy of Lemma 2.9] For the VC-dimension $\Delta$ of $\Re_{\mathcal{F}}$ we have

$$
\begin{aligned}
\Delta & \leq(d+1)\left(\log _{2}\left(\mu c \varepsilon^{-2}\right)+2\right) \\
& =O(d \log (\mu / \varepsilon))
\end{aligned}
$$

Proof of Lemma 2.9/C.8. First note that for any $G \subseteq \mathcal{F}_{1}, \beta \in \mathbb{R}^{d}$ and $r \in \mathbb{R}$ it holds that range $\mathcal{F}_{\mathcal{F}_{1}}(\beta, r) \cap G=$ $\operatorname{range}_{G}(\beta, r)$. We show that the VC-dimension of $\mathfrak{R}_{\mathcal{F}_{1}}$ is at most $d+1$. Indeed, it holds that range $\mathcal{F}_{\mathcal{F}_{1}}(\beta, r)=\mathcal{F}_{1}$ if $r \leq 0$ since all weights are positive and $G^{+}$is also positive. Otherwise we have that

$$
\begin{aligned}
\operatorname{range}_{\mathcal{F}_{1}}(\beta, r) & =\left\{w_{i} G_{i}^{+} \in \mathcal{F}_{1} \mid w_{i} G_{i}^{+}(\beta) \geq r\right\} \\
& =\left\{w_{i} G_{i}^{+} \in \mathcal{F}_{1} \mid w_{i}\left(x_{i} \beta\right)^{p} / p \geq r \wedge x_{i} \beta>0\right\}
\end{aligned}
$$

$$
=\left\{w_{i} G_{i}^{+} \in \mathcal{F}_{1} \left\lvert\, x_{i} \beta \geq\left(\frac{p r}{w_{i}}\right)^{1 / p}\right.\right\}
$$

We conclude that for $G \subseteq \mathcal{F}_{1}$ it holds that

$$
\begin{aligned}
\left|\left\{G \cap R \mid R \in \operatorname{ranges}\left(\mathcal{F}_{1}\right)\right\}\right| & =\left|\left\{\operatorname{range}_{G}(\beta, r) \mid \beta \in \mathbb{R}^{d}, r \in \mathbb{R}_{>0}\right\} \cup\left\{\operatorname{range}_{G}(\beta, r) \mid \beta \in \mathbb{R}^{d}, r \in \mathbb{R}_{\leq 0}\right\}\right| \\
& =\left|\left\{\left\{w_{i} G_{i}^{+} \in G \mid x_{i} \beta \geq\left(p r / w_{i}\right)^{1 / p}\right\} \mid \beta \in \mathbb{R}^{d}, r \in \mathbb{R}_{\geq 0}\right\} \cup\{G\}\right| \\
& \leq\left|\left\{\left\{w_{i} G_{i}^{+} \in G \mid x_{i} \beta-s \geq 0\right\} \mid \beta \in \mathbb{R}^{d}, s \in \mathbb{R}\right\}\right|
\end{aligned}
$$

which corresponds to a set of affine hyperplane classifiers $\beta \mapsto \mathbf{1}_{x_{i} \beta-s \geq 0}$, which have VC-dimension $d+1$ (Kearns and Vazirani, 1994). Thus, the induced range space $\mathfrak{R}_{\mathcal{F}_{1}}$ has VC-dimension at most $d+1$.
Next consider $\mathcal{F}_{2}$. Note that $g$ is a strictly monotonic and thus also invertible function. First fix a weight $v \in \mathbb{R}_{>0}$ and let $\mathcal{F}_{v}=\left\{w_{i} g_{i} \mid w_{i}=v\right\}$. We have

$$
\begin{aligned}
\operatorname{range}_{\mathcal{F}_{v}}(\beta, r) & =\left\{w_{i} g_{i} \in \mathcal{F}_{v} \mid w_{i} g_{i}(\beta) \geq r\right\} \\
& =\left\{w_{i} g_{i} \in \mathcal{F}_{v} \left\lvert\, x_{i} \beta \geq g^{-1}\left(\frac{r}{v}\right)\right.\right\}
\end{aligned}
$$

