Published in final edited form as:

Commun Stat Theory Methods. 2017 ; 46(7): 3290–3296. doi:10.1080/03610926.2013.851241.

Orthogonality of the Mean and Error Distribution in Generalized Linear Models

ALAN HUANG¹ and PAUL J. RATHOUZ²

University of Technology Sydney and University of Wisconsin–Madison

Abstract

We show that the mean-model parameter is always orthogonal to the error distribution in generalized linear models. Thus, the maximum likelihood estimator of the mean-model parameter will be asymptotically efficient regardless of whether the error distribution is known completely, known up to a finite vector of parameters, or left completely unspecified, in which case the likelihood is taken to be an appropriate semiparametric likelihood. Moreover, the maximum likelihood estimator of the mean-model parameter will be asymptotically independent of the maximum likelihood estimator of the error distribution. This generalizes some well-known results for the special cases of normal, gamma and multinomial regression models, and, perhaps more interestingly, suggests that asymptotically efficient estimation and inferences can always be obtained if the error distribution is nonparametrically estimated along with the mean. In contrast, estimation and inferences using misspecified error distributions or variance functions are generally not efficient.

Keywords

regression model; nuisance tangent space; semiparametric model

1 Introduction

It is well-known that in the normal linear regression model,

$$Y_i|X_i \sim X_i^T \beta + \epsilon_i, \quad \epsilon_i \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma^2),$$

the mean-model parameter β is orthogonal to the error variance σ^2 (e.g. Cox & Reid, 1987, Section 3.3). There are two important implications of this. First, the maximum likelihood estimator (MLE) of β is asymptotically efficient regardless of whether σ^2 is known or estimated simultaneously from the data. Second, the MLE of β is independent of the MLE of σ^2 , which is central to deriving the usual t -tests for inferences on β . (Note that orthogonal parameters are only asymptotically independent in general; finite-sample independence is

¹Corresponding Author. Address: 15 Broadway, Ultimo, AUSTRALIA 2007. Phone: +612 9514-2267. Alan.Huang@uts.edu.au.

²Address: K6/446a, Clinical Science Center, 600 Highland Ave, Madison, WI, USA, 53792. Phone: 608-263-1706. rathouz@biostat.wisc.edu.

special to the normal distribution.) When interest lies primarily in the mean-model, the error variance is often deemed a nuisance parameter.

Similar orthogonality results hold for other generalized linear models (GLMs). Specific examples include the gamma regression model, in which the mean-model parameter is orthogonal to a nuisance shape parameter (Cox & Reid, 1987, Section 3.2), and multinomial models for polytomous data, in which the mean-model parameter is orthogonal to a nuisance vector of baseline probability masses (Rathouz & Gao, 2009). Note that the orthogonality property holds for any link function.

In each of the above settings, the error distribution is characterized by a finite vector of nuisance parameters and orthogonality is established by showing that the Fisher information matrix is block-diagonal, with the blocks corresponding to the vector of mean-model parameters and the vector of nuisance parameters. It is usually straightforward to perform these calculations on a case-by-case basis, working with the specific family of distributions under consideration, but a general result for parametric GLMs can be found in Jørgensen & Knudsen (2004).

When we move away from specific parametric families to consider the class of all GLMs, we find that a general orthogonality property, although expected, may not be so easy to establish. This is because such a class constitutes a semiparametric model, and it is no longer feasible to compute and examine the Fisher information matrix in the presence of an infinite-dimensional parameter.

In this note, we show that the mean-model parameter is always orthogonal to the error distribution in GLMs, even when the error distribution is treated as an infinite-dimensional parameter, belonging to the space of all distributions having a Laplace transform in some neighborhood of zero. This class includes, as special cases, the classical normal, Poisson, gamma and multinomial distributions, as well as many interesting and non-standard distributions, such as the generalized Poisson distribution of Wang & Famoye (1997) for overdispersed counts and the class of all exponential dispersion models with constant dispersion (Jørgensen, 1987). We note in Section 2 that this class of distributions is as large as possible for GLMs, so the result here is indeed the most general possible.

