Logistic Regression Likelihood Ratio Test Analysis and Harnessing Graphics to Explore Safety Data in the Vaccine Adverse Event Report System (VAERS)

Kijoeng Nam * Nicholas C. Henderson † Estelle Russek-Cohen ‡
Emily Jane Woo §

Abstract

In contrast with drugs, vaccines are often given to healthy individuals according to a standardized schedule. Passive surveillance of vaccine safety data (e.g. using VAERS) presents several challenges not unlike those with other passive surveillance systems. In spite of these challenges, passive surveillance is useful in identifying unanticipated adverse events and in particular those that are temporally associated with vaccine administration. Empirical Bayes (Dumouchel (1999)) and likelihood ratio test approaches (Huang et al. (2011)) provide ways of developing presentations for assessing many product by adverse event combinations in the context of passive surveillance. In conjunction with traditional surveillance, these newer tools can identify vaccine-event combinations that might warrant further exploration as potential safety signals. Background rates of adverse events may vary by demographics group (e.g. cardiovascular events in the elderly or intussusception in young infants). Approaches that permit adjustment for potential confounders and effect modifiers, such as age, gender, and concomitant vaccines, may be particularly valuable. Graphical displays that facilitate and delineate these differences may aid medical reviewers in assessing and prioritizing vaccine adverse events for further investigation. We present several graphical approaches for examining VAERS data.

Key Words: Vaccine safety, Adverse event, Passive surveillance

1 Vaccine Safety Surveillance

VAERS is a passive surveillance system that allows patients, physicians and vaccine manufactures to report adverse events (AE) that occur after a vaccine exposure. These data are challenging to analyze because the number of AE reports is known but the number exposed to the vaccine is not. Pediatrics is especially challenging because many children receive more than one vaccine at a doctor's visit and children change rapidly at young ages. In passive surveillance systems there is a reliance on disproportionality methods. Huang et al. (2011) introduced a likelihood ratio test (LRT) method for passive surveillance. We introduce a novel LRT method, denoted LR-LRT, that utilizes logistic regression once strata are defined. For strata based on age, we use a method that only allows adjacent ages to be combined into a small number of age classes. We compare both methods using a real data set. While our initial example involves adjusting for age and gender, this method could be expanded to analyze vaccine-vaccine interactions. Graphics highlight the applicability of our method.

2 Methods of Disproportionality Analysis

Consider a large frequency table formed by vaccines, AE combinations, and other sources of information. To be more specific, suppose that AE (having I levels, $i = 1, \ldots, I$)

*Division of Biostatistics, Office of Biostatistics and Epidemiology, CBER, FDA
†Department of Statistics, University of Wisconsin - Madison
‡Division of Biostatistics, Office of Biostatistics and Epidemiology, CBER, FDA
§Division of Epidemiology, Office of Biostatistics and Epidemiology, CBER, FDA

and vaccine (having J levels, $j = 1, \ldots, I$) are two variables of interest. In the database n_{ij} represents the number of reports for the i^{th} AE and j^{th} vaccine. The total number of reports for the i^{th} AE is $n_{i.}$, and the total number of reports for the j^{th} vaccine is $n_{.j}$. The total number of reports is $n_{..}$. The reports for the i^{th} AE and the j^{th} vaccine may be summarized in a 2×2 contingency table as in table 1.

Table 1: 2×2 contingency table of distribution of the number of reports

	j^{th} Vaccine	Other Vaccines	Total
i^{th} AE	n_{ij}	$n_{i.} - n_{ij}$	$n_{i.}$
Other AEs	$n_{.j} - n_{ij}$	$n.. - n_{i.} - n_{.j} + n_{ij}$	$n_{..} - n_{i.}$
Total	$n_{.j}$	$n_{..} - n_{.j}$	$n_{..}$

In this section, we review several methods of disproportionality analysis including the proportional reporting ratio (PRR) (Evans et al. (2001)), the likelihood ratio test (LRT) method (Huang et al. (2011)), the multi-item gamma Poisson shrinker method (MGPS) (Dumouchel (1999)), and the regression-adjusted GPS algorithm (RGPS) (DuMouchel and Harpaz (2012)).

