

# AVERAGED STOCHASTIC APPROXIMATION FOR SEQUENTIAL FIXED-WIDTH CONFIDENCE INTERVALS UNDER RIGHT CENSORING

BOLTAEVA NARGIZA

NATIONAL UNIVERSITY OF UZBEKISTAN NAMED AFTER M. ULUGBEK, TASHKENT, UZBEKISTAN  
boltayevan27@gmail.com

**ABSTRACT.** We construct sequential fixed-width confidence intervals for a parameter of a right-censored lifetime distribution observed as a stream, where the Fisher information is unknown and must be learned online. The procedure combines the Polyak–Ruppert averaged Robbins–Monro estimator with a recursively maintained information estimate, both updated in  $O(1)$  time and memory. Under standard regularity and the gain condition  $cI > 1/2$ , we prove asymptotic consistency and Chow–Robbins efficiency ( $N_\varepsilon/n_\varepsilon \rightarrow 1$  a.s. and in mean). The analysis relies on an asymptotically linear representation of the averaged estimator, whose remainder is controlled almost surely via the martingale structure of the score. We show that stopping-time efficiency holds for every gain, while the gain condition governs coverage alone. Simulations on censored exponential and Weibull models confirm the theory and demonstrate that using the information in place of its reciprocal in the stopping rule collapses the procedure.

**MSC (2020):** 62L12; 62N02; 62L20; 62F12.

**Key words:** sequential estimation, fixed-width confidence intervals, right-censoring, stochastic approximation, Polyak–Ruppert averaging, Anscombe’s theorem.

## 1. Introduction

Fixed-width sequential confidence intervals, introduced by Chow and Robbins [4], determine the sample size adaptively as data accrue. The classical theory assumes i.i.d. observations and an asymptotic variance that is known or consistently estimable from a pilot. In survival and reliability analysis, however, observations are right-censored and the Fisher information depends on the unknown censoring law; it cannot be computed a priori and must be estimated *in real time*, which forces the sampling rule itself to be adaptive. For streaming data a further constraint is computational: maximum likelihood, EM, or estimating-equation solvers reprocess the sample and cost  $O(n)$  or worse per update. Stochastic approximation [23] updates in  $O(1)$  time and memory, and Polyak–Ruppert averaging [21, 22, 25] gives the averaged iterate the optimal asymptotic covariance  $I^{-1}$ , independently of the step size. It is thus a natural vehicle for sequential interval estimation: efficient in the Hájek–Le Cam sense, cheap, and carrying a recursively updatable information estimate that can drive a data-dependent stopping rule.

**Related work.** Fixed-width sequential intervals originate with [4]; coverage rests on Anscombe’s random-index theorem [1] and its nonlinear-renewal refinements [28, 8]. A recent literature treats online and streaming estimation at a *fixed* sample size: renewable regression [15], stochastic-gradient estimation for censored streams [29], and, for the Polyak–Ruppert average, variance estimation via batch-means [3], the online bootstrap [7], or random scaling [13, 26]. All fix  $n$  and report a shrinking-width interval. To our knowledge none couples a recursive information estimate with a data-driven stopping time to deliver a fully sequential *fixed-width* interval that is consistent and Chow–Robbins efficient under censoring; this is the gap we fill.

Given a stream of right-censored observations we run an averaged Robbins–Monro recursion, maintain an online estimate  $\hat{I}_n$  of the Fisher information, and stop when the implied half-width reaches  $\varepsilon$ . Our contributions are: **(i)** an asymptotically consistent, Chow–Robbins efficient procedure (Corollary 1): coverage tends to  $\gamma$  (Theorem 2) and  $N_\varepsilon/n_\varepsilon \rightarrow 1$  a.s. and in mean (Proposition 2); **(ii)** a corrected asymptotically linear representation (Lemma 2) whose frozen-score correction, not amenable to a direct Cauchy–Schwarz bound ( $O(n^{-1/2})$ ), attains  $o(n^{-1/2})$  via the martingale-difference structure of the score; **(iii)** a separation of design requirements:

stopping-time efficiency holds for *every* gain  $c$ , whereas the gain condition  $cI > \frac{1}{2}$  governs coverage alone; **(iv)** a cautionary identity: the controlling scale is the variance  $I^{-1}$ , not  $I$ , and placing  $\hat{I}_n$  where  $\hat{I}_n^{-1}$  belongs stops a factor  $I^2$  too early, catastrophic under heavy censoring ( $I < 1$ ).

The paper is organized as follows. Section 2 fixes the model and assumptions; Section 3 develops the asymptotic theory, culminating in properties (H1)-(H3); Sections 4 and 5 define the procedure and prove stopping-time efficiency and coverage; Section 6 reports simulations; Section 7 discusses extensions and open problems. Longer proofs are collected in the Appendix.

## 2. Model, recursion, and assumptions

**Censoring model and score.** Let  $X_1, X_2, \dots$  be i.i.d. failure times with density  $f(\cdot; \theta_0)$  and c.d.f.  $F(\cdot; \theta_0)$ ,  $\theta_0 \in \Theta \subseteq \mathbb{R}^k$  open, and let  $C_1, C_2, \dots$  be i.i.d. censoring times with distribution  $G$ , independent of the  $X_i$  and not depending on  $\theta$ . One observes

$$Y_i = \min(X_i, C_i), \quad \delta_i = \mathbf{1}\{X_i \leq C_i\}, \quad i = 1, 2, \dots,$$

with filtration  $\mathcal{F}_n = \sigma(Y_1, \delta_1, \dots, Y_n, \delta_n)$ . The single-observation censored score is

$$\psi(y, \delta, \theta) = \delta \frac{\partial_\theta f(y; \theta)}{f(y; \theta)} - (1 - \delta) \frac{\partial_\theta F(y; \theta)}{1 - F(y; \theta)}, \quad (1)$$

the conditional expectation of the full-data score given  $(Y, \delta)$  [19, 14]. In particular  $\mathbb{E}_{\theta_0} \psi(Y, \delta, \theta_0) = 0$ , and the Fisher information equals

$$I = \mathbb{E}_{\theta_0} [\psi(Y, \delta, \theta_0) \psi(Y, \delta, \theta_0)^\top].$$

Throughout this paper  $k = 1$ , so  $I > 0$  is a scalar; the multivariate case is discussed in Remark 5.