That a general orthogonality property should hold is alluded to in Jørgensen & Knudsen (2004, Section 6.2). In that paper, the notion of orthogonality between a finite-dimensional and infinite-dimensional parameter being considered is that orthogonality holds, in the usual Fisher information matrix sense, for every finite-dimensional submodel. In this note, we use a slightly stronger notion of orthogonality, namely, that the score function for the finite-dimensional parameter is orthogonal to the nuisance tangent space of the infinite-dimensional parameter. Recalling that the nuisance tangent space is the closure of all finite-dimensional submodel tangent spaces, we see that the notion of orthogonality here implies the notion considered in Jørgensen & Knudsen (2004).

Our proof proceeds along the following lines. We first use an exponential tilt representation of GLMs, introduced in Rathouz & Gao (2009) and expanded upon in Huang (2013), to index any GLM by just two parameters, namely, a finite-dimensional mean-model parameter

β and an infinite-dimensional error distribution parameter F . The orthogonality of the two parameters is then characterized by the orthogonality of the score function for β to the nuisance tangent space for F . As it turns out, the nuisance tangent space for F is rather difficult to work with directly, but by embedding the model into a larger class of models, we find that the required calculations become particularly simple.

The exponential tilt representation is derived in Section 2 and the general orthogonality property is proven in Section 3. A connection with mean-variance models is outlined in Section 4. A brief discussion of the theoretical and practical implications of our findings is given in Section 5, which concludes the note.

2 An exponential tilt representation of generalized linear models

Recall that a GLM (McCullagh & Nelder, 1989) for independent data pairs $(X_1, Y_1), \dots, (X_n, Y_n)$ is defined by two components. First, there is a conditional mean model for the responses,

$$E(Y_i|X_i) = \mu(X_i^T \beta), \quad (1)$$

where μ is a user-specified inverse-link function and $\beta \in \mathbb{R}^q$ is a vector of unknown regression parameters. Second, the conditional distributions F_i of each response Y_i given covariate X_i are assumed to come from some exponential family. Assuming the distributions F_i have densities dF_i with respect to some dominating measure, the second component can be written in the exponential tilt form

$$dF_i(y) = \exp\{b(X_i; \beta, F) + \theta(X_i; \beta, F)y\} dF(y) \quad (2)$$

for some reference distribution F , where

$$b(X_i; \beta, F) = -\log \int \exp\{\theta(X_i; \beta, F)y\} dF(y) \quad (3)$$

is a normalizing function and, in order to satisfy (1), the tilt $\theta(X_i; \beta, F)$ is implicitly defined as the solution to the mean constraint

$$\mu(X_i^T \beta) = \int y \exp\{b(X_i; \beta, F) + \theta(X_i; \beta, F)y\} dF(y). \quad (4)$$

It is easy to see that the exponential tilt representation (1)–(4) encompasses all classical GLMs. For example, normal, Poisson and gamma regression models can be recovered by choosing dF to be a Gaussian, Poisson or gamma kernel, respectively. The main advantage of this representation is that it naturally allows for the reference distribution F to be considered as an infinite-dimensional nuisance parameter, along with the finite-dimensional

parameter β , in the model. It is this novel representation that allows us to conveniently characterize any GLM using just the two parameters β and F .

As with any GLM, the reference distribution F is required to have a Laplace transform in some neighborhood of the origin so that the cumulant generating function (3) is well-defined. Thus, the parameter space for F is the class of all distributions that have a Laplace transform in some neighborhood of the origin. Note that this class of distribution functions is as large as it can be for GLMs, because any distribution outside this class cannot be used to generate a valid model.

The exponential tilt representation (1)–(4) was first introduced in Rathouz & Gao (2009). In that paper, the representation is used to derive a useful alternative parametrization of the multinomial regression model for polytomous responses. The representation is also used in Huang (2013) to motivate a semiparametric extension of GLMs for arbitrary responses.