which corresponds to a set of points shattered by the affine hyperplane classifier $\beta \mapsto \mathbf{1}_{x_{i} \beta-g^{-1}\left(\frac{r}{v}\right) \geq 0}$ and thus the VC-dimension of the induced range space $\mathfrak{R}_{\mathcal{F}_{v}}$ is at most $d+1$. Let $W$ be the set of all weights for functions in $\mathcal{F}_{2}$. Since all weights are powers of 2 , and we have $\frac{S}{n} \leq v \leq \frac{\mu S c \ln \left(p \varepsilon^{-1}\right)}{\varepsilon n}$ it holds that $|W| \leq \log _{2}\left(\frac{\mu c \ln \left(p \varepsilon^{-1}\right)}{\varepsilon}\right) \leq$ $\left(\log _{2}\left(\mu c \varepsilon^{-2}\right)+2\right)$. Now we claim that the VC-dimension of $\mathfrak{R}_{\mathcal{F}}$ is at $\operatorname{most}(|W|+1)(d+1)$ as $\mathcal{F}=\mathcal{F}_{1} \cup \bigcup_{v \in W}^{\varepsilon} \mathcal{F}_{v}$. Assume for the sake of contradiction that there exists $G \subset \mathcal{F}$ such that $|G|>(|W|+1)(d+1)$ and $G$ is shattered by the ranges of $\mathcal{F}$. Then by the pigeonhole principle $G^{\prime}=G \cap \mathcal{F}^{\prime}>d+1$ for some $\mathcal{F}^{\prime} \in\left\{F_{1}\right\} \cup \bigcup_{v \in W}\left\{\mathcal{F}_{v}\right\}$. But due to the pairwise disjointness of all of $\mathcal{F}_{1}$ and $\mathcal{F}_{v}, G^{\prime}$ must be shattered by the ranges of $\mathcal{F}^{\prime}$, which contradicts that their VC-dimension is bounded by $d+1$, cf. Lemma 11 in (Munteanu et al., 2018).

## C. 2 Bounding the Sensitivities

We define the $\ell_{p}$-leverage scores of $X$ by $u_{j}=\sup _{\beta \in \mathbb{R}^{d} \backslash\{0\}} \frac{\left|x_{j} \beta\right|^{p}}{\sum_{i=1}^{n}\left|x_{i} \beta\right|^{p}}$, cf. (Dasgupta et al., 2009). We note that the supremum is attained by some $\beta \in \mathbb{R}^{d}$ since

$$
\sup _{\beta \in \mathbb{R}^{d} \backslash\{0\}} \frac{\left|x_{j} \beta\right|^{p}}{\sum_{i=1}^{n}\left|x_{i} \beta\right|^{p}}=\sup _{\beta \in \mathbb{R}^{d} \backslash\{0\}} \frac{\|\beta\|_{2}^{p} \cdot\left|x_{j} \beta /\|\beta\|_{2}\right|^{p}}{\|\beta\|_{2}^{p} \cdot \sum_{i=1}^{n}\left|x_{i} \beta /\|\beta\|_{2}\right|^{p}}=\sup _{\beta \in \mathbb{R}^{d},\|\beta\|_{2}=1} \frac{\left|x_{j} \beta\right|^{p}}{\sum_{i=1}^{n}\left|x_{i} \beta\right|^{p}}
$$

and $\left\{\beta \in \mathbb{R}^{d} \mid\|\beta\|_{2}=1\right\}$ is a compact set. We also note that $u_{p} \leq 1$ always holds.
Lemma C.9. [Copy of Lemma 2.10] There is a constant $c_{s}$ such that the sensitivity $\zeta_{i}$ of $x_{i}, i \in[n]$ for $\tilde{f}$ is bounded by

