

Length minimization for Poisson confidence procedures

Mark F. Schilling & Bret Holladay

To cite this article: Mark F. Schilling & Bret Holladay (2017) Length minimization for Poisson confidence procedures, Communications in Statistics - Theory and Methods, 46:2, 861-873, DOI: [10.1080/03610926.2015.1006782](https://doi.org/10.1080/03610926.2015.1006782)

To link to this article: <http://dx.doi.org/10.1080/03610926.2015.1006782>



Accepted author version posted online: 25 Feb 2016.
Published online: 25 Feb 2016.



Submit your article to this journal [↗](#)



Article views: 49



View related articles [↗](#)



View Crossmark data [↗](#)

Length minimization for Poisson confidence procedures

Mark F. Schilling^a and Bret Holladay^b

^aDepartment of Mathematics, California State University Northridge, Northridge, CA, USA; ^bDepartment of Statistics and Applied Probability, University of California, Santa Barbara, CA, USA

ABSTRACT

We study Poisson confidence procedures that potentially lead to short confidence intervals, investigating the class of all minimal cardinality procedures. We consider how length minimization should be properly defined, and show that Casella and Robert's (1989) criterion for comparing Poisson confidence procedures leads to a contradiction. We provide an alternative criterion for comparing length performance, identify the unique length optimal minimal cardinality procedure by this criterion, and propose a modification that eliminates an important drawback it possesses. We focus on procedures whose coverage never falls below the nominal level and discuss the case in which the nominal level represents mean coverage.

ARTICLE HISTORY

Received 9 July 2014
Accepted 5 January 2015

KEYWORDS

Confidence interval;
coverage probability; Poisson
mean.

MATHEMATICS SUBJECT

CLASSIFICATION 2010

Primary 62F25

1. Introduction

Interval estimation for the rate parameter λ of a Poisson distribution based on the observed value x of a Poisson count X is an important statistical problem, and numerous solutions have been proposed. See Kabaila and Byrne (2005) for an extensive list. We focus in this paper initially and primarily on methods that produce a *strict level* $1 - \alpha$ confidence procedure C : an infinite set of intervals $\{[l_X, u_X), x = 0, 1, \dots\}$ for which $\inf_{\lambda} P(\lambda \in [l_X, u_X)) \geq 1 - \alpha$. (We use half-open intervals to avoid certain technical difficulties that arise with fully open or fully closed intervals.) The approach we take is based on careful analysis of the *coverage probability function* (*cpf*) of a Poisson procedure, which we describe in depth in Section 2.

Since accuracy is of paramount concern in estimation, it makes sense to focus on confidence procedures that produce intervals that are as short as possible. Specifying what this means in the Poisson case is somewhat problematic however, since a Poisson confidence procedure produces an infinite number of intervals, one for each possible value of x . One definition of shortness is by Kabaila and Byrne (2005), who introduced the “Inability to be Shortened” property. A confidence procedure possesses this attribute if increasing any lower interval endpoint or decreasing any upper interval endpoint causes the above coverage condition to be violated, that is, results in $\inf_{\lambda} P_{\lambda}(\lambda \in [l_X, u_X)) < 1 - \alpha$. Casella and Robert (1989) gave an alternative definition of length optimality which we present in Section 3 after describing several specific Poisson confidence procedures. In Section 4 we show that Casella and Robert's criterion leads to a paradox, and provide our own recommendation for comparing the length performance of Poisson confidence procedures. We then identify the unique length optimal

CONTACT Mark F. Schilling  mark.schilling@csun.edu  Department of Mathematics, California State University Northridge, Northridge, 91330, USA.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/lsta.

© 2017 Taylor & Francis Group, LLC

minimal cardinality procedure with respect to this criterion. As this procedure has a serious drawback, in Section 5 we provide a simple modification that resolves this shortcoming while achieving near optimality with respect to length. Finally, in Section 6 we provide a brief treatment of how to adapt a strict confidence procedure for use as an approximate procedure that allows coverage to fall below $1 - \alpha$ for some λ but has mean coverage equal to $1 - \alpha$. This produces confidence intervals that are somewhat shorter than for a strict procedure.