Write  $Z_i = (Y_i, \delta_i)$  for the observations,  $P_\theta$  for their law,  $\mathbb{P}_n = n^{-1} \sum_{i=1}^n \delta_{Z_i}$  for the empirical measure,  $Pf = \int f dP_\theta$ , and  $\mathbb{G}_n f = \sqrt{n}(\mathbb{P}_n - P_\theta)f$  for the empirical process; thus  $M(\theta) = P_\theta \psi(\cdot, \theta)$ .

**Remark 1 ( $I^{-1}$  is the information bound).** Under Assumption 2 the censored experiments are locally asymptotically normal with central sequence  $\mathbb{G}_n \psi(\cdot, \theta_0)$  and information  $I$  [27, Ch. 7]; the censored score (1) is the efficient score and  $\tilde{\psi}_{\theta_0} := I^{-1} \psi(\cdot, \theta_0)$  the efficient influence function. By the Hájek-Le Cam convolution theorem  $I^{-1}$  is the information bound for any regular estimator [9, 11, 12, 2]; the averaged estimator attains it (Theorem 1), so  $n_\varepsilon = a^2 I^{-1} / \varepsilon^2$  is the intrinsic optimal sample size.

**Averaged Robbins-Monro recursion** For a compact convex  $K \subset \Theta$  with  $\theta_0 \in \text{int } K$  and Euclidean projection  $\Pi_K$ , and gain  $\gamma_n = c/n$ , define

$$\theta_{n+1} = \Pi_K[\theta_n + \gamma_n \psi(Y_{n+1}, \delta_{n+1}, \theta_n)], \quad \bar{\theta}_n = \frac{1}{n} \sum_{k=1}^n \theta_k. \quad (2)$$

The projection controls the boundary singularity of  $\psi$ . Since  $\theta_n \xrightarrow{\text{a.s.}} \theta_0$  (proved in Theorem 1) and  $\theta_0 \in \text{int } K$ , the projection  $\Pi_K$  is eventually inactive almost surely, so it does not affect the limit theory. Write  $M(\theta) = \mathbb{E}_{\theta_0} \psi(Y, \delta, \theta)$  for the drift.

### Regularity conditions

We use the following standard conditions.

**Assumption 1 (identifiability).**  $\theta_0$  is the unique zero of  $M$  in  $\Theta$ ,  $M$  is  $C^1$  near  $\theta_0$ , and  $M'(\theta_0) = -I$ . (The last equality is the second Bartlett identity for the censored score (1), recorded here for reference.)

**Assumption 2 (smoothness).**  $\theta \mapsto \psi(y, \delta, \theta)$  is  $C^1$  on a neighbourhood  $U_0$  of  $\theta_0$ , with  $\mathbb{E}_{\theta_0}[\sup_{U_0} \psi^2] < \infty$  and  $\mathbb{E}_{\theta_0}[\sup_{U_0} (\partial_\theta \psi)^2] < \infty$ ; in particular  $\theta \mapsto \mathbb{E}_{\theta_0}[\psi^2]$  is continuous on  $U_0$ .

**Assumption 3 (moments and information).**  $\mathbb{E}_{\theta_0} |\psi(Y, \delta, \theta_0)|^{2+\eta} < \infty$  for some  $\eta > 0$  (with  $\eta = 2$ , i.e. a finite fourth moment, where the  $L^2$  control in Lemma 2 requires it), and  $I = \mathbb{E}_{\theta_0} [\psi(Y, \delta, \theta_0)^2] > 0$ .

**Assumption 4 (projection).**  $K$  is compact convex,  $\theta_0 \in \text{int } K$ , and  $\Pi_K$  is the Euclidean projection onto  $K$ .

**Assumption 5 (local strong monotonicity).**  $(\theta - \theta_0)M(\theta) \leq -\lambda(\theta - \theta_0)^2$  on a neighbourhood of  $\theta_0$  for some  $\lambda > 0$ . This is not independent: it follows from Assumptions 2-3 for any  $0 < \lambda < I$  on a sufficiently small neighbourhood  $U$ , by the mean value theorem (the constant  $\lambda$  may be taken arbitrarily close to, but not equal to,  $I$ ).

These conditions hold for the exponential and Weibull censored families of Section 6 and, more generally, for regular parametric lifetime models under a censoring law with sufficiently many moments. The *gain condition*  $cI > \frac{1}{2}$  is invoked wherever the efficient rate is claimed; its role is delineated in Sections 3 and 5.

### 3. Asymptotic theory of the averaged estimator

Modulo two classical inputs – the Polyak-Juditsky averaging principle and a moment bound for harmonic-gain stochastic approximation – this section delivers the three properties used thereafter:

- (H1) asymptotic linearity  $\bar{\theta}_n - \theta_0 = \mathbb{P}_n \tilde{\psi}_{\theta_0} + R_n$ ,  $\tilde{\psi}_{\theta_0} = I^{-1}\psi(\cdot, \theta_0)$ ;
- (H2) negligible remainder  $\sqrt{n} R_n \xrightarrow{\text{a.s.}} 0$ ;
- (H3) consistent scale  $V_n \xrightarrow{\text{a.s.}} \sigma^2 := I^{-1}$ ,  $0 < V_n \leq \iota^{-1}$ .