$$
\zeta_{i} \leq c_{s} \mu\left(\frac{1}{n}+u_{i}\right)
$$

Proof of Lemma 2.10/C.9. First note that by Lemma 2.8 it holds that $f(X \beta) \geq \frac{n}{\mu}$ and thus by Lemma 2.7 $\tilde{f}(X \beta) \geq f(X \beta)-\varepsilon \sum_{i \in I_{1}} G^{+}\left(x_{i} \beta\right)-\varepsilon \cdot \frac{n}{\mu} \geq \frac{f(X \beta)}{2} \geq \frac{n}{2 \mu}$ holds for small enough $\varepsilon \leq 1 / 4$. Thus for any $\beta$ with $x_{i} \beta \leq 1$ we have $\frac{G^{+}\left(x_{i} \beta\right)}{\tilde{f}(X \beta)} \leq \frac{g\left(x_{i} \beta\right)}{\tilde{f}(X \beta)} \leq \frac{g(1)}{n / 2 \mu}=2 g(1) \frac{\mu}{n}$.
Further for $\beta$ with $x_{i} \beta>1$ it holds that $g\left(x_{i} \beta\right) \leq c_{3}\left(x_{i} \beta\right)^{p}$ for some constant $c_{3} \leq 2 g(1)+1$ since by Lemma 2.6 and using $\frac{p-1}{t} \leq p-1$ for $t \geq 1$ it holds that

$$
\begin{aligned}
g\left(x_{i} \beta\right) & =g(1)+\int_{1}^{x_{i} \beta} g^{\prime}(t) d t \\
& \leq g(1)+\int_{1}^{x_{i} \beta} t^{p-1}+\frac{p-1}{t} d t
\end{aligned}
$$

$$
\begin{aligned}
& \leq g(1)+\int_{1}^{x_{i} \beta} t^{p-1}+p-1 d t \\
& \leq g(1)+\int_{1}^{x_{i} \beta} t^{p-1}+(p-1) t^{p-1} d t \\
& \leq g(1)+\int_{1}^{x_{i} \beta} p t^{p-1} d t \\
& =g(1)+\left(x_{i} \beta\right)^{p}-1 \leq(g(1)+1)\left(x_{i} \beta\right)^{p} .
\end{aligned}
$$

Also note that by definition of $\mu$ it holds that

$$
\frac{1}{\sum_{x_{j} \beta>0}\left|x_{j} \beta\right|^{p}(1+\mu)} \leq \frac{1}{\sum_{x_{j} \beta>0}\left|x_{j} \beta\right|^{p}+\sum_{x_{j} \beta<0}\left|x_{j} \beta\right|^{p}}=\frac{1}{\sum_{j=1}^{n}\left|x_{j} \beta\right|^{p}}
$$

and thus

$$
\frac{1}{\sum_{x_{j} \beta>0}\left|x_{j} \beta\right|^{p}} \leq \frac{1+\mu}{\sum_{j=1}^{n}\left|x_{j} \beta\right|^{p}} .
$$

Now setting $c_{3}=2 g(1)+1$ and using $\tilde{f}(X \beta) \geq \sum_{x_{j} \beta>0} \frac{\left|x_{j} \beta\right|^{p}}{p}=\sum_{j=1}^{n} G^{+}\left(x_{j} \beta\right)$ we get

$$
\frac{G^{+}\left(x_{i} \beta\right)}{\tilde{f}(X \beta)} \leq \frac{g\left(x_{i} \beta\right)}{\tilde{f}(X \beta)} \leq \frac{c_{3}}{1 / p} \cdot \frac{\left|x_{i} \beta\right|^{p}}{\sum_{x_{j} \beta>0}\left|x_{j} \beta\right|^{p}} \leq p c_{3}(1+\mu) u_{i} \leq 2 p c_{3} \mu u_{i}:=c_{s} \mu u_{i} .
$$

Combining both bounds gives us the bound for $\zeta_{i}$.

## C. 3 Well Conditioned Bases and Approximate Leverage Scores

In order to approximate the leverage scores we will need well conditioned bases:
An $(\alpha, \beta, p)$-well-conditioned basis $V$ is a basis that preserves the norm of each vector well, in the sense that its entry-wise $p$ norm $\|V\|_{p} \leq \alpha$ and for all $z \in \mathbb{R}^{d}:\|z\|_{q} \leq \beta\|V z\|_{p}$, where $q$ denotes the dual norm to $p$, i.e., $\frac{1}{p}+\frac{1}{q}=1$, see Definition 2.11. We will first state the properties of the ( $\alpha, \beta, p$ )-well-conditioned basis and then we describe how to compute the basis.
Lemma C.10. [Copy of Lemma 2.12] Let $V$ be an $(\alpha, \beta, p)$-well-conditioned basis for the column space of $X$. Then it holds for all $i \in[n]$ that $u_{i} \leq \beta^{p}\left\|v_{i}\right\|_{p}^{p}$. As a direct consequence we have $\sum_{i=1}^{n} u_{i} \leq \beta^{p}\|V\|_{p}^{p} \leq(\alpha \beta)^{p}=d^{O(p)}$.