3 The orthogonality of parameters

In parametric models, orthogonality of parameters can be characterized by the Fisher information matrix being block-diagonal. In semiparametric models however, orthogonality between a finite-dimensional parameter β and an infinite-dimensional parameter F cannot be characterized in this way. Rather, it is characterized through the score function for β being orthogonal to the nuisance tangent space for F . Intuitively speaking, the projection of the score function for β on to the nuisance tangent space is a measure of the loss of information about β due to the presence of the nuisance parameter F —this is zero if and only if the score function is orthogonal to the nuisance tangent space. Note that this general notion of orthogonality reduces to the Fisher information matrix criterion when the nuisance parameter is finite-dimensional.

Now, the loglikelihood function corresponding to model (1)–(4) is $l(\beta, F; X, Y) = \log dF(Y) + b(X; \beta, F) + \theta(X; \beta, F)Y$. Thus, the score function for β is given by

$$S_{\beta, F}(X, Y) := \frac{\partial}{\partial \beta} l(\beta, F) = \frac{\partial}{\partial \beta} b(X; \beta, F) + \frac{\partial}{\partial \beta} \theta(X; \beta, F) Y.$$

Implicit differentiation of the defining equations for b and θ leads to the identities

$$\frac{\partial}{\partial \beta} b(X; \beta, F) = -\mu(X^T \beta) \frac{\partial}{\partial \beta} \theta(X; \beta, F) \quad \text{and} \quad \frac{\partial}{\partial \beta} \theta(X; \beta, F) = \frac{\mu'(X^T \beta)}{V(X; \beta, F)} X,$$

where $V(X; \beta, F) = E_{\beta, F}[(Y - \mu(X^T \beta))^2 | X]$ is the conditional variance of Y given X under parameter value (β, F) . The score function for β therefore reduces to

$$S_{\beta, F}(X, Y) = X \frac{\mu'(X^T \beta)}{V(X; \beta, F)} (Y - \mu(X^T \beta)) \quad (5)$$

which is of the same weighted least-squares form as for a parametric GLM. The difference here is that the variance function $V(X; \beta, F)$ is not known because F is not specified.

The orthogonality between β and F now reduces to the score function (5) being orthogonal to the nuisance tangent space for F . Although it is not hard to derive a score function for F (e.g. Huang, 2013, Section 3.3), it turns out to be rather difficult to compute the nuisance tangent space explicitly. This is also noted in Jørgensen & Knudsen (2004, Section 6.2). We can work around this, however, by embedding model (1)–(4) into a more general class of “semiparametric restricted moment models” (e.g. Tsiatis, 2006, Section 4.5) for which the required calculations are much easier. This class is given by

$$Y = \mu(X, \beta) + \epsilon, \quad (6)$$

where the conditional distribution of E given X is specified only up to the moment condition $E(e|X) = 0$. It is clear that the semiparametric extension (1)–(4) is a subclass of the restricted moment model, with $\mu(X, \beta) = \mu(X^T \beta)$ and $E(e|X) = E(Y - \mu(X^T \beta)|X) = 0$ by construction. The nuisance tangent space for F in the semiparametric model (1)–(4) must therefore be a subspace of the nuisance tangent space for the restricted moment model.

Elementary calculations (see Tsiatis, 2006, pp.81–83) show that the nuisance tangent space for the restricted moment model is given by

$$\Lambda = \left\{ \text{all } q \times 1 \text{ functions } a(X, Y) \text{ such that } E_{\beta, F}[(Y - \mu(X, \beta)) a(X, Y) | X] = 0 \right\}$$

and the projection operator $\Pi_{\beta, F}$ onto this nuisance tangent space is given by

$$\Pi_{\beta, F} s = s - \frac{E_{\beta, F}[(Y - \mu(X, \beta)) s | X]}{V(X; \beta, F)} (Y - \mu(X, \beta)).$$

Applying this operator to the score function (5), with $\mu(X, \beta) = \mu(X^T \beta)$, gives

$$\begin{aligned} \Pi_{\beta, F} S_{\beta, F} &= X \frac{\mu'(X^T \beta)}{V(X; \beta, F)} (Y - \mu(X^T \beta)) \\ &\quad - X \frac{\mu'(X^T \beta)}{V(X; \beta, F)} E_{\beta, F} \left[\frac{(Y - \mu(X^T \beta))^2}{V(X; \beta, F)} | X \right] (Y - \mu(X^T \beta)) \\ &= 0 \quad \text{for all } (\beta, F), \end{aligned}$$

because $V(X; \beta, F) = E_{\beta, F}[(Y - \mu(X^T \beta))^2 | X]$ by definition. Thus, the score function (5) is orthogonal to the nuisance tangent space in the restricted moment model (6) and therefore necessarily orthogonal to the nuisance tangent space in the semiparametric model (1)–(4) also. We summarize as follows:

Proposition 1 (Orthogonality) The mean-model parameter β and the error distribution F in any generalized linear model are orthogonal.