Here  $\mathbb{P}_n \tilde{\psi}_{\theta_0} = \frac{1}{n} \sum_i \phi_i$  with  $\phi_i = I^{-1}\psi(Y_i, \delta_i, \theta_0)$  i.i.d. of mean 0 and variance  $I^{-1} = \sigma^2$ ; (H1) states that  $\bar{\theta}_n$  is asymptotically linear with the efficient influence function [27, Ch. 25].

**Consistency, normality, and the non-asymptotic bias** Under Assumptions 1-4,  $\theta_n \xrightarrow{\text{a.s.}} \theta_0$  by the Robbins-Siegmund theorem [24], and the averaged iterate is efficient.

**Theorem 1 (efficiency of the averaged estimator).** Under Assumptions 1-4 and  $cI > \frac{1}{2}$ ,

$$\sqrt{n}(\bar{\theta}_n - \theta_0) \xrightarrow{d} \mathcal{N}(0, I^{-1}).$$

**Proof.** The proof proceeds by linearising the drift. By Assumption 2,  $M(\theta_n) = -I d_n + O(d_n^2)$ , so the error  $d_n = \theta_n - \theta_0$  obeys a linear recursion with drift  $-cI$  and martingale noise  $\xi_{n+1} = \psi(Y_{n+1}, \delta_{n+1}, \theta_n) - M(\theta_n)$ . Averaging the iterates cancels the gain dependence of the asymptotic covariance, leaving  $\Gamma^{-1} S \Gamma^{-\top} = I^{-1}$  with  $\Gamma = -cI$  and  $S = c^2 I$ ; the boundary condition  $cI > \frac{1}{2}$  ensures the drift matrix  $-cI$  is Hurwitz [6, 18], which is what makes the rate  $\gamma_n = c/n$  admissible. The projection  $\Pi_K$  is inactive eventually, since  $\theta_0 \in \text{int } K$  and  $\theta_n \xrightarrow{\text{a.s.}} \theta_0$ , so the unprojected limit applies [10, Thm. 5.2.1]. The full argument is that of Polyak-Juditsky [22, 10]; see the Appendix. Equivalently  $\sqrt{n}(\bar{\theta}_n - \theta_0) = \mathbb{G}_n \tilde{\psi}_{\theta_0} + o_p(1)$ : the averaged estimator is asymptotically linear with the efficient influence function, reproducing the MLE's limiting variance  $I^{-1}$  with no matrix inversion.  $\square$

The finite-sample behaviour is governed by a non-asymptotic expansion, which explains both the coverage results and the role of  $cI$ .

**Proposition 1 (non-asymptotic mean squared error).** Under Assumptions 1-4 and  $\gamma_n = c/n$ ,

$$\mathbb{E}[(\bar{\theta}_n - \theta_0)^2] = \frac{1}{In} + B_n, \quad B_n = \begin{cases} O(n^{-2cI}), & cI \neq \frac{1}{2}, \\ O(n^{-1}(\log n)^2), & cI = \frac{1}{2}, \end{cases} \quad (3)$$

so the efficient term  $1/(In)$  dominates iff  $cI > \frac{1}{2}$ . Consequently the root-mean-square transient bias is of order  $n^{-cI}$ .

**Proof idea.** The proof uses the variation-of-constants (Duhamel) representation of the linearised recursion: the deterministic transient and the nonlinearity remainder contribute  $O(n^{-cI})$ , while the stochastic term gives the leading  $1/(In)$ . At  $cI = \frac{1}{2}$  the logarithmic factor arises because the two competing exponents coincide. The last sentence is the fact used below: the *bias*, not its square, enters the linear representation, so the remainder there is  $O(n^{-cI})$  and  $\sqrt{n}$  times it is  $O(n^{1/2-cI})$ , which vanishes exactly when  $cI > \frac{1}{2}$ .

**Lemma 1 (moment bound).** Under Assumptions 1-4,  $\gamma_n = c/n$  and  $cI > \frac{1}{2}$ ,

$$\sup_n n \mathbb{E} d_n^2 < \infty; \quad \text{if moreover } \mathbb{E}_{\theta_0} |\psi(Y, \delta, \theta_0)|^4 < \infty, \text{ then } \sup_n n^2 \mathbb{E} d_n^4 < \infty.$$

This is the standard  $L^2/L^4$  moment bound for harmonic-gain SA [18, Ch. 4], [17, Lemma 1]; see the harmonic-gain moment recursion.

## The asymptotically linear representation

**Lemma 2 (linear representation with a.s.-negligible remainder).** Under Assumptions 1-4, the fourth-moment strengthening of Assumption 3, and  $cI > \frac{1}{2}$ , (H1)-(H2) hold. In addition  $\mathbb{E}[(\sqrt{n}R_n)^2] = O(n^{-1} \log n)$ .

**Proof Sketch.** Write  $d_n = \theta_n - \theta_0$ ,  $\xi_{n+1} = \psi(Y_{n+1}, \delta_{n+1}, \theta_n) - M(\theta_n)$  and  $\xi_{n+1}^0 = \psi(Y_{n+1}, \delta_{n+1}, \theta_0)$ , both martingale differences, with  $\phi_{n+1} = I^{-1}\xi_{n+1}^0$ .

*Step 1 (Polyak-Juditsky).* By Assumption 2,  $M(\theta_n) = -I d_n + O(d_n^2)$ ; the averaging principle [22, Thm. 2] for the linear recursion with drift  $-cI$  gives  $\bar{\theta}_n - \theta_0 = \frac{1}{n} \sum_{k=1}^n I^{-1} \xi_k + P_n$  with  $\sqrt{n}P_n \xrightarrow{\text{a.s.}} 0$ . Its deterministic-transient and nonlinearity parts are  $O(n^{-cI})$  pathwise (via  $\pi_{n,1} = \prod_{k=1}^{n-1} (1 - cI/k) = \Theta(n^{-cI})$ ), hence  $o(n^{-1/2})$  since  $cI > \frac{1}{2}$ ; the stochastic averaging discrepancy is  $o(n^{-1/2})$  a.s. by the refinements of [20, 16, 5], the sole external inputs beyond Step 2. (If only  $\sqrt{n}P_n \xrightarrow{P} 0$  is available, (H2) holds in probability and all conclusions persist.)