2.1 The Proportional Reporting Ratio

For a fixed i^{th} AE, the estimated proportional reporting ratio (PRR) for the j^{th} vaccine is defined as

$$\hat{PRR}_{ij} = \frac{n_{ij}/n_{.j}}{(n_{i.} - n_{ij})/(n_{..} - n_{.j})}. \tag{1}$$

An approximate 95% confidence interval for the proportional reporting ratio is

$$\exp\left\{ln(\hat{PRR}_{ij}) \pm 1.96\sqrt{\frac{1}{n_{ij}} - \frac{1}{n_{.j}} + \frac{1}{n_{i.} - n_{ij}} - \frac{1}{n_{..} - n_{.j}}}\right\}. \tag{2}$$

It has been suggested to report a particular vaccine as being a signal of disproportional reporting if the lower bound of the above confidence interval (i.e. PRR025) is greater than one (Evans et al. (2001)).

2.2 A likelihood Ratio Test Based (LRT) Method

The Likelihood Ratio Test (LRT) method for signal detection was proposed in Huang et al. (2011). This method works as follows: for each fixed row i (i.e the i^{th} AE) the LRT method assumes that $n_{ij} \sim$ Poisson$(n_{.j}p_{ij})$ and $n_{i.} - n_{ij} \sim$ Poisson$((n_{..} - n_{.j}) \times q_{ij})$, where $j = 1, \ldots, J$. Here, p_{ij} is the reporting rate for the i^{th} AE and the j^{th} vaccine, and q_{ij} is the reporting rate for the i^{th} AE and for all other vaccines except for the j^{th} vaccine. They considered testing the null hypothesis $H_0 : p_{ij} = q_{ij} = p_{i0}$ for all columns (vaccines), $j = 1, \ldots, J$, versus the alternative hypothesis $H_a : p_{ij} > q_{ij}$ for at least one column (i.e. the j^{th} vaccine). This one-sided alternative hypothesis is due to the fact that in vaccine safety signal detection, the vaccine-AE combinations with higher reporting rates are of more interest.

To develop the likelihood ratio procedure, the authors in Huang et al. (2011) first consider the two-sided alternative $H_a : p_{ij} \neq q_{ij}$. The maximum likelihood estimates of p_{ij} and q_{ij} are given by

$$\hat{p}_{ij} = \frac{n_{ij}}{n_{.j}} \quad \text{and} \quad \hat{q}_{ij} = \frac{n_{i.} - n_{ij}}{n_{..} - n_{.j}}. \tag{3}$$

With the i^{th} AE fixed, the likelihood ratio test statistic corresponding to the hypothesis test $H_0 : p_{ij} = q_{ij} = p_{i0}$ vs. $H_a : p_{ij} > q_{ij}$ is given by

$$LR_{ij} = \frac{\max_{H_a} L(p_{ij}, q_{ij})}{\max_{H_0} L(p_{ij}, q_{ij})} = \frac{L_a(\hat{p}_{ij}, \hat{q}_{ij})}{L_0(\hat{p}_{i0}, \hat{p}_{i0})} = \left(\frac{n_{ij}}{E_{ij}}\right)^{n_{ij}} \left(\frac{n_{i.} - n_{ij}}{n_{i.} - E_{ij}}\right)^{n_{i.} - n_{ij}}, \quad (4)$$

where $E_{ij} = n_{i.}n_{.j}/n_{..}$ is the expected number of reports for i^{th} AE and j^{th} vaccine under H_0. With the i^{th} AE fixed, the desired test statistic is the maximum of the log-likelihood ratios $\log(\text{LR}_{ij})$ over the J columns:

$$\text{MaxLogLR}_i = \max_{1 \le j \le J} \log(\text{LR}_{ij}). \quad (5)$$

The distribution of MaxLogLR_i under H_0 is not tractable, and the authors in Huang et al. (2011) suggest a Monte-Carlo approach for obtaining p-values associated with the test $H_0 : p_{ij} = q_{ij}$ vs. $H_a : p_{ij} \neq q_{ij}$.

2.3 The Multi-item Gamma Poisson Shrinker (MGPS) Method and the Regression-Adjusted GPS Algorithm (RGPS)

The multi-item gamma Poisson shrinker (MGPS) method (Dumouchel (1999)) is an empirical Bayes approach that compares the ratio of observed-to-expected reporting events, relative reporting rate $R_{ij} = n_{ij}/E_{ij}$. The MGPS assumes that the distribution of the number of events in the $(i, j)^{th}$ cell is given by $n_{ij}|\lambda_{ij} \sim \text{Poisson}(\lambda_{ij}E_{ij})$, where $E_{ij} = n_{i.}n_{.j}/n_{..}$ is the expected number of counts in the $(i, j)^{th}$ cell, and the λ_{ij}'s are the relative reporting rates. The λ_{ij}'s are assumed to be independent random variables with a distribution that is assumed to be a mixture of two Gamma distributions. From this model, one may directly obtain the posterior distribution of each λ_{ij} and compute the posterior .05-percentile ($EB05_{ij}$) of λ_{ij}. Often, vaccine j is reported as being associated with the i^{th} AE if $EB05_{ij} > 2$ (Szarfman et al. (2002)).

