

# Effective sample size: a measure of individual uncertainty in predictions

Doranne Thomassen<sup>1</sup>, Saskia le Cessie<sup>1,2</sup>, Hans van Houwelingen<sup>1</sup>, Ewout Steyerberg<sup>1</sup>

<sup>1</sup>Biomedical Data Sciences, Leiden University Medical Center, The Netherlands

<sup>2</sup>Clinical Epidemiology, Leiden University Medical Center, The Netherlands

September 19, 2023

## Abstract

Clinical prediction models are estimated using a sample of limited size from the target population, leading to uncertainty in predictions, even when the model is correctly specified. Generally, not all patient profiles are observed uniformly in model development. As a result, sampling uncertainty varies between individual patients' predictions. We aimed to develop an intuitive measure of individual prediction uncertainty.

The variance of a patient's prediction can be equated to the variance of the sample mean outcome in  $n_*$  hypothetical patients with the same predictor values. This hypothetical sample size  $n_*$  can be interpreted as the number of similar patients  $n_{\text{eff}}$  that the prediction is effectively based on, given that the model is correct. For generalised linear models, we derived analytical expressions for the effective sample size. In addition, we illustrated the concept in patients with acute myocardial infarction.

In model development,  $n_{\text{eff}}$  can be used to balance accuracy versus uncertainty of predictions. In a validation sample, the distribution of  $n_{\text{eff}}$  indicates which patients were more and less represented in the development data, and whether predictions might be too uncertain for some to be practically meaningful. In a clinical setting, the effective sample size may facilitate communication of uncertainty about predictions.

We propose the effective sample size as a clinically interpretable measure of uncertainty in individual predictions. Its implications should be explored further for the development, validation and clinical implementation of prediction models.

## 1 Introduction

When using a statistical model to predict a patient's outcome, there is uncertainty in predictions, even when the model is correctly specified. This uncertainty arises from the fact that a prediction model is estimated using a random sample of limited size from the target population. Typically, sampling uncertainty about predictions varies between individual patients, as not all patient profiles are represented uniformly in model development.

Paradoxically, although a rarer type of patient stands to benefit most from personalised prediction, their predictions are usually more uncertain compared to the average patient. Yet between-patient variation in prediction uncertainty is typically not viewed as a target of optimisation in model development. Neither is it explored as a performance metric in model validation. The TRIPOD guideline for prediction model reporting does not mention any explicit metrics for individual prediction uncertainty [1].

In the risk communication literature, it has been recommended to present absolute risks in terms of frequencies with a common denominator, preferably with pictograms [2], [3]. For example, a patient’s predicted risk can be expressed as a number out of “100 people like you” [4]. This presentation fails to capture the uncertainty of the prediction. There is no guideline on the communication of uncertainty around predictions, or whether to communicate uncertainty at all. Statisticians tend to present sampling uncertainty about estimates with a 95% confidence interval or a standard error. However, communicating sampling uncertainty to patients in the form of confidence intervals has been suggested to have little effect on risk perceptions [5], [6] and to increase patient worry [6], [7]. Another suggestion has been to use verbal qualifiers, such as “it is most likely that your risk is 30%” or “our best estimate is that you have a 40% risk, though we are not certain about this exact number”. It has also been suggested to communicate numbers only when they are estimated with sufficient certainty [8].

As an alternative to standard errors and confidence intervals, we aimed to translate the standard error into an “effective sample size” being close in spirit to a “number of patients like you”. This could serve as a more intuitive measure of individual sampling uncertainty in predictions. The effective sample size can then be used as a communication tool in clinical practice and as a performance measure in model validation. In Section 2, we define effective sample size in this context and we derive explicit formulas for effective samples sizes in linear and logistic regression. In addition, we illustrate the concept in an example dataset of patients with acute myocardial infarction (Section 3). Finally, we discuss the potential applications of effective sample size in prediction modelling and we suggest directions for further investigation (Section 4).

## 2 Effective sample size and relative variance

### 2.1 Data setting

Suppose that for patients  $i \in \{1, \dots, n\}$ , we have observed the  $n \times 1$  outcome vector  $Y$ . We have also observed baseline covariates and coded these as desired in the  $n \times p$  design matrix  $X$ .  $Y$  and  $X$  constitute our *development data*. To the development data, a generalised linear model (GLM)  $\mathbb{E}[Y] = g^{-1}(X\beta)$  was fitted. The estimated linear predictor for patient  $i$  is denoted as  $x_i^\top \hat{\beta}$ , where  $\hat{\beta}$  is an estimate of  $\beta$ . After fitting the model to the development sample, it can be used to predict the expected value of  $y_{\text{new}}$  for a new patient, given their covariates  $x_{\text{new}}$ .