Proof of Lemma 2.12/C.10. We have by a change of basis

$$
u_{i}=\sup _{z \in \mathbb{R}^{d} \backslash\{0\}} \frac{\left|(X z)_{i}\right|^{p}}{\|X z\|_{p}^{p}}=\sup _{z \in \mathbb{R}^{d} \backslash\{0\}} \frac{\left|(V z)_{i}\right|^{p}}{\|V z\|_{p}^{p}} .
$$

Now assume that $z$ attains the value $\sup _{z \in \mathbb{R}^{d} \backslash\{0\}} \frac{\mid(V z)_{i}^{p}}{\|V z\|_{p}^{p}}$. Then we get by using Hölder's inequality and the properties of $V$ that

$$
u_{i}=\frac{\left|(V z)_{i}\right|^{p}}{\|V z\|_{p}^{p}} \leq \frac{\beta^{p}\left|(V z)_{i}\right|^{p}}{\|z\|_{q}^{p}} \leq \frac{\beta^{p}\left\|v_{i}\right\|_{p}^{p}\|z\|_{q}^{p}}{\|z\|_{q}^{p}}=\beta^{p}\left\|v_{i}\right\|_{p}^{p} .
$$

An $(\alpha, \beta, p)$-well-conditioned basis can be computed using sketching techniques.
Lemma C.11. (Woodruff and Zhang, 2013; Clarkson and Woodruff, 2017)[Copy of Lemma 2.13] There exists a random embedding matrix $\Pi \in \mathbb{R}^{n^{\prime} \times n}$ and $\gamma=O(d \log (d))$ such that

$$
\forall \beta \in \mathbb{R}^{d}: \frac{1}{\gamma^{1 / p}}\|X \beta\|_{p} \leq\|\Pi X \beta\|_{q} \leq \gamma^{1 / p}\|X \beta\|_{p}
$$

holds with constant probability, where

$$
\left(q, n^{\prime}\right)= \begin{cases}\left(2, O\left(d^{2}\right)\right) & \text { if } p \in[1,2] \\ \left(\infty, O\left(n^{1-\frac{2}{p}} \log n(d \log d)^{1+\frac{2}{p}}+d^{5+4 p}\right)\right) & \text { if } p \in(2, \infty) .\end{cases}
$$

For $p=2$ we have $\gamma=2$. Further $\Pi X$ can be computed in $O(\mathrm{nnz}(X))$ time.

The sketching matrix $\Pi$ can be constructed as follows: First let $D \in \mathbb{R}^{n \times n}$ be the diagonal matrix with $D_{i i}=1$ or $D_{i i}=-1$ each with probability $1 / 2$. Further let $h:[n] \rightarrow\left[n^{\prime}\right]$ be a random map where $h$ hashes each entry of $[n]$ to one of $n^{\prime}$ buckets uniformly at random. Set $\Psi \in \mathbb{R}^{n^{\prime} \times n}$ to be the matrix where $\Psi_{h(i) i}=1$ and $\Psi_{j i}=0$ if $j \neq h(i)$. For $p=2$ it suffices to take $\Pi=\Psi D$ (Clarkson and Woodruff, 2017). Otherwise if $p \neq 2$ let $E$ be a diagonal matrix with $E_{i i}=1 / \lambda_{i}^{1 / p}$ where $\lambda_{i} \sim \exp (1)$ is drawn from a standard exponential distribution and set $\Pi=\Psi D E$ (Woodruff and Zhang, 2013).
Lemma C.12. [Copy of Lemma 2.14] If $\Pi$ satisfies Lemma 2.13 and $\Pi X=Q R$ is the $Q R$-decomposition of $\Pi X$ then $V=X R^{-1}$ is an $(\alpha, \beta, p)$-well-conditioned basis for the columnspace of $X$, where for $\gamma=O(d \log (d))$ we have