*Step 2 (replacing the frozen score).* Put  $D_n = \frac{1}{n} \sum_{k \leq n} I^{-1}(\xi_k - \xi_k^0)$ , so  $R_n = P_n + D_n$ . The increments are martingale differences, and by the mean value theorem, using the boundedness of  $\partial_\theta \psi$  on  $U_0$  (Assumption 2),  $\mathbb{E}[(\xi_k - \xi_k^0)^2 | \mathcal{F}_{k-1}] = O(d_{k-1}^2)$ , so with  $M_n = \sum_{k \leq n} I^{-1}(\xi_k - \xi_k^0)$  one has  $\mathbb{E}[(\Delta M_k)^2] = O(\mathbb{E}d_{k-1}^2) = O(k^{-1})$  by Lemma 1. Thus  $\mathbb{E}[(\sqrt{n}D_n)^2] = \frac{1}{n} \sum_{k \leq n} O(k^{-1}) = O(n^{-1} \log n)$ , giving  $\sqrt{n}D_n \xrightarrow{L^2} 0$ ; and since  $\sum_k k^{-1} \mathbb{E}[(\Delta M_k)^2] < \infty$ , the  $L^2$ -bounded martingale  $\sum_k k^{-1/2} \Delta M_k$  converges a.s. (by the  $L^2$ -martingale convergence theorem [10]), so Kronecker's lemma yields  $\sqrt{n}D_n \xrightarrow{\text{a.s.}} 0$ . A direct Cauchy-Schwarz bound would give only  $O(n^{-1/2})$ ; centring on the martingale difference is what recovers the sharper rate. Combining,  $\sqrt{n}R_n \xrightarrow{\text{a.s.}} 0$  and  $\mathbb{E}[(\sqrt{n}R_n)^2] = O(n^{-1} \log n)$ .  $\square$

**Online variance estimate.** Fix a floor  $\iota \in (0, I)$  and estimate the information recursively by

$$\hat{I}_n = \frac{1}{n} \sum_{i=1}^n \psi(Y_i, \delta_i, \bar{\theta}_{i-1})^2, \quad V_n = (\hat{I}_n \vee \iota)^{-1}. \quad (4)$$

**Lemma 3 (strong consistency of the variance estimate).** Under Assumptions 1-5,  $\hat{I}_n \xrightarrow{\text{a.s.}} I$ ; hence  $V_n \xrightarrow{\text{a.s.}} I^{-1} = \sigma^2$  and  $0 < V_n \leq \iota^{-1}$  for all  $n$ . Thus (H3) holds.

**Proof.** Since  $\theta_n \xrightarrow{\text{a.s.}} \theta_0$ , also  $\bar{\theta}_{i-1} \xrightarrow{\text{a.s.}} \theta_0$ . With  $g(\theta) = \mathbb{E}_{\theta_0}[\psi(Y, \delta, \theta)^2]$  continuous near  $\theta_0$  and  $g(\theta_0) = I$ , write  $\hat{I}_n - I = (a) + (b)$  where (a)  $= \frac{1}{n} \sum_i [\psi(Y_i, \delta_i, \bar{\theta}_{i-1})^2 - g(\bar{\theta}_{i-1})]$  and (b)  $= \frac{1}{n} \sum_i [g(\bar{\theta}_{i-1}) - I]$ . Term (b)  $\rightarrow 0$  a.s. by continuity and the Cesàro lemma. For term (a), on a compact sub-neighbourhood  $\bar{U} \subset U_0$  of  $\theta_0$  the class  $\{\psi(\cdot, \theta)^2 : \theta \in \bar{U}\}$  is  $P_{\theta_0}$ -Glivenko-Cantelli—it has an integrable envelope and is continuous in  $\theta$  (Assumptions 2-3), hence [27, Thm. 19.4]  $\sup_{\theta \in \bar{U}} |(\mathbb{P}_n - P)\psi(\cdot, \theta)^2| \xrightarrow{\text{a.s.}} 0$ ; since  $\bar{\theta}_{i-1} \in \bar{U}$  eventually, term (a)  $\xrightarrow{\text{a.s.}} 0$ . The continuous mapping theorem gives  $V_n \xrightarrow{\text{a.s.}} I^{-1}$  ( $I > \iota$ ), and the bound is immediate.  $\square$

## 4. Sequential procedure and stopping-time efficiency

**The procedure.** Fix a confidence level  $\gamma \in (0, 1)$  and half-width  $\varepsilon > 0$ , and set  $a = \Phi^{-1}(\frac{1+\gamma}{2})$ , so  $2\Phi(a) - 1 = \gamma$ . Because  $\sqrt{n}(\bar{\theta}_n - \theta_0) \xrightarrow{d} \mathcal{N}(0, I^{-1})$ , the error law is  $\mathcal{N}(0, \sigma^2/n)$  with  $\sigma^2 = I^{-1}$ , and the Chow-Robbins optimal (but infeasible) nonrandom sample size is

$$n_\varepsilon = \frac{a^2 \sigma^2}{\varepsilon^2} = \frac{a^2 I^{-1}}{\varepsilon^2} = \frac{a^2}{I \varepsilon^2}. \quad (5)$$

Replacing  $\sigma^2$  by  $V_n$  from (4) and choosing a warm-up  $n_0 = n_0(\varepsilon)$  with  $n_0 \rightarrow \infty$ ,  $n_0 = o(n_\varepsilon)$  (e.g.  $n_0 = \lceil \varepsilon^{-\alpha} \rceil$ ,  $0 < \alpha < 2$ ), define

$$N_\varepsilon = \inf \left\{ n \geq n_0 : n \geq \frac{a^2 V_n}{\varepsilon^2} = \frac{a^2}{(\hat{I}_n \vee \iota) \varepsilon^2} \right\}, \quad I(N_\varepsilon) = [\bar{\theta}_{N_\varepsilon} - \varepsilon, \bar{\theta}_{N_\varepsilon} + \varepsilon]. \quad (6)$$

Both  $\hat{I}_n$  and  $V_n$  update by a single rank-one addition per step; because  $\hat{I}_n$  is  $\mathcal{F}_n$ -measurable,  $\{N_\varepsilon \leq n\} \in \mathcal{F}_n$ , so  $N_\varepsilon$  is a bona fide stopping time and the whole procedure runs in  $O(1)$  memory.