### 2.2 Effective sample size

The overall predictive performance of the model can be assessed by comparing predicted and observed values in the development sample (internal validation) or in an external sample of observed patients. A standard validation procedure assesses discrimination and calibration [9]. It does not give any personal measure of uncertainty for a new individual for whom we make a prediction.

Say there is such a new patient, named Sam. We use the prediction model to predict their

expected outcome  $y_{\text{new}}$  based on their covariates  $x_{\text{new}}$ . We denote the predicted outcome as  $\hat{y}_{\text{new}}$ . To express the individual uncertainty in Sam’s prediction, we draw a parallel between the variance of this prediction and the variance of the sample mean  $\bar{y}_*$  in a hypothetical independent sample of  $n_*$  patients like Sam. Here, being ‘like Sam’ is defined with respect to the model, so having the same design matrix entries as Sam. Suppose that we observe  $\bar{y}_* = \hat{y}_{\text{new}}$  in this hypothetical sample. We define Sam’s effective sample size as  $n_*$ , such that we have as much certainty about Sam’s prediction as we would have had, had we directly observed the predicted value in the hypothetical sample of  $n_*$  patients like Sam. In other words, given that the model is correct, Sam’s prediction is effectively based on  $n_*$  similar patients in the development of the model.

The concept of effective sample size can be used as a clinically interpretable measure of individual uncertainty. If the effective sample size for a new patient is small, their covariate profile is exceptional relative to the development data. The development data did not contain much information about patients like them, so the prediction should be interpreted with caution. By contrast, if the effective sample size for a new patient is large, we can have more confidence in their prediction.

We show in detail how  $n_*$  can be derived for linear and logistic regression models and give a general formulation for other generalized linear models.

### 2.3 Linear regression model

In the linear model,  $y_i$  are assumed to be independent samples from a normal distribution with mean  $\mu_i = x_i^\top \beta$  and variance  $\sigma^2$ . The least squares and maximum likelihood estimator for  $\beta$  is  $\hat{\beta} = (X^\top X)^{-1} X^\top Y$ , whose covariance matrix is  $\text{Cov}(\hat{\beta}) = \sigma^2 (X^\top X)^{-1}$ . Therefore, the variance of a prediction for a new patient, say, Sam, is  $\text{Var}(\hat{\mu}_{\text{new}}) = \text{Var}(x_{\text{new}}^\top \hat{\beta}) = \sigma^2 x_{\text{new}}^\top (X^\top X)^{-1} x_{\text{new}}$ .

Now suppose there was an independent sample of  $n_*$  patients who all have the same design matrix entries  $x_{\text{new}}$  and the same prediction  $\hat{\mu}_{\text{new}}$  as Sam. In addition, suppose that we had observed the sample mean  $\bar{y}_*$  to be equal to Sam’s prediction  $\hat{\mu}_{\text{new}}$ . We could use the sample mean  $\bar{y}_*$  as an estimator for  $\mu_{\text{new}}$ , which would have variance  $\sigma^2/n_*$ . Solving  $\frac{\sigma^2}{n_*} = \text{Var}(\hat{\mu}_{\text{new}}) = \sigma^2 x_{\text{new}}^\top (X^\top X)^{-1} x_{\text{new}}$  for  $n_*$  yields

$$n_* = \frac{1}{x_{\text{new}}^\top (X^\top X)^{-1} x_{\text{new}}}. \quad (1)$$

In the linear model case, the effective sample size is fully determined by the design matrix and does not depend on the outcomes  $Y$  through the fitted model. The inverse effective sample sizes for the patients in the development data are equal to the diagonal elements  $h_{ii}$  of the hat (projection) matrix, sometimes called leverages. It follows that the harmonic mean of the effective sample sizes in the development data is equal to  $n/p$  (where  $p$  is the number of parameters in the model).

Usually, leverages are examined as part of model diagnostics. High leverage points have a large potential influence on the fitted model, due to their relatively ‘extreme’ position in the  $X$ -space. Analogously, the effective sample size of a new prediction expresses how exceptional this new patient’s profile is relative to the development data and the fitted model.