$$
(\alpha, \beta)= \begin{cases}(\sqrt{2 d}, \sqrt{2}), & \text { for } p=2 \\ \left(d \gamma^{1 / p}, \gamma^{1 / p}\right), & \text { for } p \in[1,2) \\ \left(d \gamma^{1 / p}, d \gamma^{1 / p}\right), & \text { for } p \in(2, \infty)\end{cases}
$$

Proof of Lemma 2.14/C.12. We are going to use the fact that $Q=\Pi X R^{-1}$ is an orthonormal basis. Let $e_{i}$ for $i \in[d]$ denote the $i$ th standard basis vector. We define $\left(R^{-1}\right)^{(i)}$ to be the $i$ th column of $R^{-1}$. We have

$$
\begin{align*}
\|V\|_{p} & =\left\|X R^{-1}\right\|_{p}=\left\|X \sum_{i=1}^{d}\left(R^{-1}\right)^{(i)} e_{i}^{T}\right\|_{p}=\left\|\sum_{i=1}^{d} X\left(R^{-1}\right)^{(i)} e_{i}^{T}\right\|_{p} \\
& \leq \sum_{i=1}^{d}\left\|X\left(R^{-1}\right)^{(i)} e_{i}^{T}\right\|_{p}=\sum_{i=1}^{d}\left\|X\left(R^{-1}\right)^{(i)}\right\|_{p} \tag{21}
\end{align*}
$$

Now suppose $p \in(2, \infty)$.

$$
\begin{aligned}
(21) & \leq \gamma^{1 / p} \sum_{i=1}^{d}\left\|\Pi X\left(R^{-1}\right)^{(i)}\right\|_{\infty} \leq \gamma^{1 / p} \sqrt{d}\left(\sum_{i=1}^{d}\left\|\Pi X\left(R^{-1}\right)^{(i)}\right\|_{\infty}^{2}\right)^{\frac{1}{2}} \\
& \leq \gamma^{1 / p} \sqrt{d}\left(\sum_{i=1}^{d}\left\|\Pi X\left(R^{-1}\right)^{(i)}\right\|_{2}^{2}\right)^{\frac{1}{2}} \leq \gamma^{1 / p} \sqrt{d}(\sum_{i=1}^{d} \underbrace{\left\|Q^{(i)}\right\|_{2}^{2}}_{=1})^{\frac{1}{2}}=\gamma^{1 / p} d
\end{aligned}
$$

For arbitrary $z \in \mathbb{R}^{d}$ it holds that

$$
\|z\|_{q} \leq \sqrt{d}\|z\|_{2}=\sqrt{d}\|Q z\|_{2}=\sqrt{d}\left\|\Pi X R^{-1} z\right\|_{2} \leq d\left\|\Pi X R^{-1} z\right\|_{\infty} \leq d \gamma^{1 / p}\|V z\|_{p}
$$

Consequently $V$ is $\left(\gamma^{1 / p} d, \gamma^{1 / p} d, p\right)$-well-conditioned.
Next suppose $p \in[1,2)$. Again we bound

$$
\begin{aligned}
(21) & \leq \gamma^{1 / p} \sum_{i=1}^{d}\left\|\Pi X\left(R^{-1}\right)^{(i)}\right\|_{2} \leq \gamma^{1 / p} \sqrt{d}\left(\sum_{i=1}^{d}\left\|\Pi X\left(R^{-1}\right)^{(i)}\right\|_{2}^{2}\right)^{\frac{1}{2}} \\
& \leq \gamma^{1 / p} \sqrt{d}(\sum_{i=1}^{d} \underbrace{\left\|Q^{(i)}\right\|_{2}^{2}}_{=1})^{\frac{1}{2}}=\gamma^{1 / p} d
\end{aligned}
$$

Also, since $p \leq 2$, the dual norm satisfies $q \geq 2$. Fix an arbitrary $z \in \mathbb{R}^{d}$. It follows that

$$
\|z\|_{q} \leq\|z\|_{2}=\|Q z\|_{2}=\left\|\Pi X R^{-1} z\right\|_{2} \leq \gamma^{1 / p}\|V z\|_{p}
$$

It follows that $V$ is even $\left(\gamma^{1 / p} d, \gamma^{1 / p}, p\right)$-well-conditioned in this case.