2. Length optimality: a coverage probability function perspective

There is a one-to-one correspondence between a level $1 - \alpha$ confidence procedure and its *cpf* $CP(\lambda) = P_\lambda(\lambda \in [l_x, u_x])$ viewed as a function of λ (except possibly on a set of measure zero). Analyzing Poisson confidence procedures through the structure of their cpfs is central to understanding their behavior and performance.

A Poisson confidence procedure can be described by its *acceptance sets* $A_\lambda = \{x: l_x \leq \lambda < u_x\}$, $0 < \lambda < \infty$; the confidence set for a given x is $\{\lambda: x \in A_\lambda\}$. Due to the unimodality of the Poisson distribution, the only reasonable acceptance sets comprise a sequence of consecutive values of x . Define the *acceptance curve* associated with A_λ as $AC(a-b) = P_\lambda(X \in A_\lambda)$ considered as a function of λ , where $a = \min\{x \in A_\lambda\}$, $b = \max\{x \in A_\lambda\}$. The cpf of a Poisson confidence procedure comprises a collection of acceptance curve segments, where the sequences of a and b values obtained as λ increases must each be non decreasing in order that the confidence sets generated for each x are intervals (contain no gaps).

Given any confidence level $1 - \alpha \in (0,1)$, let $M(\lambda) = \min\{k: P_\lambda(a \leq X \leq a + k - 1) \geq 1 - \alpha \text{ for some } a \in Z^*\}$, where $Z^* = \{0,1,2, \dots\}$. We call $M(\lambda)$ the minimal cardinality at λ . Choosing for each λ an acceptance curve whose acceptance set is of minimal cardinality naturally tends to lead to short confidence intervals. We say that a confidence procedure whose acceptance sets have minimal cardinality for all λ has the *Minimal Cardinality Property*.

Theorem 1. *If confidence procedure C has the Minimal Cardinality Property, then it has the Inability to be Shortened property.*

Proof. Suppose a strict confidence procedure C has the Minimal Cardinality Property. Now define a new confidence procedure C' that is identical to C except that for the confidence interval for a fixed x_0 we increase the lower endpoint by ε ; that is, $l'_x = l_x$ for all $x \neq x_0$ and $u'_x = u_x$ for all x , but $l'_{x_0} = l_{x_0} + \varepsilon$. Then for all $\lambda \in [l_{x_0}, l'_{x_0})$, $CP_{C'}(\lambda) = CP_C(\lambda) - P(X = x_0)$, hence the segment of the cpf of C' corresponding to $\lambda \in [l_{x_0}, l'_{x_0})$ is of cardinality one less than that of C . However, this implies all such λ will be included in one less confidence interval; therefore the coverage of C' must drop below the confidence level since by the Minimal Cardinality Property each λ was already contained in the fewest possible intervals. Similarly, decreasing any upper endpoint also causes the coverage probability to fall below the confidence level. In either case a contradiction to the assumption that C is a strict confidence procedure occurs. Hence, C has the Inability to be Shortened property. \square

In view of Theorem 1, it makes sense to limit attention to procedures that possess the Minimal Cardinality Property. It can be shown that methods based on closed-form formulas (typically based on normal or other approximations) do not satisfy this property; consequently such procedures are length inadmissible as they produce intervals that can be shortened without their cpfs falling below the confidence level.

As an illustration of the key elements involved, consider the cpf of a confidence procedure with level 95% that satisfies the Minimal Cardinality Property. For $\lambda < 20$ this cpf must be




Figure 1 All minimal cardinality curves for a strict 95% confidence procedure that allow the sequences of a and b values obtained as λ increases to be non decreasing.

constructed from portions of the acceptance curves shown in Figure 1, which correspond to all acceptance sets of minimal cardinality for each $\lambda \in [0, 20)$ except those whose use would force either the sequence of a values or the sequence of b values obtained as λ increases to be decreasing. In the leftmost portion of the figure, where $0 \leq \lambda \leq 3.28$, the acceptance curve of minimal cardinality for each λ is unique. Thus in this region the lower confidence limits of a Minimal Cardinality confidence procedure are uniquely determined for $x = 0, 1, \dots, 7$. For $\lambda > 3.28$, there are many regions in which only one curve of minimal cardinality appears. When λ is not within one of those regions, and thus two or more acceptance curves of minimal cardinality are present, there arises the phenomenon Casella (1986) termed *coincidental endpoints*.