### 2.4 Logistic regression model

We now turn to binary outcome variables  $y_i$  and assume  $y_i \sim \text{Bin}(n, p_i)$ ,  $p_i = \frac{e^{x_i^\top \beta}}{1 + e^{x_i^\top \beta}} = g^{-1}(x_i^\top \beta)$ , such that  $g$  is the logit function. To derive the effective sample size for a new prediction, again

for Sam, we require the variance of the prediction  $g^{-1}(x_{\text{new}}^\top \hat{\beta})$ . Using the delta method (local linearisation), we have

$$\text{Var}(g^{-1}(x_{\text{new}}^\top \hat{\beta})) \approx \text{Var}(x_{\text{new}}^\top \hat{\beta}) \cdot \left( \frac{\delta g^{-1}}{\delta(x_{\text{new}}^\top \hat{\beta})} \right)^2. \quad (2)$$

By general properties of the variance, the first factor on the right hand side of (2) can be expressed as

$$\text{Var}(x_{\text{new}}^\top \hat{\beta}) = x_{\text{new}}^\top \text{Cov}(\hat{\beta}) x_{\text{new}}. \quad (3)$$

When  $\hat{\beta}$  is the maximum likelihood estimator,  $\text{Cov}(\hat{\beta})$  is usually estimated as  $(X^\top V X)^{-1}$ , where  $V$  is a diagonal matrix with elements  $v_{ii} = \hat{p}_i(1 - \hat{p}_i)$ . For the second factor on the right hand side of (2), we have

$$\frac{\delta g^{-1}}{\delta(x_{\text{new}}^\top \hat{\beta})} = \frac{e^{x_{\text{new}}^\top \hat{\beta}}}{(1 + e^{x_{\text{new}}^\top \hat{\beta}})^2} \quad (4)$$

$$= g^{-1}(x_{\text{new}}^\top \hat{\beta})(1 - g^{-1}(x_{\text{new}}^\top \hat{\beta})) = \hat{p}_{\text{new}}(1 - \hat{p}_{\text{new}}). \quad (5)$$

Substituting (3) and (4) into (2) yields

$$\text{Var}(g^{-1}(x_{\text{new}}^\top \hat{\beta})) \approx x_{\text{new}}^\top \text{Cov}(\hat{\beta}) x_{\text{new}} \cdot (\hat{p}_{\text{new}}(1 - \hat{p}_{\text{new}}))^2 \quad (6)$$

for the variance of Sam's prediction.

For a binary variable, the variance is fully determined by the mean. Again, we imagine we have an independent sample of  $n_*$  patients with the same design matrix entries  $x_{\text{new}}$  and prediction  $\hat{p}_{\text{new}}$  as Sam. In this sample, we assume to have observed a sample mean of  $\bar{y}_* = \hat{p}_{\text{new}}$ . The sample mean  $\bar{y}_*$  as an estimator for  $\mathbb{E}[y_{\text{new}}] = p_{\text{new}}$  would have estimated variance  $\frac{\hat{p}_{\text{new}}(1 - \hat{p}_{\text{new}})}{n_*}$ . Equating this expression to (6) and solving for  $n_*$  results in an estimate for Sam's effective sample size:

$$n_* \approx \frac{\hat{p}_{\text{new}}(1 - \hat{p}_{\text{new}})}{x_{\text{new}}^\top \text{Cov}(\hat{\beta}) x_{\text{new}} \cdot (\hat{p}_{\text{new}}(1 - \hat{p}_{\text{new}}))^2} \quad (7)$$

$$= \left( x_{\text{new}}^\top \text{Cov}(\hat{\beta}) x_{\text{new}} \cdot (\hat{p}_{\text{new}}(1 - \hat{p}_{\text{new}})) \right)^{-1}. \quad (8)$$

Here, the inverse of the estimated effective sample size for patients in the development data is equal to their corresponding diagonal element of the approximated 'hat' matrix (approximated leverage) in logistic regression:  $V^{\frac{1}{2}} X (X^\top V X)^{-1} X^\top V^{\frac{1}{2}}$  [10].

In contrast with the linear model case, the effective sample size depends explicitly on  $\hat{\beta}$ , and thereby on  $Y$ , for logistic regression models. It also depends on the predicted risk  $\hat{p}$ , which will draw the effective sample size towards infinity for patients whose predicted risk is close to 0% or 100%. On one hand, it is counterintuitive that extreme extrapolation can lead to exceedingly large effective sample sizes. On the other hand, if a monotonic relation between  $x$  and  $p$  is assumed, and  $p$  is bounded, then it makes sense that certainty about  $\hat{p}$  being equal to 0 or 1 increases when  $x$  becomes more extreme and the model is correct.

## 2.5 Relative variance

The inverse of the effective sample size,  $n_*^{-1}$ , can be interpreted as a *relative variance*, i.e., the prediction variance normalised by the residual or conditional variance. For the linear model,  $n_*^{-1}$  is equal to the variance of the prediction  $\hat{\mu}_{\text{new}}$  divided by the residual variance of  $y_{\text{new}}$ :

$$\text{RelVar} = n_*^{-1} = x_{\text{new}}^\top (X^\top X)^{-1} x_{\text{new}} = \text{Var}(\hat{\mu}_{\text{new}}) / \sigma^2. \quad (9)$$

For the logistic regression model, we have

$$\text{RelVar} = n_*^{-1} \approx x_{\text{new}}^\top \text{Cov}(\hat{\beta}) x_{\text{new}} \cdot (\hat{p}_{\text{new}}(1 - \hat{p}_{\text{new}})). \quad (10)$$

Though less suitable as a communication tool, the relative variance can be used in lieu of effective sample size in model validation.