Finally suppose $p=2$, where the entry-wise matrix norm is the Frobenius norm $\|\cdot\|_{F}$. We have

$$
\begin{aligned}
\|V\|_{F}^{2}=\sum_{i=1}^{d}\left\|V^{(i)}\right\|_{2}^{2}=\sum_{i=1}^{d}\left\|\left(X R^{-1}\right)^{(i)}\right\|_{2}^{2} & \leq \sum_{i=1}^{d} 2\left\|\left(\Pi X R^{-1}\right)^{(i)}\right\|_{2}^{2} \\
& =2 \sum_{i=1}^{d}\left\|Q^{(i)}\right\|_{2}^{2}=2 d
\end{aligned}
$$

Thus we have $\|V\|_{F}=\sqrt{2 d}$. Since $p=q=2$ we have for any $\beta \in \mathbb{R}^{d}$ that

$$
\|z\|_{q}=\|z\|_{2}=\|Q z\|_{2}=\left\|\Pi X R^{-1} z\right\|_{2} \leq \sqrt{2}\|V z\|_{p}
$$

Consequently $V$ is a $(\sqrt{2 d}, \sqrt{2}, p)$-well-conditioned in this case.

## C. 4 Proof of Main Results

The main results, Theorem 1 and the improvement for $p=2$, Corollary 2.3 , are completely contained in the main body of the paper. It remains to prove Corollary 2.4 that handles the approximation factor of the solution obtained from the coreset with respect to the original loss function on the full data.
Corollary C.13. [Copy of Corollary 2.4] Let $\left(X^{\prime}, w\right)$ be a weighted $\varepsilon$-coreset for $f$. Let $\tilde{\beta} \in \operatorname{argmin}_{\beta \in \mathbb{R}^{d}} f_{w}\left(X^{\prime} \beta\right)$. Then it holds that $f(X \tilde{\beta}) \leq(1+3 \varepsilon) \min _{\beta \in \mathbb{R}^{d}} f(X \beta)$.

Proof of Corollary 2.4/C.13. Let $\beta^{*} \in \operatorname{argmin}_{\beta \in \mathbb{R}^{d}} f(X \beta)$. Since $\left(X^{\prime}, w\right)$ is a coreset we have by Definition 2.1 and using the optimality of $\tilde{\beta}$ for the coreset that

$$
\begin{aligned}
f(X \tilde{\beta}) & \leq f_{w}\left(X^{\prime} \tilde{\beta}\right) /(1-\varepsilon) \leq f_{w}\left(X^{\prime} \beta^{*}\right) /(1-\varepsilon) \\
& \leq f\left(X \beta^{*}\right)(1+\varepsilon) /(1-\varepsilon) \leq f\left(X \beta^{*}\right)(1+3 \varepsilon)
\end{aligned}
$$

## D EXPERIMENTS, PLOTS AND PSEUDO CODE

Here we provide further material deferred from the experimental section in the main body.
We briefly introduce the data sets that we used. Table D provides a summary. The Webspam ${ }^{8}$ data consists of 350000 unigrams with 128 features from web pages which have to be classified as spam or normal pages. The Covertype ${ }^{9}$ data consists of 581012 cartographic observations of different forests with 54 features. The task is to predict the type of trees at each location. The Kddcup ${ }^{10}$ data consists of 494021 network connections with 33 features and the task is to detect network intrusions. The Example-2D data consists of 175 synthetic data points ( 80 per class plus 15 outliers in one class) in 2 dimensions to visualize the results of $p$-probit regressions for different values of $p$.