Coincidental endpoints occur when $u_x = l_y$ for some pair (x, y) with $x < y$. Note that increasing (decreasing) their common value lengthens (shortens) the confidence interval for x while shortening (lengthening) the confidence interval for y by an equal amount. Coincidental endpoints are generated by the transition of the cpf from one acceptance curve $AC(a-b)$ to another, $AC(a'-b')$, with $a < a'$ and $b < b'$. Such a transition point represents the upper limit of the confidence interval for $x = a$ and the lower limit of the confidence interval for $x = b'$. Consider for example the first such instance for confidence level 95% in which multiple minimal cardinality acceptance curves occur, where $AC(0-7)$ and $AC(1-8)$ overlap for $3.29 < \lambda \leq 3.98$. Any point within this interval can serve as both the upper confidence limit for $x = 0$ and the lower confidence limit for $x = 8$.

For given λ and k let $S_k(\lambda) = \max_a P_\lambda(a \leq X \leq a + k - 1)$. Crow and Gardner (1959, p. 442; see also their Figure 1) proved that for fixed k the function $S_k(\lambda)$ is a non increasing function of λ . It follows that the minimal cardinality $M(\lambda)$ is non increasing, as shown in Figure 2.

We exploit this fact to prove the following.

Theorem 2. Let C be any procedure having the Minimal Cardinality Property. Then all upper endpoints generated by C are coincidental.

Proof. Let $\{a\}$ and $\{b\}$ be the sequences of values for the acceptance curves $\{AC(a-b)\}$ that the cpf of C uses as λ increases. As already mentioned, to avoid gaps these sequences must be non decreasing. Thus an upper endpoint occurs whenever there is an increase in a , while a lower endpoint occurs whenever there is an increase in b . But since C has the Minimal Cardinality Property and $M(\lambda)$ is non decreasing, an increase in a must be accompanied by an increase in b . Thus every upper endpoint is coincidental with at least one lower endpoint. \square




Figure 1 Minimal cardinality values $M(\lambda)$ for a level 95% confidence procedure.

Not all lower endpoints are necessarily coincidental, however. Non coincidental lower endpoints arise when an increase in b is not accompanied by an increase in a , which occurs precisely at those values of λ where $M(\lambda)$ increases. Let λ^* be such a value. Just before λ^* there is a unique acceptance curve $AC(a-b)$ of minimal cardinality (see Figure 1). Since $M(\lambda)$ increases at λ^* , this point establishes the lower endpoint for $x = b + 1$. Thus all confidence procedures satisfying the Minimal Cardinality Property produce identical lower endpoints at these values. Note that $\lim_{\lambda \uparrow \lambda^*} CP(\lambda) = 1 - \alpha$ for such procedures. The various procedures we describe below are distinguished from each other entirely by the values of their coincidental endpoints.

3. Some specific minimal cardinality Poisson confidence procedures

We begin by briefly considering the procedure described by Garwood (1936), although it does not satisfy the Minimal Cardinality Property, because of its prevalence as one of the most common methods for generating Poisson confidence intervals—for example, it is the sole exact method provided by the *StatXact*® software. The Garwood intervals are the Poisson analog to the Clopper–Pearson intervals for the binomial case, obtained by inverting a two-sided test whose rejection region includes no more than $\alpha/2$ probability in each tail of the null distribution. The effect of restricting each tail probability rather than limiting the total of both tails to $\leq \alpha$ makes the resulting confidence procedure highly conservative, as can be seen from its cpf, shown in Figure 3.

Note that the cpf for the 95% Garwood procedure jumps rapidly between different acceptance curves, and typically has values far above the confidence level. The result is that the




Figure 3 Coverage probability function for Garwood's method, $1 - \alpha = 95\%$.




Figure 4 Coverage probability function for the Crow–Gardner method, $1 - \alpha = 95\%$.

Garwood intervals are much wider than those obtainable from a Minimal Cardinality confidence procedure, as will be seen later.