Note that the nuisance tangent space for any parametric model (that is, a model in which F is characterized by a finite number of parameters) is necessarily a subspace of the semi-parametric nuisance tangent space. We therefore have the following corollary:

Corollary 1 (Orthogonality in parametric models) If the error distribution is characterized by a finite vector of nuisance parameters ϕ , then the mean-model parameter β is orthogonal to ϕ .

4 A connection with quasilielihood models

A popular extension of GLMs is the class of quasilielihood (QL) models, also known mean-variance models (e.g. Wedderburn, 1974). These models make the assumption that $E(Y|X) = \mu(X, \beta)$ for some mean function μ and $\text{Var}(Y|X) = v(\mu)$ for some positive variance function v . Such models can be characterized by their quasi-score functions for β ,

$$\frac{\partial \mu}{\partial \beta} \frac{Y - \mu}{v(\mu)}. \quad (7)$$

In classical QL literature, the functional forms of both μ and v are usually specified, although there is growing literature on adaptive estimation in which the variance function is left unspecified and estimated nonparametrically from data (e.g. Dewanji & Zhao, 2002). By comparing score equations (5) and (7), note that GLMs form a subset of QL models.

For models characterized by (7), Jørgensen & Knudsen (2004) showed that the mean-model parameter β is orthogonal to the variance function v whenever the latter can be characterized by a finite number of parameters. A general orthogonality result for arbitrary, infinite-dimensional v remains elusive, however, perhaps because of the fact that not all QL score functions (7) correspond to actual probability models. Indeed, such correspondences are atypical. Nevertheless, there is an interesting connection between QL models with unspecified variance functions and GLMs with unspecified error distributions in a certain asymptotic sense made more precise below.

The connection is based on a rather remarkable, but relatively obscure, result from Hiejima (1997), who showed that GLMs can be considered “dense” in the class of QL models in the following asymptotic sense: for any mean-variance relationship, there exists an exponential family of distributions (i.e. a GLM with some error distribution F) whose score equations for β admit roots that are arbitrarily close to the roots of the corresponding QL score equation, as the sample size increases. Thus, for large enough sample sizes, any adaptive QL model with unspecified variance function v can be approximated arbitrarily well by a GLM with unspecified error distribution F , with the latter possessing the orthogonality property 1. We conjecture that this connection may be the best possible for adaptive QL models, mainly because of the aforementioned fact that QL score functions typically do not correspond to actual probability models.

5 Practical implications

The orthogonality property 1 naturally suggests the idea of estimating the error distribution nonparametrically and simultaneously with the mean-model, leading to a kind of “adaptive GLM”. Indeed, if the joint estimation procedure is based on maximum semiparametric likelihood, then the estimator for β is guaranteed to be asymptotically efficient and asymptotically independent of the estimated error distribution. In other words, both estimation and inferences on β are asymptotically unaffected by having to also estimate the error distribution. In contrast, estimation and inferences in GLMs with misspecified error distributions, or QL models with misspecified variance functions, are generally not efficient.

The idea of jointly estimating the mean and error distribution in GLMs is considered in more detail in Huang (2013). In that paper, it is demonstrated that inferences on β based on profiling out the error distribution F in the likelihood can be more accurate than inferences based on QL methods. Here, we focus our attention on the accuracy of the point estimates of β . The results here complement those found in Huang (2013, Section 6).