## 2.6 Other generalised linear models

For further generalised linear models, an approximation of the effective sample size can be obtained as in Section 2.4. Suppose again that we have a model prediction  $\hat{y}_{\text{new}} = g^{-1}(x_{\text{new}}^\top \hat{\beta})$  for Sam, based on a fitted GLM, where  $g$  is the link function of the GLM. As before, we then assume that we have observed Sam's prediction as the sample mean  $\bar{y}_*$  in a hypothetical sample of  $n_*$  patients with the same  $x$ -values as Sam. Given that the model is correct, the sample mean  $\bar{y}_*$  has variance  $\frac{\text{Var}(Y|x_{\text{new}}^\top \hat{\beta})}{n_*}$ .

To approximate the prediction variance, we can apply the delta method in (2):  $\text{Var}(\hat{y}_{\text{new}}) \approx x_{\text{new}}^\top (X^\top V X)^{-1} x_{\text{new}} \cdot \left( \frac{\delta g^{-1}}{\delta(x_{\text{new}}^\top \hat{\beta})} \right)^2$ . Here,  $V$  is a diagonal matrix with diagonal elements  $v_{ii} = \left( \frac{\delta g^{-1}}{\delta(x_i^\top \hat{\beta})} \right)^2 \left( \text{Var}(Y_i|x_i^\top \hat{\beta}) \right)^{-1}$ . An expression for Sam's effective sample size follows:

$$n_* \approx \frac{\text{Var}(Y|x_{\text{new}}^\top \hat{\beta})}{x_{\text{new}}^\top (X^\top V X)^{-1} x_{\text{new}} \cdot \left( \frac{\delta g^{-1}}{\delta(x_{\text{new}}^\top \hat{\beta})} \right)^2} \quad (11)$$

For patients in the development data, this expression reduces to the inverse of their approximated leverages  $V^{\frac{1}{2}} X (X^\top V X)^{-1} X^\top V^{\frac{1}{2}}$ .

## 3 Application to GUSTO data

### 3.1 Data description

We will illustrate the effective sample size in a large data set of patients with acute myocardial infarction (GUSTO-I) [11], publicly available from <https://hbiostat.org/data/gusto.rda>. This data set has been used to illustrate prediction modelling methodology before [12]–[16]. The dataset contains data on 40,830 patients, of whom 2851 (7.0%) died within the first 30 days after myocardial infarction.

## 3.2 Illustration with two predictors

To illustrate the concept of effective sample size, we apply the theory of the previous section to a smaller subset of the GUSTO-I dataset ( $n = 1214$ , comprising study sites from one US region and one non-US region). The variables that we used are DAY30 (a binary indicator of death within 30 days), age, height and shock (a binary indicator of whether the patient was in shock at hospital admission). The variables age and height were centered at their marginal means for the analyses. On average, patients were 61 (SD 11) years old and stood 170 (SD 10) centimeters tall. A minority of 28 (2.3%) patients presented with shock at admission. 76 patients (6.3%) died within 30 days after admission. In previous analyses, age and shock were strong predictors of death within 30 days. Height was uncorrelated with age.

For illustrative purposes, we calculated effective sample sizes for two linear regression models. In the first model, the predictors were age and height; in the second model, age and shock were predictors. As effective sample sizes are independent of the outcome in linear regression, the outcome is not relevant here. We also determined effective sample sizes for the following logistic regression models: (1)  $\text{DAY30} \sim \text{age}$ ; (2)  $\text{DAY30} \sim \text{age} + \text{height}$ ; (3)  $\text{DAY30} \sim \text{age} + \text{shock}$ .

For the linear regression models, lower effective sample sizes were observed in patients whose covariate values are more ‘rare’ relative to the joint distribution of the covariates (Figures 1a, 1b). Patients whose covariate values are closer to the average have higher effective sample sizes. Note that in the linear regression case, none of the effective sample sizes exceed the actual total sample size of 1214.