The first idea for a confidence procedure that possesses the Minimal Cardinality Property is due to Sterne (1954) (although his proposal was for the one-sample binomial problem): Enter x values into A_λ in decreasing order of their probabilities until $P(X \in A_\lambda) \geq 1 - \alpha$. Unfortunately, this sometimes results in the $\{a\}$ sequence of the procedure not being monotone non decreasing; this yields confidence sets for certain x that do not comprise a single interval. We discuss this phenomenon further below.

Crow and Gardner (1959) developed the first Minimal Cardinality Poisson confidence procedure that does not produce gaps in its confidence sets. Their approach, expressed in terms of the cpf perspective, is to always use the curve $AC(a-b)$ with the largest values of a and b when multiple curves of minimal cardinality are available. Figure 4 shows the cpf of the Crow–Gardner procedure (CG) for confidence level 95%.

Figure 5 shows the interplay between the cpf and the confidence intervals for the 95% Crow–Gardner procedure. While solving one problem (gaps in the confidence sets), Crow




Figure 5 Relationship of the cpf to the confidence intervals for the 95% CG procedure.

and Gardner’s method creates another: often, the lower or upper confidence limits for consecutive values of x are the same. For example (see Figure 4), the cpf of CG transitions from $AC(0-6)$ directly to $AC(1-8)$; $AC(0-7)$ is never used. The consequence is that the transition point between $AC(0-6)$ and $AC(1-8)$ becomes the lower confidence limit for both $x = 7$ and $x = 8$. In general, equal endpoints for consecutive x ’s occur only at values of λ where the minimal cardinality increases, when CG transitions from the unique minimal cardinality curve $AC(a-b)$ to a curve $AC((a+k)-(b+k+1))$ for $k \geq 1$. This circumstance results in $k+1$ consecutive values of x having the same lower endpoint and, if $k \geq 2$, k consecutive values having the same upper endpoint. Figure 5 shows several such instances ($l_7 = l_8, l_{10} = l_{11}, l_{12} = l_{13}$ and $l_{14} = l_{15}$).

Kabaila and Byrne (2001) proposed an approach that is the opposite of CG. When choosing among curves of minimal cardinality, use the curve $AC(a-b)$ having the *smallest* values of a and b that keep the sequence of $\{a\}$ and $\{b\}$ values monotone non decreasing. Thus Kabaila and Byrne’s procedure (KB) transitions between curves of equal cardinality as late as possible and Crow and Gardner’s transitions as early as possible. KB resolves the problems of both Sterne’s approach (gaps) and Crow and Gardner’s approach (interval endpoints that are not strictly monotone increasing in x). Figure 6 shows the cpf of KB for confidence level 95%.

Of the methods discussed so far that possess the Minimal Cardinality Property, Sterne’s procedure has a particularly attractive feature in that it amounts to using the highest acceptance curve of minimal cardinality at each λ , thereby producing greater coverage than other procedures. As indicated above, though, Sterne’s method occasionally produces a gap in a confidence set. Figure 7 shows a portion of the cpf of Sterne’s procedure at 95% confidence that illustrates a specific instance of this phenomenon.

Notice that the cpf transitions from $AC(23-45)$ to $AC(24-46)$ but then jumps to $AC(23-46)$. The result is that the confidence set for $x = 23$ has a gap corresponding to the interval where the cpf follows $AC(24-46)$.

Schilling and Doi (2014) developed a new confidence procedure for the one-sample binomial problem that resolves the gap problem while maximizing coverage among all confidence procedures having minimal average length. We now describe an analogous approach for the Poisson case.

- (1) For each λ , choose the highest acceptance curve among those having minimal cardinality, except:
- (2) Whenever step (1) gives a curve $AC(a-b)$ that results in a decrease in the $\{a\}$ sequence of the procedure, substitute for $AC(a-b)$ the curve $AC((a+1)-(b+1))$.




Figure 6 Coverage probability function for the Kabaila–Byrne method, $1 - \alpha = 95\%$.




Figure 7 Illustration of a phenomenon that can cause gaps in confidence intervals. The inset shows how the problem can be remedied.