DuMouchel and Harpaz (2012) developed a hybrid of Extended Logistic Regression (ELR) and the Multi-item Gamma Poisson Shrinker (MGPS) which they labeled as the Regression-Adjusted GPS Algorithm (RGPS). This method estimates the relative reporting ratios for each vaccine in three steps.

In the first first step, all covariates are converted into categorical variables and then further grouped into various strata so that the event rate within each stratum is similar. This allows them to use a extended logistic regression model to estimate each strata and vaccine effect. Using the results of this logistic regression fit, they compute covariate-adjusted E_{ij} for each vaccine. Then using these $\{E_{ij}\}$, they use a Poisson-Gamma model to generate shrinkage estimates of the relative reporting ratios (EBRRR)

$$EBRRR_{ij} = \frac{n_{ij} + \gamma}{E_{ij} + \delta}, \quad (6)$$

where (γ, δ) are the hyperparameters in the Gamma prior distribution.

3 Logistic Regression-Adjusted LRT Algorithm (LR-LRT)

This section describes a new algorithm that is a hybrid of Logistic Regression (LR) and the LRT method. We call this new algorithm the Logistic Regression-Adjusted LRT Algorithm (LR-LRT). The LR-LRT method assumes that the data are assembled into a series of reports with each report containing an adverse event-vaccine combination. The LR-LRT method processes one response (AE) at a time, and the computation of each response is independent of other responses so that several responses may be processed in parallel.

3.1 Selecting Stratum Groups for a given AE

The LR-LRT method allows for extra covariates such as age, gender, or concomitant vaccines. Although it is possible to use the LR-LRT method with continuous covariates, we have usually chosen to discretize these variables. With discretization, all the covariates are categorical, and we can then think of each category combination as representing a separate strata. For the subsequent parts of the LR-LRT procedure, we only consider the strata labels to be the covariates. For example, if we had the covariates (gender, age) with age discretized into three groups, then we would have the following six strata: (male,female)$\times$(6-12 month, 12-18 month, 18-24 month).

We would like to select the stratum groups based on covariates such as age and gender. For gender, it is straightforward to define the strata (female, male). However, for age, the correct stratification is not as obvious since the age category cut-offs are somewhat arbitrary. Another difficulty is that we can not just group similar ages together while ignoring the fact that each age group should be an interval. Although there are many reasonable ways to stratify by age, our approach seeks to create the age strata so that the relative reporting rate is similar within each stratum. This is discussed in Section 3.2.

3.2 Selecting Age Groups

Let $\{a_1, \ldots, a_T\}$ denote the possible age values and suppose that we want to discretize the age variable into $K+1$ age categories. Choosing the age strata so that each category is an interval may be formulated as choosing K cut-points $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)$ with the corresponding age strata $(S_1(\boldsymbol{\theta}), \ldots, S_{K+1}(\boldsymbol{\theta}))$ being determined by

$$S_1(\boldsymbol{\theta}) = \{a_k : a_k \leq \theta_1\}, \quad S_i(\boldsymbol{\theta}) = \{a_k : \theta_{i-1} < a_k \leq \theta_i\}, \quad S_{K+1}(\boldsymbol{\theta}) = \{a_k : \theta_K < a_k\}.$$

When selecting $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)$, we assume that each $\theta_i = a_k$ for some k, and that $a_1 \leq \theta_1 < \theta_2 < \ldots, \theta_K < a_T$.