In the logistic models, similar properties appear, though now combined with the impact of the predicted risk being close to zero for some patients (Figures 1c-1d,2). When the predicted risk is close to 0 or 100%, the effective sample sizes increase, even when the prediction was based on rarer covariate values. This is clearly visible in the upper left corner of Figure 1c. Younger patients have a lower 30-day mortality risk, which is so close to 0 that their effective sample sizes are actually larger than those of patients whose age was closer to average. Some effective sample sizes even exceed 1214. This illustrates that the effective sample size is not a ‘real’ sample size, but a re-expression of the uncertainty of the prediction. In contrast to linear regression, effective sample sizes for logistic regression may first decrease with extrapolation and then increase again as covariate values become so extreme that predictions approach 0 or 100%.

The models with shock illustrate effective sample sizes when the model contains a rare binary covariate. From 1214 GUSTO patients, only 28 presented with shock at hospital admission. The effect of shock on death within 30 days is large but uncertain, due to the small number of patients. This is reflected in effective sample sizes for patients with shock, ranging from 25.2 to 28.0 in the linear regression case (Figure 1b). In the logistic regression case, effective sample sizes for these patients ranged from 22.1 to 55.7, as shock is a strong predictor for the outcome and the outcome influences  $n_{\text{eff}}$  here (Figures 1d, 2b). Telling a patient with shock that their prediction is based on a study of 1214 patients does not express that there were effectively about 28 patients like them in that study.

Finally, we visualised a comparison of effective sample sizes in the logistic model with age and shock to a model with age only (Figure 3). For the majority of patients without shock, the effective sample size is increased slightly when shock is added to the model, whereas for patients with shock  $n_{\text{eff}}$  decreases dramatically. There is a bias-variance tradeoff here: adding shock to the model may bring predictions closer to the truth for some patients, though at the cost of increased uncertainty.




Figure 1: Heatmap of effective sample sizes in the GUSTO subsample ( $n = 1214$ ) when two-predictor regression models are fitted to predict the outcome of 30-day mortality. Effective sample sizes in linear regression are independent of the outcome, hence the outcome is denoted more generally as  $Y$  for these models. In all models, the variables Age and Height were centered at their respective marginal means (61 years and 170 cm). Only a minority of patients in the data presented with shock at admission to hospital.

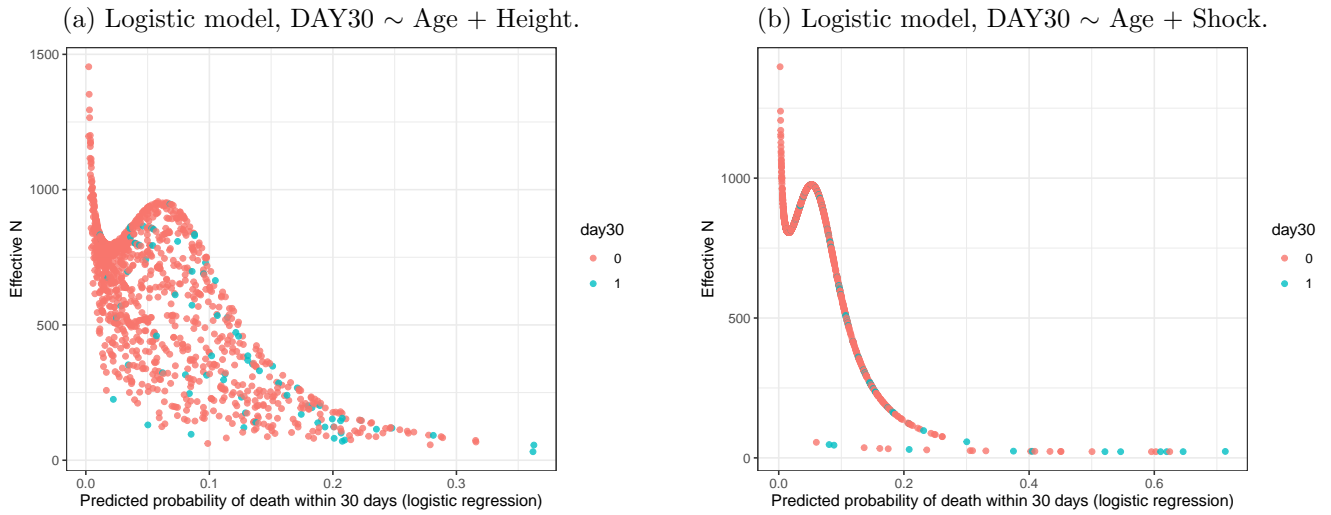


Figure 2: Effective sample sizes and the corresponding predicted probabilities of 30-day mortality for patients in the GUSTO subsample ( $n = 1214$ ), based on two logistic regression models. In all models, the variables Age and Height were centered at their respective marginal means (61 years and 170 cm). Only a minority of patients in the data presented with shock at admission to hospital.