To be more explicit about Step (2), Figure 7 shows how a sequence of highest minimal cardinality acceptance curves can occasionally transition from $AC(a-b)$ to $AC((a-1)-b)$ when $M(\lambda)$ increases. This causes a gap in the confidence set for $x = a - 1$, since the previously generated confidence interval for $a - 1$ then has the subinterval $\{\lambda: CP(\lambda) = AC((a-1)-b)\}$ appended to it. In such cases (which are rare), transitioning from $AC(a-b)$ to $AC(a-(b+1))$ rather than to $AC((a-1)-b)$ preserves monotonicity of the $\{a\}$ sequence and eliminates the above subinterval from the confidence set for $a - 1$. For the specific case shown in Figure 7, the resolution given in Step (2) is to transition from $AC(24-46)$ when it hits the confidence level directly to $AC(24-47)$, onto the portion shown as a dashed line in the inset of Figure 7.

Except for those regions of the parameter space where Step (2) applies, the resulting procedure is equivalent to that of Sterne's. Figure 8 shows the level 95% cpf obtained from this approach. Note that the transitions in Figure 8 between the acceptance curves of equal cardinality occur at the cusps where the curves meet. Because this method maximizes the cpf




Figure 8 Coverage probability function for the Optimal Coverage method, $1 - \alpha = 95\%$.

over the class of all minimal cardinality procedures that produce intervals for all x , we call this method the *OC (Optimal Coverage)* procedure.

The differences among CG, KB, and OC can be summarized as follows. When multiple acceptance curves of minimal cardinality are available, CG always uses the curve $AC(a-b)$ for which a and b are maximal, KB always uses the curve $AC(a-b)$ for which a and b are minimal, provided that the $\{a\}$ sequence remains non decreasing, and OC transitions between curves of equal cardinality at their points of intersection, typically near the middle of the possible range of such transitions. Because CG transitions as early as possible, its upper endpoints are as small as possible for a minimal cardinality procedure, whereas since KB transitions as late as possible, its upper endpoints are as large as possible.

Casella and Robert (1989) were the first to identify the class of all gapless Poisson confidence procedures having the Minimal Cardinality Property and presented a “refinement algorithm” that, when applied to any procedure not in that class, successively shortened its intervals until a procedure within that class is obtained. They also introduced the following asymptotic criterion for assessing the length performance of a Poisson confidence procedure:

A $1 - \alpha$ confidence procedure $C' = \{[l'_x, u'_x], x = 0, 1, \dots\}$ dominates a competing procedure $C = \{[l_x, u_x], x = 0, 1, \dots\}$ if there exists an N_0 such that for all $N > N_0$, either

$$\sum_{x=0}^{\infty} (u'_x - l'_x) < \sum_{x=0}^{\infty} (u_x - l_x)$$

or

$$\sum_{x=0}^{\infty} (u'_x - l'_x) = \sum_{x=0}^{\infty} (u_x - l_x)$$

and $P(\lambda \in [l'_x, u'_x]) \geq P(\lambda \in [l_x, u_x]) \forall \lambda$, with strict inequality for at least one λ .

Thus C' dominates C if the sum of its interval lengths eventually remains smaller than the corresponding sum for C , with coverage used as a tiebreaker. Note that Casella and Robert's criterion does not address how two procedures compare over their initial N intervals.

We also discuss here a procedure described by Blaker (2000). Define the tail probability of an observed value x to be the minimum of $P(X \leq x)$ and $P(X \geq x)$. Then for any given λ , define x as *rare* if the probability of observing a value with a tail probability as small as that of x does not exceed α . The acceptance set for each λ for Blaker's method is those x that are not rare for that λ .

Blaker's method performs quite well with regard to length; its coverage is also high. However, Blaker's method does not have the Minimal Cardinality Property, as at its non coincidental lower endpoints its cpf typically exceeds $1 - \alpha$; thus Blaker's procedure can be refined to create a strict minimal cardinality procedure that produces smaller intervals. The refinement consists of simply increasing each non coincidental lower endpoint as much as possible, to the point of making the value of the cpf equal to $1 - \alpha$ at that endpoint.