Table 1: Summary of the used data sets and their dimensions: The given values of $d$ do not include the intercept. The data sets are downloaded or generated automatically by our open Python implementation that is available at https://github.com/cxan96/efficient-probit-regression.

| data set | $\mathbf{n}$ | $\mathbf{d}$ |
| :--- | ---: | ---: |
| Webspam | 350000 | 128 |
| Covertype | 581012 | 54 |
| Kddcup | 494021 | 33 |
| Example-2D | 175 | 2 |

```
Algorithm 1 Coreset algorithm for \(p\)-generalized probit regression.
    Input: data \(X \in \mathbb{R}^{n \times d}\), number of rows \(k\).;
    Output: coreset \(C=\left(X^{\prime}, w\right) \in \mathbb{R}^{k \times d}\) with \(k\) rows.;
    Initialize sketch \(X^{\prime \prime}=\mathbf{0} \in \mathbb{R}^{n^{\prime} \times d}\), (where \(n^{\prime}=O\left(d^{2}\right)\) for \(p \leq 2\) or \(n^{\prime}=O\left(n^{1-\frac{2}{p}} \log n \cdot \operatorname{poly}(d)\right)\) for \(\left.p>2\right)\);
    for \(i=1 \ldots n\) do
        Draw a random number \(B_{i} \in\left[n^{\prime}\right]\); \(\triangleright\) hash to bucket \(B_{i}\)
        Draw a random number \(\sigma_{i} \in\{-1,1\}\); \(\triangleright\) random sign
        if \(p \neq 2\) then
            Draw a random number \(\lambda_{i} \sim \exp (1) ; \quad \triangleright \ell_{p}\) embedding
            \(\sigma_{i}=\sigma_{i} / \lambda_{i}^{1 / p}\).
        \(X_{B_{i}}^{\prime \prime}=X_{B_{i}}^{\prime \prime}+\sigma_{i} \cdot x_{i} . \quad \triangleright\) sketch
    Compute the QR-decomposition of \(X^{\prime \prime}=Q R\).; \(\triangleright\) well-conditioned basis
    Initialize coreset \(X^{\prime}=\mathbf{0} \in \mathbb{R}^{k \times d} \quad \triangleright\) coreset points
    Initialize weights \(w=0 \in \mathbb{R}^{k}\); \(\triangleright\) coreset weights
    Initialize \(k\) independent weighted reservoir samplers \(S_{j}\), sampling row \(X_{j}^{\prime}\), for each \(j \in[k]\);
    Initialize \(G=I \in \mathbb{R}^{d \times d} ; \quad \triangleright\) Identity matrix
    if \(p=2\) and \(\ln n<d\) then
        Draw \(G \in \mathbb{R}^{d \times \ln n}\) with \(G_{i j} \sim N\left(0, \frac{1}{\ln n}\right) ; \quad \triangleright\) JL-embedding
    for \(i=1 \ldots n\) do
        Compute \(q_{i}=\left\|x_{i}\left(R^{-1} G\right)\right\|_{p}^{p} ; \quad \triangleright \ell_{p}\)-leverage score approximation
        for \(j=1 \ldots k\) do
            Feed \(s_{i}=q_{i}+1 / n\) to \(S_{j} ; \quad \triangleright\) unnormalized sampling probabilities
            if \(S_{j}\) samples \(x_{i}\) then
                \(w_{j}=1 /\left(k \cdot s_{i}\right) ; \quad \triangleright\) unnormalized weights
                \(X_{j}^{\prime}=x_{i} ; \quad \triangleright\) save row identity in the coreset
    \(w=w \cdot \sum_{i=1}^{n} s_{i} ;\)
                            \(\triangleright\) normalize weights
    return \(C=\left(X^{\prime}, w\right)\);
```