Figure 3: Effective sample sizes and the corresponding predicted probabilities for patients in the GUSTO dataset. Two logistic regression models were fitted to the data to obtain predictions:  $\text{DAY30} \sim \text{Age}$  and  $\text{DAY30} \sim \text{Age} + \text{Shock}$ . Predictions for the same patient are connected by a solid line.






Figure 4: Histograms of effective sample sizes in the GUSTO US data (development sample) for a 7- and 15-predictor model for 30-day mortality.

### 3.3 Development and validation of published prediction models

To illustrate the use of effective sample sizes in model development and external validation, we used the full GUSTO dataset. We split the data into a development sample of data that was collected within the United States (US,  $n=23,034$ ) and an external validation sample of data that was collected elsewhere (non-US,  $n=17,796$ ).

Two previously published prediction models were fitted to the US data: one model with 7 predictor variables and 13 model parameters [17]; and a more complex model with 15 predictor variables and 29 model parameters [18]. We then calculated effective sample sizes for each patient in the development (US) data and summarized their distribution (Figure 4). Lower effective sample sizes are more frequent for the 15 predictor model compared to the 7 predictor model. The total sample size of 23,034 is large, however, hence rarer patient profiles are still quite well-represented in absolute numbers. There were five and sixteen patients with  $n_{\text{eff}}$  below 30 for the 7- and 15-predictor models, respectively.

We then calculated effective sample sizes for the previously fitted models in the external validation sample (Figure 5). For the 7-predictor model, the distribution of effective sample sizes in the validation sample was very similar to that in the development sample. So, with respect to the 7-predictor model the non-US patients overall seem to have a similar prediction uncertainty to the US patients and the occurrence of rare or common covariate patterns seems similar as well. For the 15-predictor model, lower effective sample sizes occur more in the validation sample than they did in the development sample. This indicates that the covariate patterns of patients in the validation sample were somewhat further away from the ‘average patient’ in the development sample.




Figure 5: Distribution of effective sample sizes in the GUSTO US data (development sample) vs the non-US data (external validation sample) for a 7- and 15-predictor model for 30-day mortality.

## 4 Discussion

We have defined the effective sample size, which can be used as a measure of individual sampling uncertainty in predictions. We derived explicit formulas for the effective sample sizes for linear and generalized linear models. With linear regression models, the effective sample sizes are the inverted leverages and they do not depend on the outcome. In the logistic regression case, effective sample sizes depend on the prediction (and therefore on the outcome) as well as the covariate values, and they increase when the predictions approach 0 or 1. Even with an exceptionally large total development sample size ( $> 23,000$ ), we have shown examples in which effective sample sizes dropped below 30 for some patients, indicating large differences in effective representation of patient types in the development of the 7- and 15-predictor models.

We see three directions for application of the concept of effective sample size: in model development and validation; to communicate predicted risks to individual patients; and as a tool to detect patient groups who were underrepresented in model development. In model development, there can be at least two aims regarding the effective sample sizes. The first is that for all patients in the (development/validation) data, their  $n_{\text{eff}}$  is not too low, adding to an overall sample size requirement for prediction model development[19]. This would require specification of a minimal  $n_{\text{eff}}$  for the prediction to be of added value in clinical practice. The minimal effective samples size in model development can be increased by refraining from adding rare binary covariates to a model, like the shock variable in the GUSTO example. There is a bias-variance tradeoff here, however, since shock is a strong predictor for the outcome. Furthermore,  $n_{\text{eff}}$  is defined relative to a specific model and provides a measure of uncertainty conditional on the model. Therefore, it is debatable whether  $n_{\text{eff}}$  can be compared between different models.

Another objective in model building could be that the effective sample size is more or less the same in all patients, or that their spread remains within certain bounds. Note that shrinkage methods such as Ridge and Lasso regression shrink predictions via the regression coefficients[20],

which are not individual to each patient. To even out effective sample sizes in a sample of patients, individual shrinkage of the covariates similar to Winsorizing [21] can be applied. This creates a nonlinear relationship between the original covariate and the linear predictor. Further investigation is needed in this direction.

In a model validation setting, the distribution of effective sample sizes in a validation sample summarises individual sampling uncertainty around predictions in a sample of interest. This provides an idea of the variability of prediction uncertainty across patients and patient profiles, as well as a measure of how different (with respect to the model) the validation sample is compared to the development sample. Therefore, the distribution of effective sample sizes could be an addition to standard performance metrics based on point estimates, such as the c-statistic and calibration summaries [1], [9].