**Remark 2 (the scale is the variance, not the information).** The information  $I$  appears in the denominator of (5): a more informative model (larger  $I$ , lighter censoring) needs fewer observations. A rule must place the variance estimate  $V_n \rightarrow I^{-1}$  in the numerator. Using  $\hat{I}_n$  itself would target  $a^2 \hat{I}_n / \varepsilon^2$ , a factor  $I^2$  of (5); under heavy censoring  $I < 1$ , so  $I^2 \ll 1$  and the rule stops far too early. Section 6 shows the resulting collapse.

### Efficiency of the stopping time

**Proposition 2 (stopping-time properties).** Assume only (H3). Then for every  $\gamma \in (0, 1)$ : (i)  $N_\varepsilon \leq \max\{n_0, \lceil a^2 \iota^{-1} / \varepsilon^2 \rceil\} < \infty$  surely; (ii)  $N_\varepsilon \xrightarrow{\text{a.s.}} \infty$  as  $\varepsilon \rightarrow 0$ ; (iii)  $N_\varepsilon / n_\varepsilon \xrightarrow{\text{a.s.}} 1$  as  $\varepsilon \rightarrow 0$ ; (iv)  $\mathbb{E} N_\varepsilon / n_\varepsilon \rightarrow 1$  as  $\varepsilon \rightarrow 0$ .

**Proof.** (i) As  $V_n \leq \iota^{-1}$ , the threshold  $a^2 V_n / \varepsilon^2 \leq a^2 \iota^{-1} / \varepsilon^2 =: m^*$  for every  $n$ , so the defining inequality holds at  $n = \max\{n_0, \lceil m^* \rceil\}$ .

(ii) Work on the probability-one event  $\{V_n > 0 \ \forall n\} \cap \{V_n \rightarrow \sigma^2\}$ . If  $N_\varepsilon \leq B$  for some finite  $B$ , then  $B \geq N_\varepsilon \geq a^2 V_{N_\varepsilon} / \varepsilon^2 \geq a^2 (\min_{n_0 \leq n \leq B} V_n) / \varepsilon^2$ ; since the minimum is strictly positive, the right side diverges as  $\varepsilon \rightarrow 0$ , a contradiction. Hence  $N_\varepsilon \xrightarrow{\text{a.s.}} \infty$ .

(iii) For  $\varepsilon$  small enough that  $N_\varepsilon > n_0$ , minimality gives

$$\frac{a^2 V_{N_\varepsilon}}{\varepsilon^2} \leq N_\varepsilon < \max\left(n_0, \frac{a^2 V_{N_\varepsilon-1}}{\varepsilon^2}\right) + 1.$$

Dividing by  $n_\varepsilon = a^2 \sigma^2 / \varepsilon^2$ ,  $\frac{V_{N_\varepsilon}}{\sigma^2} \leq \frac{N_\varepsilon}{n_\varepsilon} < \frac{V_{N_\varepsilon-1}}{\sigma^2} + \frac{1+n_0}{n_\varepsilon}$ . As  $\varepsilon \rightarrow 0$ ,  $n_\varepsilon \rightarrow \infty$  and  $n_0 / n_\varepsilon \rightarrow 0$ ; by (ii)  $N_\varepsilon \rightarrow \infty$ , so  $V_{N_\varepsilon}, V_{N_\varepsilon-1} \rightarrow \sigma^2$  by the random-index principle ( $v_n \rightarrow v$  and integer  $N_\varepsilon \rightarrow \infty$  imply  $v_{N_\varepsilon} \rightarrow v$ ); both bounds tend to 1.

(iv) By (i),  $N_\varepsilon / n_\varepsilon \leq \max\{n_0, \lceil m^* \rceil\} / n_\varepsilon \rightarrow \iota^{-1} / \sigma^2 = I / \iota$ , a deterministic constant; bounded convergence upgrades (iii) to convergence in mean.  $\square$

**Remark 3 (efficiency needs no gain condition).** Proposition 2 uses only (H3), i.e. consistency of  $\hat{I}_n$ , which holds under Assumptions 1–5 *without* the gain condition. Thus  $N_\varepsilon / n_\varepsilon \rightarrow 1$  for every  $c > 0$ . The gain condition  $cI > \frac{1}{2}$  enters only in the coverage analysis below, through (H2). Efficiency and coverage are therefore governed by different requirements and must be tuned separately—a point confirmed numerically in the simulation study, where  $N_\varepsilon / n_\varepsilon \approx 1$  across all gains while coverage varies with  $cI$ .

## 5. Consistency of the fixed-width interval

**Theorem 2 (asymptotic consistency).** Assume (H1)–(H3). Then for every  $\gamma \in (0, 1)$ ,  $\lim_{\varepsilon \rightarrow 0} \mathbb{P}\{\theta_0 \in I(N_\varepsilon)\} = \gamma$ .

**Proof.** Let  $S_m = \sum_{i=1}^m \phi_i$ , i.i.d. with mean 0 and variance  $\sigma^2 = I^{-1}$ . By (H1),  $\sqrt{m}(\bar{\theta}_m - \theta_0) = S_m / \sqrt{m} + \sqrt{m}R_m$ .