Let $\hat{\boldsymbol{R}}_c = (\hat{R}_{c,1}, \ldots, \hat{R}_{c,n_c})$ be the estimated relative reporting rate for those in strata $S_c(\boldsymbol{\theta})$, for $c \in \{1, \ldots, K+1\}$. If we assume that $\hat{R}_{c,i} \sim N(\mu_c, \sigma^2)$, then the log-likelihood function is

$$\begin{aligned} \log L(\hat{R}_1, \ldots, \hat{R}_{K+1} | \boldsymbol{\theta}, \mu, \sigma^2) &= \sum_{c=1}^{K+1} \sum_{i=1}^{n_c} \log(p(\hat{R}_{c,i}; \mu_c, \sigma^2)) \\ &= \frac{-(K+1)n_c \log(2\pi\sigma^2)}{2} - \frac{1}{2\sigma^2} \sum_{c=1}^{K+1} \sum_{i=1}^{n_c} (\hat{R}_{c,i} - \mu_c)^2. \end{aligned}$$

Given a choice of cut-points $\boldsymbol{\theta}$, estimation of μ_c and σ^2 is straightforward (i.e., $\hat{\mu}_c = \bar{R}_c = \frac{1}{n_c} \sum_{i=1}^{n_c} R_{c,i}$ is the sample mean in the stratum c). So, given $\boldsymbol{\theta}$, let $\hat{\mu}_S$ and $\hat{\sigma}_S^2$ denote the corresponding MLE's and define $LL(S(\boldsymbol{\theta}))$ as

$$\begin{aligned} LL(S(\boldsymbol{\theta})) &= \log L(R | \theta, \hat{\mu}_S, \hat{\sigma}_S^2) \\ &= -\frac{1}{2}(K+1)n_c \log(2\pi\hat{\sigma}^2) - \frac{1}{2\hat{\sigma}_S^2} \sum_{c=1}^{K+1} \sum_{i=1}^{n_c} (\hat{R}_{c,i} - \bar{R}_c)^2. \end{aligned} \tag{7}$$

Since $\hat{\sigma}_S$ is the same for each stratum, maximizing $LL(S(\boldsymbol{\theta}))$ just means minimizing the within-strata sum of squares (WSSS),

$$\text{WSSS} = \sum_{c=1}^{K+1} \sum_{i=1}^{n_c} (\hat{R}_{c,i} - \bar{R}_c)^2. \tag{8}$$

Computation of $LL(S(\boldsymbol{\theta}))$ for every possible choice of $(\theta_1, \ldots, \theta_K)$ is possible if K is relatively small, and in these cases we let $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \ldots, \hat{\theta}_K)$ denote the cut-points which maximize $LL(S(\boldsymbol{\theta}))$.

3.3 Fitting the Logistic Regression

Let $Y_{jk}^i \in \{0, 1\}$ be an indicator of the i^{th} AE in the k^{th} report for vaccine j, where $Y_{jk}^i = 1$ if the k^{th} report for vaccine j contains the i^{th} AE and equals 0 otherwise. Let $g(k, j)$ be the stratum associated with the k^{th} report for vaccine j.

We fit a logistic regression assuming the calculated strata are the covariates for the logistic regression model. This model assumes that under the null hypothesis of no differences in the vaccine effects, $Y_{jk}^i \sim \text{Bernoulli}(p_{jk}^{i,0})$ where $p_{jk}^{i,0}$ is given by

$$p_{jk}^{i,0} = logit^{-1}\Big(\beta_{0,g(k,j)}\Big) = \frac{\exp(\beta_{0,g(k,j)})}{1 + \exp(\beta_{0,g(k,j)})}, \tag{9}$$

where $logit^{-1}(x) = e^x/(1 + e^x)$ and $\beta_{0,g(k,j)}$ is the intercept term for reports in the $g(k, j)^{th}$ stratum. The expression for $p_{jk}^{i,0}$ in (9) is equivalent to writing

$$\log\Big(\frac{p_{jk}^{i,0}}{1 - p_{jk}^{i,0}}\Big) = \beta_{0,g(k,j)}. \tag{10}$$

In general, when the null hypothesis does not hold, we assume that $Y_{jk}^i \sim \text{Bernoulli}(p_{jk}^{i,A})$ where

$$\log\Big(\frac{p_{jk}^{i,A}}{1 - p_{jk}^{i,A}}\Big) = \beta_{0,g(k,j)} + \beta_j. \tag{11}$$

Under H_0, the log-likelihood function is

$$\begin{aligned} \log L(\boldsymbol{\beta}|Y) &= \sum_{j=1}^{J}\sum_{k=1}^{n_{\cdot j}} Y_{jk}^i \log(p_{jk}^{i,0}) + \sum_{j=1}^{J}\sum_{k=1}^{n_{\cdot j}}(1 - Y_{jk}^i)\log(1 - p_{jk}^{i,0}) \\ &= \sum_{j=1}^{J}\sum_{k=1}^{n_{\cdot j}} Y_{jk}^i \beta_{0,g(k,j)} - \sum_{j=1}^{J}\sum_{k=1}^{n_{\cdot j}} \log\Big(1 + \exp(\beta_{0,g(k,j)})\Big). \end{aligned} \tag{12}$$

The estimated strata coefficients $(\hat{\beta}_{0,1}, \ldots, \hat{\beta}_{0,S})$ are computed by maximizing $\log L(\boldsymbol{\beta}|Y)$.