In addition to the statistical applications of effective sample sizes in prediction, we propose  $n_{\text{eff}}$  as a communication tool in clinical practice. The effective sample size could be used to inform both patients and clinicians about sampling uncertainty in predictions, with statements as “this number is effectively based on 55 patients like you”. Further empirical research is needed to determine whether  $n_{\text{eff}}$  indeed has beneficial effects on the communication between clinician and patient, and between clinician and model developer.

Furthermore, effective sample sizes can detect groups of patients to whom we would like to apply a prediction model in clinical practice, but who were effectively underrepresented in the development of the model. We have shown that effective sample sizes can be very small for certain patients, even when the total development sample size was as large as 23,000. Most prediction models are developed with far smaller sample sizes, however. In such a case, effective sample sizes function as a warning sign that the model may not be sufficiently trustworthy to apply in practice for specific groups of patients. When such groups of patients are defined by characteristics historically subject to social biases, detection of their underrepresentation is important in light of recent discussions on algorithmic fairness [22]–[25].

In this paper, we focused our derivations on linear and logistic regression models, and provided a generalisation for GLMs. To make the concept of effective sample size more widely applicable, expressions for other types of regression models need to be obtained. In the context of clinical prediction models, the Cox model is a very relevant example. As the variance around a predicted risk from a Cox model is based not only on observed covariates but also on a censoring process, a sample of hypothetical patients “like” our new patient is not immediately clearly defined.

Another direction is to develop a generic numerical algorithm to obtain effective sample sizes when an analytical expression is not available and/or in cases of more complex and flexible models. We suspect that some modelling assumptions are necessary to move from the intuitive definition of  $n_{\text{eff}}$  as ‘hypothetical sample with the same variance’ to a formulation as ‘prediction variance divided by outcome variance’. If these assumptions can be made, then bootstrap or other sampling-based estimates of the prediction and outcome variances may be used to approximate effective sample sizes. We intend to explore this further, with links to a recently proposed ‘instability index’[26].

All in all, the concept of effective sample size may prove useful as an intuitive measure of individual prediction uncertainty. Its implications should be explored further for the development, validation and clinical implementation of prediction models.

## Financial disclosure

None reported.

## **Conflict of interest**

The authors declare no potential conflict of interests.

## **Data availability statement**

The data that support the results of this paper were loaded from a public repository and are publicly available from <https://hbiostat.org/data/gusto.rda>.

## **Supporting information**

All R[27] code used for the illustrations on the GUSTO data and corresponding output is provided as supporting information.

## References

- [1] K. G. Moons, D. G. Altman, J. B. Reitsma, *et al.*, “Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (tripod): Explanation and elaboration,” *Annals of internal medicine*, vol. 162, no. 1, W1–W73, 2015.
- [2] L. J. Trevena, B. J. Zikmund-Fisher, A. Edwards, *et al.*, “Presenting quantitative information about decision outcomes: A risk communication primer for patient decision aid developers,” *BMC medical informatics and decision making*, vol. 13, no. 2, pp. 1–15, 2013.
- [3] A. Fagerlin, B. J. Zikmund-Fisher, and P. A. Ubel, “Helping patients decide: Ten steps to better risk communication,” *Journal of the National Cancer Institute*, vol. 103, no. 19, pp. 1436–1443, 2011.
- [4] M. W. Kattan, J. A. Eastham, T. M. Wheeler, *et al.*, “Counseling men with prostate cancer: A nomogram for predicting the presence of small, moderately differentiated, confined tumors,” *The Journal of urology*, vol. 170, no. 5, pp. 1792–1797, 2003.
- [5] I. M. Lipkus, W. M. P. Klein, and B. K. Rimer, “Communicating Breast Cancer Risks to Women Using Different Formats1,” *Cancer Epidemiology, Biomarkers & Prevention*, vol. 10, no. 8, pp. 895–898, Aug. 2001, ISSN: 1055-9965. eprint: <https://aacrjournals.org/cebpa/article-pdf/10/8/895/1938598/ce0801000895.pdf>.
- [6] P. K. Han, W. M. Klein, T. Lehman, B. Killam, H. Massett, and A. N. Freedman, “Communication of uncertainty regarding individualized cancer risk estimates: Effects and influential factors,” *Medical Decision Making*, vol. 31, no. 2, pp. 354–366, 2011.
- [7] M. W. Kattan, “Doc, what are my chances? a conversation about prognostic uncertainty,” *European urology*, vol. 2, no. 59, p. 224, 2011.
- [8] D. Spiegelhalter, “Risk and uncertainty communication,” *Annual Review of Statistics and Its Application*, vol. 4, no. 1, pp. 31–60, 2017. DOI: 10.1146/annurev-statistics-010814-020148.
- [9] E. W. Steyerberg, A. J. Vickers, N. R. Cook, *et al.*, “Assessing the performance of prediction models: A framework for some traditional and novel measures,” *Epidemiology (Cambridge, Mass.)*, vol. 21, no. 1, p. 128, 2010.
- [10] P. MacCullagh, *Generalized linear models* (Monographs on statistics and applied probability 842820124), eng. London [etc.]: Chapman and Hall, 1983, ISBN: 0412238500.
- [11] GUSTO Investigators, “An international randomized trial comparing four thrombolytic strategies for acute myocardial infarction,” *New England Journal of Medicine*, vol. 329, no. 10, pp. 673–682, 1993.
- [12] E. W. Steyerberg, *Clinical Prediction Models: A Practical Approach to Development, Validation, and Updating* (Statistics for biology and health), eng. Cham: Springer International Publishing AG, 2019, ISBN: 3030163989.
- [13] E. W. Steyerberg, G. J. Borsboom, H. C. van Houwelingen, M. J. Eijkemans, and J. D. F. Habbema, “Validation and updating of predictive logistic regression models: A study on sample size and shrinkage,” *Statistics in medicine*, vol. 23, no. 16, pp. 2567–2586, 2004.
- [14] E. W. Steyerberg, F. E. Harrell Jr, G. J. Borsboom, M. Eijkemans, Y. Vergouwe, and J. D. F. Habbema, “Internal validation of predictive models: Efficiency of some procedures for logistic regression analysis,” *Journal of clinical epidemiology*, vol. 54, no. 8, pp. 774–781, 2001.