4. A refinement paradox

With the above criterion, Casella and Robert show that any confidence procedure that has the same lower or upper endpoint for consecutive values of x can be dominated (see their Proposition 2.2). In particular, CG can be dominated. Now suppose we try to refine the 95% CG confidence procedure as suggested by Casella and Robert by increasing coverage without increasing asymptotic interval length, that is, without changing the value of

$\frac{1}{N} \sum_{x=0}^N (u_x - l_x)$ for sufficiently large N . To accomplish this we can increase each CG coincidental endpoint one by one to the place that yields the most gain in coverage. First we increase $u_0 = l_8$ from CG's value, 3.285, to 3.764, the value used for OC. This creates an increase in coverage for all $\lambda \in (3.285, 3.764)$. Moreover, this move increases the interval length for $x = 0$ by $3.764 - 3.285 = 0.479$ and decreases the interval for $x = 8$ by the same amount. After refining the first coincidental endpoint we have increased the sum of interval lengths $\frac{1}{N} \sum_{x=0}^N (u_x - l_x)$ for $N < 8$, but not changed the sum for all $N \geq 8$. Thus this refined confidence procedure dominates CG since we improved overall coverage without increasing asymptotic length.

Now suppose we make an analogous refinement for the next coincidental endpoint, increasing $u_1 = l_{11}$ from its value for CG to the OC value. These first two moves have increased coverage without increasing the asymptotic length since $\frac{1}{N} \sum_{x=0}^N (u_x - l_x)$ is unchanged for $N \geq 11$, although it is larger than before for $N < 11$.

Since for all x , u_x is coincidental to l_y for some $y > x$, we can continue this refinement process indefinitely. In the limit, this sequence of refinements yields the OC procedure. Yet $\frac{1}{N} \sum_{x=0}^N (u_x - l_x)$ is larger for OC than the CG method for all N . Thus, comparing two strict confidence procedures' interval lengths using Casella and Robert's criterion can create a contradiction—an infinite sequence of improvements leading to a procedure that is inferior to the original one with respect to length.

We propose an alternative criterion by which to compare the length performance of Poisson confidence procedures. We say that C' is superior to C on length if there exists N_0 such that

$$\frac{1}{N} \sum_{x=0}^N (u'_x - l'_x) \leq \frac{1}{N} \sum_{x=0}^N (u_x - l_x) \text{ for all } N, \text{ with } \frac{1}{N} \sum_{x=0}^N (u'_x - l'_x) < \frac{1}{N} \sum_{x=0}^N (u_x - l_x) \text{ for all } N > N_0. \tag{1}$$

Using this measure, CG dominates OC, which in turn dominates KB. In fact, since CG always transitions to the curve $AC(a-b)$ with the highest possible values of a and b , CG is the unique length minimizing minimal cardinality procedure with respect to the above criterion. To see this, let C be any minimal cardinality confidence procedure, let C' be CG, and note that (i) the above sums have the same terms for all forced lower endpoints (places where $M(\lambda)$ increases) and (ii) all coincidental endpoints for which both lower and upper endpoints are included in the sums contribute 0 to those sums. Thus the distinction between the sums for C and C' reduces to a comparison of those upper endpoints that do not have a lower endpoint match in the sums, because $x > N$ for those l_x . Since CG transitions as early as possible, its lower endpoints fall at the smallest possible values.

Figure 9 compares the values of average interval length $\frac{1}{N} \sum_{x=0}^N (u_x - l_x)$ for $0 \leq N \leq 100$ for CG, OC, and KB to the corresponding values for Garwood's method. As expected, each of the minimal cardinality procedures described above yields intervals that are shorter on average than those of Garwood's method, except in some cases for the smallest values of x . Among those procedures, the length ranking $KB > OC > CG$ is consistent, with CG yielding a significant advantage in average interval length.

5. A recommendation for a near-length optimal Poisson confidence procedure

We have seen that CG is optimal with respect to length when judged for all $x \leq N$ for any positive integer N , yet it possesses the counterintuitive property of failing to have endpoints




Figure 14 Average 95% confidence interval lengths $\frac{1}{N} \sum_{x=0}^N (u_x - l_x)$ for $0 \leq N \leq 100$ for KB, OC, and CG relative to Garwood.

that are strictly increasing in x . Any modification of CG that satisfies the Minimal Cardinality Property must have endpoints that are larger than those of CG, at least for some x ; however, the effect of increasing the endpoints is to increase the length as measured by Equation (1); therefore it is prudent to keep such increases to a minimum. We therefore propose the following *Modified Crow–Gardner* (MCG) procedure that achieves strictly increasing endpoints:

Initially let $l_x(\text{MCG}) = l_x(\text{CG}) \forall x$. Then, beginning at $x = 0$, whenever $l_{x+1}(\text{MCG}) < ml(x) = l_x(\text{MCG}) + \min(.01l_x(\text{MCG}), 0.1)$, increase $l_{x+1}(\text{MCG})$ to $ml(x)$. Each modified endpoint represents an endpoint that is coincidental with some upper endpoint; thus, those upper endpoints change as well. All other endpoints remain unchanged.