*Leading term.* By Proposition 2(ii),  $N_\varepsilon / n_\varepsilon \xrightarrow{\text{P}} 1$  with  $n_\varepsilon \rightarrow \infty$  nonrandom. Anscombe’s random-sum central limit theorem for i.i.d. finite-variance summands [1], whose uniform-continuity-in-probability condition holds automatically for such partial sums [8, Thm. 3.1], gives  $S_{N_\varepsilon} / \sqrt{N_\varepsilon} \xrightarrow{d} \mathcal{N}(0, \sigma^2)$ .

*Remainder.* By (H2),  $\sqrt{m}R_m \xrightarrow{\text{a.s.}} 0$ , and  $N_\varepsilon \xrightarrow{\text{a.s.}} \infty$  (Proposition 2(ii)); the random-index principle gives  $\sqrt{N_\varepsilon}R_{N_\varepsilon} \xrightarrow{\text{a.s.}} 0$ , hence  $\xrightarrow{\text{P}} 0$ . (In the in-probability form of (H2), the same follows from uniform negligibility of  $\sqrt{m}R_m$  on windows  $[(1 \mp \eta)n_\varepsilon]$  via Doob’s  $L^2$  maximal inequality and the finite-time bounds of [16, 5].)

*Combination.* By Slutsky’s lemma,  $\sqrt{N_\varepsilon}(\bar{\theta}_{N_\varepsilon} - \theta_0) \xrightarrow{d} \mathcal{N}(0, \sigma^2)$ , equivalently  $T_\varepsilon := \sqrt{N_\varepsilon}I(\bar{\theta}_{N_\varepsilon} - \theta_0) \xrightarrow{d} \mathcal{N}(0, 1)$ . By Proposition 2(iii) and  $\varepsilon^2 n_\varepsilon I = a^2$ ,  $\varepsilon^2 N_\varepsilon I = a^2 (N_\varepsilon / n_\varepsilon) \xrightarrow{\text{a.s.}} a^2$ , so  $b_\varepsilon := \varepsilon \sqrt{N_\varepsilon}I \xrightarrow{\text{a.s.}} a$ . Then

$$\mathbb{P}\{\theta_0 \in I(N_\varepsilon)\} = \mathbb{P}\{|\bar{\theta}_{N_\varepsilon} - \theta_0| \leq \varepsilon\} = \mathbb{P}\{|T_\varepsilon| \leq b_\varepsilon\}.$$

Since  $(T_\varepsilon, b_\varepsilon) \xrightarrow{d} (Z, a)$  with  $Z \sim \mathcal{N}(0, 1)$  and  $\mathbb{P}\{|Z| = a\} = 0$ , the portmanteau theorem gives  $\mathbb{P}\{\theta_0 \in I(N_\varepsilon)\} \rightarrow \mathbb{P}\{|Z| \leq a\} = 2\Phi(a) - 1 = \gamma$ .  $\square$

**Corollary 1 (efficient consistent procedure).** Under Assumptions 1–4, the fourth-moment strengthening of Assumption 3, and  $cI > \frac{1}{2}$ , the family  $\{I(N_\varepsilon)\}_{\varepsilon > 0}$  of (6) is asymptotically consistent and asymptotically efficient in the Chow-Robbins sense ( $N_\varepsilon / n_\varepsilon \rightarrow 1$  a.s. and in mean).

**Proof.** Lemmas 2-3 give (H1)-(H3); combine Proposition 2 and Theorem 2.  $\square$

**Remark 4 (the gain condition and coverage).** The gain condition enters only through (H2). If  $cI \leq \frac{1}{2}$ , the transient bias is  $O(n^{-cI})$  (Proposition 1), so  $\sqrt{n}R_n = O(n^{1/2-cI}) \not\rightarrow 0$ ; the Anscombe step fails and coverage falls below  $\gamma$ . (*This is the corrected order: the bias is  $n^{-cI}$ , not  $n^{-2cI}$ ; the latter is the order of its square  $B_n$ .*) Quantifying the deficit as  $P\{\theta_0 \in I(N_\varepsilon)\} = \gamma + O(\varepsilon^\beta)$  requires the nonlinear renewal / Edgeworth machinery of Woodroffe [28] applied to the rates of  $\hat{I}_n^{-1} \rightarrow I^{-1}$  and of  $\sqrt{n}(\bar{\theta}_n - \theta_0)$ , and is left to future work.

**Remark 5 (multivariate case).** For  $k > 1$ , replace the interval by a fixed maximal-axis confidence ellipsoid. With  $\hat{I}_n = n^{-1} \sum_i \psi(Y_i, \delta_i, \theta_{i-1}) \psi(\cdot)^\top$  and  $E_n = \{\theta : n(\bar{\theta}_n - \theta)^\top \hat{I}_n (\bar{\theta}_n - \theta) \leq \chi_{k,\gamma}^2\}$ , the largest semi-axis is  $\{\chi_{k,\gamma}^2 \lambda_{\max}(\hat{I}_n^{-1})/n\}^{1/2}$ ; requiring it  $\leq \varepsilon$  gives  $N_\varepsilon = \inf\{n \geq n_0 : n \geq \chi_{k,\gamma}^2 \lambda_{\max}(\hat{I}_n^{-1})/\varepsilon^2\}$ . The vector analogue of Theorem 2 holds under  $c \lambda_{\min}(I) > \frac{1}{2}$ ; the least-informative direction, through  $\lambda_{\max}(\hat{I}_n^{-1}) = \lambda_{\min}(\hat{I}_n)^{-1}$ , governs the sample size. Preconditioning by  $\hat{I}_{n-1}^{-1}$  (an  $O(k^2)$  Sherman-Morrison update) reduces the gain condition to  $c > \frac{1}{2}$  regardless of the conditioning of  $I$ .