3.4 Computing Expected Counts E_{ij} for Every Vaccine

The next step is to compute a set of baseline or expected counts E_{ij} for every vaccine available, including those that were used as strata predictors in the LR-LRT model of Section 3.3. We define E_{ij} to be the expected value of n_{ij} under the null hypothesis

$$E_{ij} = E_{H_0}(n_{ij}) = \sum_{k=1}^{n_{\cdot j}} logit^{-1}\Big(\beta_{0,g(k,j)}\Big) = \sum_{k=1}^{n_{\cdot j}} \frac{\exp(\beta_{0,g(k,j)})}{1 + \exp(\beta_{0,g(k,j)})}, \tag{13}$$

and hence the estimated E_{ij} are defined as

$$\hat{E}_{ij} = \sum_{k=1}^{n_{\cdot j}} \frac{\exp(\hat{\beta}_{0,g(k,j)})}{1 + \exp(\hat{\beta}_{0,g(k,j)})} = \sum_{s=1}^{S} n_{\cdot j}^s \frac{\exp(\hat{\beta}_{0,s})}{1 + \exp(\hat{\beta}_{0,s})}, \tag{14}$$

where $n_{\cdot j}^s$ is the number of reports for vaccine j and strata s. Due to the consistency of $\hat{\beta}_{0,s}$ as an estimate of $\beta_{0,s}$, it is not difficult to see that $\hat{E}_{ij}$ is asymptotically equivalent to E_{ij} in the sense that $(\hat{E}_{ij} - E_{ij})/n_{\cdot\cdot}$ converges to zero in probability as $n_{\cdot\cdot}$ goes to infinity.

3.5 Employing the Novel LRT method by using E_{ij}

Because $n_{ij} = \sum_{k=1}^{n_{.j}} Y_{jk}^i$ and $Y_{jk}^i \sim$ Bernoulli$(p_{jk}^{i,A})$, n_{ij} is the sum of independent Bernoulli random variables with different probabilities (commonly referred to as Poisson binomial) The distribution of n_{ij} is closely approximated by assuming that $n_{ij} \sim$ Poisson$(\sum_k p_{jk}^{i,A})$ (See Hodges et al. (1960)). Now, note that

$$\sum_k p_{jk}^{i,A} = \sum_{k=1}^{n_{.j}} logit^{-1}\Big(\beta_{0,g(k,j)} + \beta_j\Big) = \Big[\frac{\sum_{k=1}^{n_{.j}} logit^{-1}\Big(\beta_{0,g(k,j)} + \beta_j\Big)}{\sum_{k=1}^{n_{.j}} logit^{-1}\Big(\beta_{0,g(k,j)}\Big)}\Big] \times E_{ij}. \quad (15)$$

So, if we define λ_{ij} to be

$$\lambda_{ij} = \Big[\frac{\sum_{k=1}^{n_{.j}} logit^{-1}\Big(\beta_{0,g(k,j)} + \beta_j\Big)}{\sum_{k=1}^{n_{.j}} logit^{-1}\Big(\beta_{0,g(k,j)}\Big)}\Big], \quad (16)$$

then we have the reasonable approximation $n_{ij} \sim$ Poisson$(\lambda_{ij} E_{ij})$ with $\lambda_{ij} = 1$ under H_0.

The two-sided test is

$$H_0 : \lambda_{ij} = 1 \quad \text{vs.} \quad H_A : \lambda_{ij} \neq 1 \quad (17)$$

and with the Poisson approximation the log-likelihood function is

$$l(\lambda_{ij}) = \log L(\lambda_{ij}) = -E_{ij}\lambda_{ij} + n_{ij}\log(\lambda_{ij}) + n_{ij}\log(n_{ij}) - \log(n_{ij}!). \quad (18)$$