The resulting procedure thus has all endpoints separated by at least 1% (if those endpoints are less than 10) or by at least 0.1 (if 10 or greater). The rationale for this specific adjustment to CG is as follows:

- (i) Increases in endpoints need to be kept small in order to keep confidence interval lengths nearly as small as for CG, the length optimal confidence procedure;
- (ii) Small endpoints (i.e., < 10) should be changed less than other endpoints so that the relative change in those endpoints is not large;
- (iii) It is a fairly common practice to round confidence interval endpoints to three significant figures, except that for endpoints above 100 a single decimal digit is retained. With such a rounding protocol, all MCG endpoints will remain different after rounding;
- (iv) A potential increase of much more than 0.1 leads to difficulties for large x . To see why, see Figure 10. All points shown in Figure 10 are CG lower endpoints, while the points in the top row of each panel are the CG upper endpoints (hence they are the coincidental endpoints). As x increases, progressively greater numbers of points are tied or nearly




Figure 15 Confidence interval endpoints for Crow–Gardner 95% confidence procedure.




Figure 11 Ratio of MCG to CG 95% confidence interval lengths for $x = 0, 1, \dots, N$

tied. For example, $l_{290} = l_{291} = l_{292} = 258.34$, and l_{293} is only slightly higher. For still larger values of x , it is regularly the case that even more endpoints are equal or nearly so. Modifying CG by spreading out the points in such a cluster by too large an amount would cause the modified points to run into the following points that are currently nearly evenly spaced out and not otherwise needing adjustment.

Note that MCG is still a minimal cardinality procedure. If $l_x(\text{CG}) = l_{x+1}(\text{CG}) = \lambda^*$, then increasing the coincidental endpoint for $x+1$ amounts to transitioning from the acceptance curve $AC(a-b)$ used just before λ^* to the curve $AC(a-(b+1))$ (as does OC) rather than to $AC((a+k)-(b+k+1))$ (as does CG); the cardinality of these latter two acceptance curves is the same.

In addition to not having identical lower or upper endpoints for different values of x , MCG has another advantage over CG in that its average coverage is higher. The intervals over which the tied endpoints of CG are increased for MCG are precisely the places where coverage is increased the most by exchanging $AC((a+k)-(b+k+1))$ for $AC(a-(b+1))$. For example, the alteration by MCG of the first tied endpoint for CG at level 95% results in the initial portion of the acceptance curve $AC(1-8)$ being replaced by $AC(0-7)$; refer to Figure 4 to see that coverage is increased over that interval from approximately 95.6% to approximately 98%.

Figure 11 shows how individual 95% MCG confidence intervals compare in length to those of CG. Observe that no MCG intervals are more than 1.2% longer than the corresponding CG intervals, and many are in fact shorter. Results for confidence levels 90% and 99% are similar.




Figure 12 Ratios of cumulative average interval length for MCG versus CG.




Figure 11 Expected lengths of KB, OC, and MCG relative to CG, $1 - \alpha = 95\%$.

Figure 12 shows for confidence levels 90%, 95%, and 99% the ratio of cumulative average interval length for MCG as compared to CG,

$$\frac{\sum_{x=0}^N (u_x(MCG) - l_x(MCG))}{N} \quad \frac{\sum_{x=0}^N (u'_x(CG) - l'_x(CG))}{N},$$

for $N = 0, 1, \dots, 25$. Cumulative average MCG interval length ranges from about 0.5% to 1.0% greater than for CG for very small N , with a rapid decrease in excess length as N increases.