## 6. Simulation study and practical guidance

We use the exponential censored model:  $X \sim \text{Exp}(\lambda_0)$ ,  $C \sim \text{Exp}(\lambda_C)$  with  $\lambda_C = p_c \lambda_0 / (1 - p_c)$ , so the censoring fraction is  $p_c$  and  $I = (1 - p_c) / \lambda_0^2$ ; the score is  $\psi(y, \delta, \lambda) = \delta / \lambda - y$ . We take  $\lambda_0 = 1$ ,  $\gamma = 0.95$  ( $a = 1.96$ ), gain  $\gamma_n = c / (n + 50)$ , floor  $\iota = 10^{-6}$ , projection radius  $R = 100$ , a 200-observation pilot MLE start, and  $M = 2500$  Monte Carlo replications. All quantities are computed in a single streaming pass per replication.

Table 2: Table 1. Effect of the gain at  $p_c = 0.5$  ( $I = 0.5$ ),  $\varepsilon = 0.03$ ,  $n_\varepsilon = 8537$ ,  $M = 2500$  replications. Efficiency  $N_\varepsilon / n_\varepsilon$  equals 1.00 to two decimals throughout (observed range 0.999–1.004); coverage is governed by  $cI$ . Monte Carlo standard errors are  $\approx 0.006$ – $0.010$  for coverage and  $\lesssim 0.003$  for the ratio.

| $cI$                            | 0.30 | 0.50 | 0.75 | 1.00 | 1.50 | 2.00 | 3.00 | 4.00 | 5.00 |
|---------------------------------|------|------|------|------|------|------|------|------|------|
| $N_\varepsilon / n_\varepsilon$ | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| coverage                        | 0.46 | 0.61 | 0.75 | 0.83 | 0.88 | 0.90 | 0.92 | 0.93 | 0.93 |

**Stopping-time efficiency holds for every gain.** Table 1 sweeps the gain at  $p_c = 0.5$  ( $I = 0.5$ ),  $\varepsilon = 0.03$ . Across the entire range—including  $cI \leq \frac{1}{2}$ —the ratio  $N_\varepsilon / n_\varepsilon$  stays at 1.00 within Monte Carlo error, confirming Remark 3: efficiency of the stopping time does not require the gain condition. Coverage, by contrast, rises monotonically through the threshold  $cI = \frac{1}{2}$  and approaches—but at these sample sizes does not reach—the nominal 0.95. The residual gap is the finite-sample efficiency deficit of the averaged estimator quantified by  $B_n$  in Proposition 1 (the averaged MSE sits at a constant factor above the Cramér–Rao bound that shrinks only slowly in  $n$ ); it narrows as  $c$  increases (coverage  $0.90 \rightarrow 0.93$  from  $cI = 2$  to  $cI = 5$ ).

Table 3: Table 2. Correct versus naive stopping scale,  $\varepsilon = 0.03$ ,  $cI \geq 3$ ,  $M = 2500$ ; coverage in brackets (Monte Carlo standard error  $\approx 0.006$ – $0.010$ ). The naive rule stops at  $\approx I^2$  of the correct size.

|                                                | $p_c = 0.3$ ( $I=0.7$ ) | $p_c = 0.5$ ( $I=0.5$ ) | $p_c = 0.7$ ( $I=0.3$ ) |
|------------------------------------------------|-------------------------|-------------------------|-------------------------|
| correct: $N_\varepsilon / n_\varepsilon$ [cov] | 1.02 [0.93]             | 1.00 [0.93]             | 0.98 [0.93]             |
| naive: $N_\varepsilon / n_\varepsilon$ [cov]   | 0.47 [0.78]             | 0.25 [0.64]             | 0.11 [0.44]             |

**The correct scale is essential.** Table 2 contrasts the correct rule ( $V_n = \hat{I}_n^{-1}$ ) with the naive one ( $\hat{I}_n$ ) at  $\varepsilon = 0.03$  and gain tuned to  $cI \geq 3$ . The correct rule attains  $N_\varepsilon / n_\varepsilon \approx 1$  and coverage near 0.93 at all censoring levels. The naive rule stops at a fraction  $N_\varepsilon / n_\varepsilon \approx I^2$  of the required size—0.47, 0.25, 0.11 for  $p_c = 0.3, 0.5, 0.7$  (against  $I^2 = 0.49, 0.25, 0.09$ ; agreement to within Monte Carlo error  $\approx 0.02$ ) and its coverage collapses to 0.78, 0.64, 0.44. The error is largest exactly where censoring is heaviest and sequential methods are most needed.

**A two-parameter check: censored Weibull.** To go beyond  $k = 1$  we check the ellipsoidal rule of Remark 5 on  $X \sim \text{Weibull}(\lambda, \gamma)$  with  $\theta = (\lambda, \gamma)$  unknown and independent exponential censoring at a 25% fraction. With

$W = (\lambda y)^\gamma$ ,  $e = \delta - W$ , the censored score (1) is  $\psi_\lambda = \frac{\gamma}{\lambda}e$ ,  $\psi_\gamma = \frac{\delta}{\gamma} + e \ln(\lambda y)$ , verified against the Bartlett identity  $-\mathbb{E}[\partial_\theta \psi] = I$  by finite differences at  $(\lambda_0, \gamma_0) = (1, 2)$ ; the information is ill-conditioned,  $\lambda_{\max}(I)/\lambda_{\min}(I) \approx 9.3$ . We run (2) from a 300-observation pilot, gain  $c = 15$  ( $c\lambda_{\min}(I) \approx 4.9$ ), floor  $\iota = 10^{-6}$ , box projection, and the ellipsoidal rule with  $\gamma = 0.95$  ( $\chi_{2,0.95}^2 = 5.99$ ),  $M = 1500$ .