The MLE of λ_{ij} is $\hat{\lambda}_{ij} = n_{ij}/E_{ij}$ and the likelihood ratio statistic is then

$$LR_{ij} = \frac{\max_{H_A} L(\lambda)}{\max_{H_0} L(\lambda)} = \frac{L(\hat{\lambda}_{ij})}{L(1)} = \frac{e^{-\hat{\lambda}_{ij}E_{ij}}(\hat{\lambda}_{ij}E_{ij})^{n_{ij}}}{e^{-E_{ij}}E_{ij}^{n_{ij}}} = e^{E_{ij}-n_{ij}}\Big(\frac{n_{ij}}{E_{ij}}\Big)^{n_{ij}}, \quad (19)$$

with the logarithm of the likelihood ratio statistic being

$$\log LR_{ij} = E_{ij} - n_{ij} + n_{ij}\Big(\log(n_{ij}) - \log(E_{ij})\Big). \quad (20)$$

The MLR test statistic is

$$MLR_i = \max_j \log LR_{ij}, \quad (21)$$

where we plug-in $\hat{E}_{ij}$ in place of E_{ij} to compute the MLR_i statistic. Since we are really interested in the one-sided test $H_0 : \lambda_{ij} \leq 1$ vs. $H_A : \lambda_{ij} > 1$, we instead use the following for the likelihood ratio

$$\tilde{LR}_{ij} = LR_{ij}\mathbf{1}\{\hat{\lambda}_{ij} > 1\} = LR_{ij}\mathbf{1}\{n_{ij} > E_{ij}\}. \quad (22)$$

3.6 Hypothesis Testing

The distribution of the MLR statistic under H_0, is not analytically tractable so it is obtained through Monte Carlo simulation (Huang et al. (2011)). Simulating the null distribution of the MLR statistic utilizes the fact that under H_0

$$n_{ij} \sim \text{Poisson}(\hat{E}_{ij}), \quad (23)$$

which implies that given the margin total $n_{i.}$,

$$(n_{i1}, \ldots, n_{iJ})|n_{i.} \sim \text{Multinomial}\Big(n_{i.}; \frac{\hat{E}_{i1}}{\hat{E}_{i.}}, \ldots, \frac{\hat{E}_{iJ}}{\hat{E}_{i.}}\Big) \quad \text{under } H_0. \quad (24)$$

By using (24), $10,000$ MLR values are calculated (with 9999 being simulated MLR values and one being the observed MLR) and then reject H_0 at the $\alpha = 0.05$ level if the observed MLR exceeds the 95^{th} percentile of the $10,000$ MLR values. We then rank the $\log(LR_{ij})$ from the observed dataset from largest to smallest. For the j^{th} vaccine, calculate the corresponding p-value $p_j = 1 - R_j/(1 + 9999)$, where R_j is the rank of $\log(LR_{ij})$ among all the $10,000$ MLR. If $p_j < 0.05$, then we consider vaccine j to be a signal. Because the p-values are determined by the null distribution of the maximum of the log-likelihood ratio statistics, this procedure naturally controls for multiplicity.

4 An Example

As noted in our introduction, passive surveillance systems such as VAERS are used to identify potential signals that may warrant further investigation. The data for our example were VAERS reports from $2005 - 2010$ in the pediatric population from 6 months to 2 years. We elected to focus on a single adverse event, namely febrile convulsions, because **?**) suggested an association with a particular flu vaccine with febrile convulsions and this event, while rare, tends to be more common in a young pediatric population. We restricted our study to children 6 months and older since the flu vaccines being evaluated in this paper are currently not licensed for children under 6 months, and we wanted to look at vaccines with a comparable age and gender distribution.

We only consider age and gender as covariates at this stage. We caution that our method is still under development and our example does not fully account for all possible confounders. For these data, the age stratification that maximizes $LL(S(\boldsymbol{\theta}))$ in (7) and assuming three age groups as discussed in Section 3.2 is shown in Figure 1. This figure shows that those with ages 0.5 or 0.6 were assigned to age stratum 1, those aged 0.7, 0.8, or 0.9 were assigned to age stratum 2 and those with aged $1.0, 1.1, \ldots, 2.0$ were assigned to age stratum 3.