When N is large, the disparity between MCG and CG with respect to cumulative average interval length becomes negligible, as shown in Table 1:

Table 1 Ratios of MCG to CG cumulative average confidence interval length for $x = 0, 1, \dots, N$

Confidence level	$N =$	10	100	200
90%		1.00416	1.00053	1.00034
95%		1.00722	1.00065	1.00032
99%		1.00488	1.00054	1.00027

An alternative way to compare the length performance of different confidence procedures produced is through *expected length*. Figure 13 shows the ratio of expected lengths for the 95% KB, OC, and MCG confidence procedures relative to that of CG for $0 \leq \lambda \leq 25$. Both KB and OC give significantly larger expected length than CG; however, the expected confidence interval length for MCG is only about 1% greater than that of CG for the smallest values of λ and is nearly identical to that of CG for larger λ .

Confidence limits of the MCG procedure for $1 - \alpha = 90\%, 95\%$, and 99% and $0 \leq x \leq 200$ can be found at www.csun.edu/~hcmth031/MCG.pdf.

6. Approximate Poisson confidence procedures

We call a confidence procedure an *approximate* procedure when the coverage of the procedure falls below the stated confidence level, but not by a large amount (at least for most λ). Approximate confidence procedures are widely used for discrete data. Allowing the cpf to fall below the nominal confidence level results in smaller confidence intervals than strict procedures produce.

A natural way to create a good approximate discrete confidence procedure is to adjust the confidence level of a high-performing strict procedure to the value that makes the mean actual coverage equal to the nominal level. Minimal cardinality confidence procedures are

Table 1 Comparison of strict and approximate 95% Poisson confidence procedures. The rightmost column gives the average ratio of interval lengths for $0 \leq x \leq 100$ relative to the strict 95% CG procedure.

	Minimum coverage	Mean coverage on $0 \leq \lambda \leq 100$	Average interval length relative to strict CG
Strict CG	95.00%	95.34%	—
Strict MCG	95.00%	95.36%	+0.12%
Approximate CG	94.62%	95.00%	−1.48%
Approximate MCG	94.60%	95.00%	−1.42%

particularly attractive for this purpose because the coverage of such a procedure achieves its minimum at numerous locations, which keeps that minimum much higher than for other approximate procedures. In other words, when a minimal cardinality approximate confidence procedure falls below the nominal level, it does not do so by much.

This approach was used for the one-sample binomial case by Reiczigel (2003) and Schilling and Doi (2014). Since the parameter space is unbounded in the Poisson case, specification of the interval for λ over which mean coverage is set equal to the nominal confidence level is required. Table 2 shows a comparison between CG and MCG approximate confidence procedures each adjusted to have mean coverage equal to 95% over the interval $0 \leq \lambda \leq 100$ and the strict CG and MCG 95% confidence procedures.

The minimum coverage of both approximate procedures is only slightly below the nominal level, which arguably could be considered a good tradeoff in order to gain a reduction in average interval length of approximately 1.5%. We recommend that when an approximate procedure is used in practice, a statement to that effect is provided along with the confidence interval.

References

- Blaker, H. (2000). Confidence curves and improved exact confidence intervals for discrete distributions. *Can. J. Stat.* 28(4):783–798.
- Casella, G. (1986). Refining binomial confidence intervals. *Can. J. Stat.* 14:113–129.
- Casella, G., Robert, C. (1989). Refining Poisson confidence intervals. *Can. J. Stat.* 17(1):45–57.
- Crow, E.L., Gardner, R.S. (1959). Confidence intervals for the expectation of a Poisson variable. *Biometrika* 46(3/4):441–453.
- Garwood, F. (1936). Fiducial limits for the Poisson distribution. *Biometrika* 28(3/4):437–442.
- Kabaila, P., Byrne, J. (2001). Exact short Poisson confidence intervals. *Can. J. Stat.* 29(1):99–106.
- Kabaila, P., Byrne, J. (2005). Comparison of Poisson confidence intervals. *Commun. Stat. – Theory Methods* 34(3):545–556.
- Reiczigel, J. (2003). Confidence intervals for the binomial parameter: some new considerations. *Stat. Med.* 22:611–621.
- Schilling, M.F., Doi, J.A. (2014). A coverage probability approach to finding an optimal binomial confidence procedure. *Am. Statistician* 68(3):133–145.
- Sterne, Theodore E. (1954). Some remarks on confidence or fiducial limits. *Biometrika* 41(1/2):275–278.