In Table 1, we compare the relative root mean-square error of a semiparametric MLE of β (with F unknown) to that of the usual MLE (with F set to the true distribution) from three sets of simulations. Recall that the relative root mean-square error of an estimator $\hat{\beta}$ is defined as the root mean-square error of $\hat{\beta}$ divided by the absolute value $|\beta|$ of the parameter. The simulation settings are based on a leukemia survival dataset from Davison & Hinkley (1997) and a mine injury dataset from Myers et al. (2010), and are described in more detail in Sections 6.1 and 6.2 of Huang (2013). The particular semiparametric estimation approach we use for estimating β is based on empirical likelihood and is described in more detail in Section 4 of Huang (2013).

We see from Table 1 that the relative root mean-square errors of the two estimators are essentially the same, even for moderately small sample sizes. This supports the claim that maximum likelihood estimation of β is asymptotically efficient regardless of whether F is known or completely unknown and estimated nonparametrically from data.

6 Conclusion

In this note, we have shown that orthogonality between the mean-model parameter and the error distribution holds for any GLM, parametric or nonparametric. This confirms, in greatest generality, what is well-known for the special cases of normal, gamma and multinomial regression. The result also has implications for applied statistical work, with our numerical results suggesting that little is lost by treating the error distribution nonparametrically, even in moderately sized problems. (It can also be said that little is gained by knowing the error distribution completely!) Nonparametric estimation of the error distribution can therefore safeguard against biases due to parametric model misspecification, without sacrificing much in terms of efficiency.

Acknowledgments

The authors thank the Associate Editor and an anonymous referee for suggestions that improved the paper. Paul. J. Rathouz was funded by NIH grant R01 HL094786.

References

- Cox DR, Reid N. Parameter orthogonality and approximate conditional inference. *J. Roy. Stat. Soc. Ser. B.* 1987; 49:1–39.
- Davison, AC., Hinkley, DV. *Bootstrap methods and their applications.* Cambridge University Press; Cambridge, UK: 1997.
- Dewanji A, Zhao LP. An optimal estimating equation with unspecified variances. *Sankhyā*. 2002; 64:95–108.
- Hiejima Y. Interpretation of the quasi-likelihood via the tilted exponential family. *J. Japan Statist. Soc.* 1997; 27:157–164.
- Huang A. Joint estimation of the mean and error distribution in generalized linear models. *J. Amer. Stat. Assoc.* 2013 to appear.
- Jørgensen B. Exponential dispersion models. *J. Roy. Stat. Soc. Ser. B.* 1987; 49:127–162. with discussion and a reply by the author.
- Jørgensen B, Knudsen SJ. Parameter orthogonality and adjustment for estimating functions. *Scand. J. Stat.* 2004; 31:93–114.
- McCullagh, P., Nelder, JA. *Generalized linear models.* Chapman & Hall; London, UK: 1989.
- Myers, RH., Montgomery, DC., Vining, GG., Robinson, TJ. *Generalized linear models.* John Wiley & Sons Inc.; Hoboken, USA: 2010.
- Rathouz PJ, Gao LP. Generalized linear models with unspecified reference distribution. *Biostatistics.* 2009; 10:205–218. [PubMed: 18824517]
- Tsiatis, AA. *Semiparametric theory and missing data.* Springer; New York, USA: 2006.
- Wang W, Famoye F. Modeling household fertility decisions with generalized Poisson regression. *J. Population. Econ.* 1997; 10:273–283.
- Wedderburn RWM. Quasi-likelihood functions, generalized linear models, and the Gauss-Newton method. *Biometrika.* 1974; 61:439–447.

Table 1

Relative root mean-square errors of a semiparametric MLE $(\hat{\beta}_{SP})$ and the usual MLE $(\hat{\beta}_{MLE})$ in three simulation settings, based on 5000 replications each.

Data	n	Parameter	Estimator	
			$\hat{\beta}_{SP}$	$\hat{\beta}_{MLE}$
Exponential	33	Intercept	0.199	0.194
		Group effect	0.361	0.354
		Common slope	0.464	0.455
	66	Intercept	0.124	0.122
		Group effect	0.247	0.243
		Common slope	0.300	0.297
Poisson	44	Intercept	0.275	0.271
		Coefficient of X_1	0.610	0.594
		Coefficient of X_2	0.205	0.201
		Coefficient of X_3	0.546	0.533