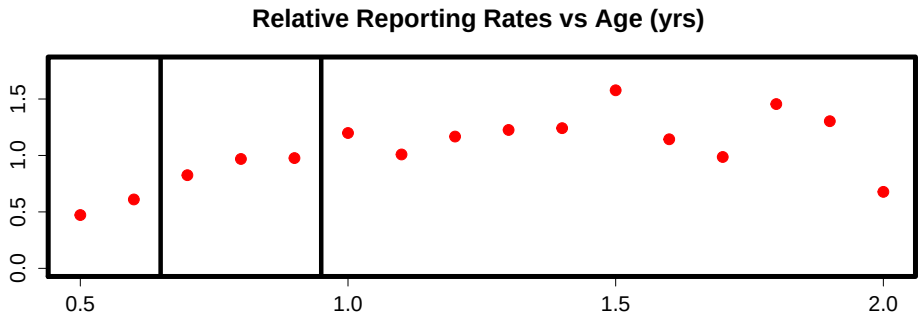

Figure 1: Best Age Stratification assuming Three Age Groups

We computed the PRR scores described in Section 5 and applied both the LRT method and the LR-LRT method to the VAERS dataset. For the PRR method, we calculated the lower confidence interval of PRR (i.e., PRR025). For the LRT method, we calculated the log-likelihood ratios (LLR) and the p-values by the method of Huang et al. (2011), and we calculated the log-likelihood ratios (LR-LLR) and the p-values for our new LR-LRT procedure using the steps discussed in Section 3.6. The results are presented in Table 2. The term "FOREIGN" is not a particular influenza brand or vaccine brand, but a collection of various influenza vaccines. We note the results for LRT and LR-LRT are similar for

INFLUENZA BRAND A while INFLUENZA FOREIGN and INFLUENZA BRAND B are significant only for the LRT method. PRR captures 13 vaccine signals including three of the similar captured by the LRT method. The reason that PRR captures so many signals is most likely due to the fact that just reporting the vaccines whose PRR025 is greater than one does not account in any way for multiplicity. In contrast, both the LRT and the LR-LRT method address the multiplicity problem by using adjusted p-values that control the overall Type I error level.

Table 2: The association of Febrile Convulsion (FC) with various vaccines in VAERS. n_{ij} is the number of reports with vaccine j and adverse event i. $n_{.j}$ is the total number of reports for vaccine j. Only those analyses with p-values $< .05$ are reported. In total, there are 76 vaccines among the 141 vaccines that have Febrile Convulsion as an AE.

Vaccine	$n_{.j}$	n_{ij}	PRR(PRR025)	LLR(p-value)	LR-LLR(p-value)
INFLUENZA BRAND A	8391	182	1.43 (1.29)	10.54 (0.0003)	10.35 (0.0197)
INFLUENZA FOREIGN	563	22	2.62 (2.20)	6.27 (0.0107)	6.12 (0.2580)
INFLUENZA BRAND B	244	12	3.34 (2.76)	5.76 (0.0201)	5.62 (0.3273)

5 Harnessing Graphics

A common use of passive surveillance is to investigate many vaccine-AE pairs in an exploratory fashion and it was logical to consider some means of displaying the results graphically. We present two possible graphs though many more could be considered. The first of which is a volcano plot as shown in Figure 2 and we borrow this idea from Zink et al. (2013). The x-axis is a metric reflecting the magnitude of the effect, namely the relative reporting rates (See Section 2.3). The y-axis reports $-\log(p)$ with p being the calculated p-value. The p-value is typically not adjusted for multiple comparisons. Smaller p-values correspond to a greater $-\log(p)$. The size of the bubble reflects the number of reports for that vaccine-AE combination.

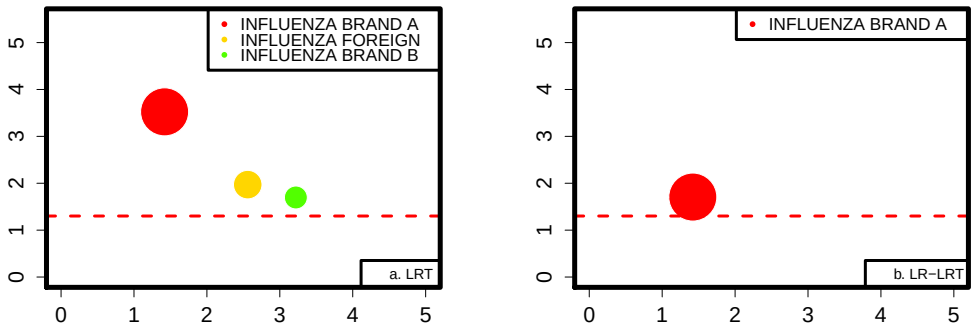

Figure 2: Volcano plots show -log(p-value) vs. the estimated relative reporting ratio for both LRT and LR-LRT. Only the significant vaccines $(p < .05)$ reported for Febrile Convulsion in Table 2 are shown.