[^8]```
Algorithm 2 Online coreset algorithm for the standard probit model ( \(p=2\) ).
    Input: data \(X \in \mathbb{R}^{n \times d}\), number of rows \(k\);
    Output: coreset \(C=\left(X^{\prime}, w\right) \in \mathbb{R}^{k \times d}\) with \(k\) rows.;
    Initialize \(M=M_{i n v}=Q=0 \in \mathbb{R}^{d \times d}\);
    Initialize \(k\) independent weighted reservoir samplers \(S_{j}\), sampling row \(X_{j}^{\prime}\), for each \(j \in[k]\);
    for \(i=1 \ldots n\) do
        \(x_{i} \in \mathbb{R}^{d}:=i\) th column of \(X^{T} ;\)
        \(M=M+x_{i} x_{i}^{T} ;\)
        if \(\left\|Q x_{i}\right\|_{2}=\left\|x_{i}\right\|_{2}\) then \(\quad \triangleright\) is \(x_{i}\) in column-space of Q ?
                \(M_{i n v}=M_{i n v}-\frac{M_{i n v} x_{i} x_{i}^{T} M_{i n v}}{1+x_{i}^{T} M_{i n v} x_{i}} ; \quad \triangleright\) adapted Sherman-Morrison formula
        else
            \(M_{i n v}=M^{\dagger} ; \quad \triangleright\) Moore-Penrose pseudoinverse
            \(Q R=M ; \quad \triangleright Q R\) decomposition of \(M\)
        \(\ell_{i}=\min \left\{x_{i}^{T} M_{i n v} x_{i}, 1\right\} ; \quad \triangleright\) approximate \(\ell_{2}\)-leverage scores
        for \(j=1 \ldots k\) do
            Feed \(s_{i}=\ell_{i}+\frac{1}{n}\) to \(S_{j} ; \quad \triangleright\) unnormalized sampling probabilities
            if \(S_{j}\) samples \(x_{i}\) then
                \(w_{j}=1 /\left(k \cdot s_{i}\right) ; \quad \triangleright\) unnormalized weights
                \(X_{j}^{\prime}=x_{i} ; \quad \triangleright\) save row in coreset
    \(w=w \cdot \sum_{i=1}^{n} s_{i} ; \quad \triangleright\) normalize weights
    return \(C=\left(X^{\prime}, w\right)\);
```



Figure 4: (Q1) \& (Q2): (top) A 2D example data set that demonstrates how different values of $p$ affect the linear separator, as well as the distribution of the residuals $X \beta /\|\beta\|_{2}$ of misclassifications. (middle \& bottom) Comparison of normed coefficients for logit vs. 1-probit model for different data sets.


Figure 5: Comparison of approximation ratios for $p \in\{1,1.5,2,3,5\}$ on the data sets Covertype, Webspam and Kddcup. The median approximation ratio denotes for each sample size the median of the approximation ratios over all repetitions. The solid line indicates the median, and the shaded area indicates the normalized interquartile range.


Figure 6: Comparison of total running times (including optimization) vs. accuracy for $p \in\{1,1.5,2,3,5\}$ on the data sets Covertype, Webspam and Kddcup. The median approximation ratio (resp. running time) denotes for each sample size the median of the approximation ratios (resp. running times) over all repetitions.


[^0]:    Proceedings of the $25^{\text {th }}$ International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

[^1]:    ${ }^{1}$ Extending to $p \in(0, \infty)$ is in principal possible, but $p<1$ leads to non-convex level sets in the multivariate case and the maximum likelihood estimation problem thus becomes inefficient to solve.

[^2]:    ${ }^{2}$ Implementations are available in the Python package scipy.stats.gennorm as well as in the $R$ package pgnorm.

[^3]:    ${ }^{3}$ The connection is not explicitly elaborated in those references.

[^4]:    ${ }^{4}$ The distributions differ more significantly near their means, i.e., around zero.

[^5]:    ${ }^{5}$ It is widely believed that $O(d)$ suffices but there is no formal proof for this, except in the case of equal weights.

[^6]:    ${ }^{6}$ Indeed, $y_{i}$ and $z_{i}$ always appear in that combination so we can assume the data consists only of $x_{i} \in \mathbb{R}^{d}$ for $i \in[n]$.

[^7]:    ${ }^{7}$ Here, $T=\operatorname{poly}(d)$ means that there exists a constant $c \geq 1$ such that $T=\Theta\left(d^{c}\right)$.

[^8]:    ${ }^{8}$ https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html\#webspam
    ${ }^{9}$ https://archive.ics.uci.edu/ml/datasets/Covertype
    ${ }^{10}$ https://kdd.ics.uci.edu/databases/kddcup99/kddcup99.html