- [15] E. W. Steyerberg, M. J. Eijkemans, F. E. Harrell Jr, and J. D. F. Habbema, “Prognostic modelling with logistic regression analysis: A comparison of selection and estimation methods in small data sets,” *Statistics in medicine*, vol. 19, no. 8, pp. 1059–1079, 2000.
- [16] M. Ennis, G. Hinton, D. Naylor, M. Revow, and R. Tibshirani, “A comparison of statistical learning methods on the gusto database,” *Statistics in medicine*, vol. 17, no. 21, pp. 2501–2508, 1998.
- [17] R. M. Califf, L. H. Woodlief, F. E. Harrell Jr, *et al.*, “Selection of thrombolytic therapy for individual patients: Development of a clinical model,” *American heart journal*, vol. 133, no. 6, pp. 630–639, 1997.
- [18] K. L. Lee, L. H. Woodlief, E. J. Topol, *et al.*, “Predictors of 30-day mortality in the era of reperfusion for acute myocardial infarction: Results from an international trial of 41 021 patients,” *Circulation*, vol. 91, no. 6, pp. 1659–1668, 1995.
- [19] R. D. Riley, J. Ensor, K. I. Snell, *et al.*, “Calculating the sample size required for developing a clinical prediction model,” *Bmj*, vol. 368, 2020.
- [20] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: With Applications in R* (Springer texts in statistics), eng. New York, NY: Springer, 2021, ISBN: 9781071614174.
- [21] D. Ruppert, “Trimming and winsorization,” in *Wiley StatsRef: Statistics Reference Online*. John Wiley & Sons, Ltd, 2014, ISBN: 9781118445112. DOI: <https://doi.org/10.1002/9781118445112.stat01887>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118445112.stat01887>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118445112.stat01887>.
- [22] J. Kleinberg, J. Ludwig, S. Mullainathan, and A. Rambachan, “Algorithmic fairness,” *AEA Papers and Proceedings*, vol. 108, pp. 22–27, May 2018. DOI: 10.1257/pandp.20181018. [Online]. Available: <https://www.aeaweb.org/articles?id=10.1257/pandp.20181018>.
- [23] A. Birhane, “Algorithmic injustice: A relational ethics approach,” *Patterns*, vol. 2, no. 2, p. 100205, 2021, ISSN: 2666-3899. DOI: <https://doi.org/10.1016/j.patter.2021.100205>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666389921000155>.
- [24] D. Pessach and E. Shmueli, “A review on fairness in machine learning,” *ACM Comput. Surv.*, vol. 55, no. 3, Feb. 2022, ISSN: 0360-0300. DOI: 10.1145/3494672. [Online]. Available: <https://doi.org/10.1145/3494672>.
- [25] High-level Expert Group on Artificial Intelligence (AI HLEG), *Ethics guidelines for trustworthy AI*, European Commission, Brussels, Belgium, Apr. 2019.
- [26] R. D. Riley and G. S. Collins, *Stability of clinical prediction models developed using statistical or machine learning methods*, <https://doi.org/10.48550/arXiv.2211.01061>, 2022. arXiv: 2211.01061.
- [27] R Core Team, *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria, 2022. [Online]. Available: <https://www.R-project.org/>.