The proportional reporting ratio (PRR) (Evans et al. (2001)) plot could provide a comparison of adverse event rates of the same AE across a class of products. We recognize

an important limitation: vaccines that may have been used more often may influence the results, since p-values are sensitive to sample size. Here we only report the three associations that are statistically significant, but many vaccines were examined as noted in Table 2. The right-hand panel of Figure 3 plots $P\hat{R}R_{ij}$ by vaccine along with the associated 95% confidence intervals. Only those highlighted in Table 2 by LRT and LR-LRT are graphed in Table 2.

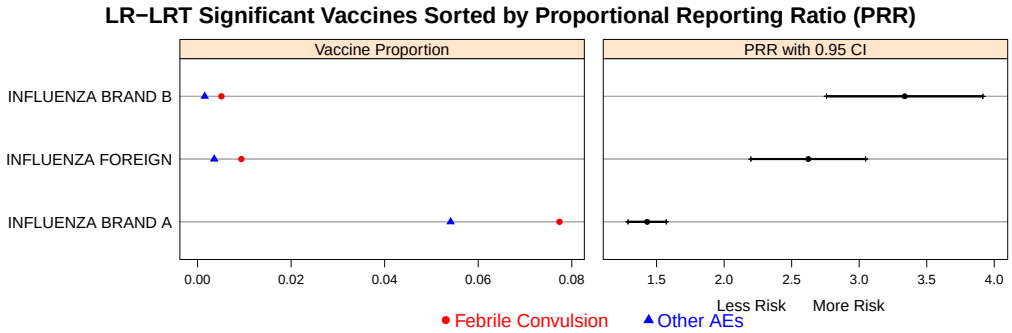

Figure 3: Graph illustrates disproportionality for Febrile Convulsion for a set of vaccines. The PRR plot could be used to compare vaccines within a certain class.

6 Conclusion and Future Work

Our LR-LRT procedure is a novel method for identifying vaccine-AE associations while adjusting for possible confounders in the context of passive surveillance data. We highlighted our method by focusing on the adverse event of febrile convulsion, but passive surveillance signals would need to be further explored using other data. Now we are developing an approach to multiplicity adjustment to facilitate data mining using false discovery rates and are developing additional graphical displays tied to our analysis procedure. Finally, to further evaluate the performance of our procedure, we plan to conduct simulations that examine the LR-LRT method on a variety measures such as Type I error, power, false discovery rate, and sensitivity.

Acknowledgement

We appreciate Dr. Tiwari providing the data in Huang et al. (2011) which enabled us to check our software implementation of their LRT method.

References

Vaccine Adverse Event Reporting System. http://vaers.hhs.gov/index.

Dumouchel, W. (1999). Bayesian Data Mining in Large Frequency Tables, with an Application to the FDA Spontaneous Reporting System. *The American Statistician* **53,** 177–190.

DuMouchel, W. and Harpaz, R. (2012). Regression-Adjusted GPS Algorithm (RGPS). *ORACLE Health Sciences* .

Evans, S. J., Waller, P. C., and Davis, S. (2001). Use of proportional reporting ratios (PRRs) for signal generation from spontaneous adverse drug reaction reports. *Pharmacoepidemiology and drug safety* **10,** 483–486.

Hodges, J., Le Cam, L., et al. (1960). The Poisson Approximation to the Poisson Binomial Distribution. *The Annals of Mathematical Statistics* **31,** 737–740.

Huang, L., Zalkikar, J., and Tiwari, R. C. (2011). A Likelihood Ratio Test Based Method for Signal Detection With Application to FDAs Drug Safety Data. *Journal of the American Statistical Association* **106,** 1230–1241.

Martin, D., Menschik, D., Bryant-Genevier, M., and Ball, R. (2013). Data mining for prospective early detection of safety signals in the Vaccine Adverse Event Reporting System (VAERS): A case study of febrile seizures after a 2010 – 2011 seasonal influenza virus vaccine. *Drug safety* **36,** 547–556.

Szarfman, A., Machado, S. G., and ONeill, R. T. (2002). Use of screening algorithms and computer systems to efficiently signal higher-than-expected combinations of drugs and events in the US FDAs spontaneous reports database. *Drug Safety* **25,** 381–392.

Zink, R. C., Wolfinger, R., and Mann, G. (2013). Summarizing the incidence of adverse events using volcano plots and time intervals. *Clin Trials* **10,** 398–406.