Table 4: Table 3. Two-parameter censored Weibull ( $\lambda_0 = 1$ ,  $\gamma_0 = 2$ , 25% censoring,  $c\lambda_{\min}(I) \approx 4.9$ ,  $M = 1500$ );  $n_\varepsilon = \chi_{2,0.95}^2 \lambda_{\max}(I^{-1})/\varepsilon^2$ .

| $\varepsilon$ | $n_\varepsilon$ | mean $N_\varepsilon$ | $N_\varepsilon/n_\varepsilon$ | coverage |
|---------------|-----------------|----------------------|-------------------------------|----------|
| 0.05          | 7381            | 7007                 | 0.95                          | 0.68     |
| 0.04          | 11533           | 11150                | 0.97                          | 0.76     |
| 0.03          | 20502           | 20132                | 0.98                          | 0.81     |
| 0.02          | 46130           | 45791                | 0.99                          | 0.89     |

As in the scalar case  $N_\varepsilon/n_\varepsilon \rightarrow 1$  regardless of coverage, confirming Remark 3 in the multivariate setting. Coverage approaches 0.95 more slowly, consistent with the larger finite-sample bias under a correlated, ill-conditioned  $I$ ; it worsens for larger  $c$  (e.g.  $c = 30, 45$  gave 0.37, 0.02), since a scalar gain tuned to  $\lambda_{\min}(I)$  over-drives the best-conditioned direction. The preconditioned recursion of Remark 5 is designed to remove this; a full study is left to future work.

**Practical guidance.** Choose  $c = \min\{\kappa/\hat{I}_0, c_{\max}\}$  with  $\kappa \geq 3$  from the pilot  $\hat{I}_0$  ( $c_{\max}$  guarding the boundary singularity of  $\psi$ ), floor  $\iota = 10^{-6}$ , and warm-up  $n_0 = \lceil \varepsilon^{-\alpha} \rceil$ ,  $\alpha \in (1, 2)$ . The gain may be self-tuned by  $c_n = \kappa/(\hat{I}_{n-1} \vee \iota)$ , so  $c_n I \rightarrow \kappa$  without knowing  $I$ . Where accurate coverage at moderate  $n$  is required, a larger gain or a bias correction to  $\bar{\theta}_n$  is advisable (Remark 4).

## 7. Discussion and open problems

The procedure enjoys asymptotic coverage consistency, first-order Chow-Robbins efficiency, and  $O(1)$  per-observation cost, under conditions verifiable for regular parametric lifetime models. A few points merit discussion. Efficiency and coverage are governed by different requirements (the simulation study), so a two-stage design—a pilot estimate of  $I$  then a refined  $c$  with  $cI \gg \frac{1}{2}$ —is natural. The projection (Assumption 4) is asymptotically inactive since  $\theta_n \rightarrow \theta_0 \in \text{int } K$ ; a generous  $K$  avoids boundary effects. The fourth-moment strengthening of Assumption 3 is needed only for the  $L^2$  control of  $P_n$ ; the martingale part  $D_n$  and consistency of  $\hat{I}_n$  need only finite variance and a  $2 + \eta$  moment, all finite for exponential and Weibull models with censoring bounded away from one.

**Sensitivity to model misspecification.** The procedure relies on a correctly specified family  $f(\cdot; \theta)$  in (1). If the true law  $P \notin \{P_\theta\}$  but  $M(\theta) = \mathbb{E}_P \psi(Y, \delta, \theta)$  still has a unique zero  $\theta^*$ , then  $\theta_n \xrightarrow{\text{a.s.}} \theta^*$ , the Kullback-Leibler projection of  $P$ ; but  $\hat{I}_n \rightarrow I^* = \mathbb{E}_P[\psi(\cdot, \theta^*)^2] \neq -M'(\theta^*)$ , so the correct variance is the sandwich  $\Sigma = M'(\theta^*)^{-1} I^* M'(\theta^*)^{-\top}$ , not  $I^{*-1}$ , and coverage is no longer protected by Theorem 2. Mitigations: enlarge the family, or use a bounded-influence (Huber-type) score—to which the martingale argument of Lemma 2 applies essentially unchanged. In practice, a goodness-of-fit test on the pilot sample can detect severe misspecification before the sequential phase begins, at which point the family can be enlarged or the sandwich variance  $\Sigma$  estimated online.

**Open problems.** (i) The exact coverage rate  $\gamma + O(\varepsilon^\beta)$  via nonlinear renewal theory [28], licensing bias-corrected intervals at moderate  $n$ ; (ii) the multivariate ellipsoidal procedure of Remark 5, with control of the least-informative direction; (iii) accelerated-failure-time and Cox regression on streams, where  $I$  may be ill-conditioned; (iv) a bounded-influence score variant; (v) a fully nonparametric online information estimate for unknown censoring.

## 8. Conclusion

We proposed and analysed a sequential fixed-width confidence procedure for right-censored data based on the averaged Robbins-Monro estimator and an online Fisher-information estimate. Under standard regularity and  $cI > \frac{1}{2}$  it is asymptotically consistent and Chow-Robbins efficient at  $O(1)$  per-observation cost. The

analysis isolates three transferable properties (H1)-(H3), corrects the linearisation step by a martingale-difference argument, and separates the two design requirements: stopping-time efficiency holds for every gain, while the gain condition controls coverage alone. Using the information rather than its reciprocal as the stopping scale is shown to collapse the procedure under heavy censoring.

### Appendix: Proof notes

*Monotonicity (Assumption 5).* Since  $M'(\theta_0) = -I < 0$  and  $M'$  is continuous,  $M' \leq -\lambda$  near  $\theta_0$  for any  $\lambda < I$ , and the fundamental theorem of calculus gives  $(\theta - \theta_0)M(\theta) \leq -\lambda(\theta - \theta_0)^2$ .

*Efficiency (Theorem 1).* Consistency is Robbins-Siegmund applied to  $(\theta_n - \theta_0)^2$ ; the projection binds finitely often since  $\theta_0 \in \text{int } K$  [10, Thm. 5.2.1]. The averaged CLT is Polyak-Juditsky [22, Thm. 2]/[10, Thm. 10.1.1] with drift  $\Gamma = -cI$ , noise  $S = c^2I$ , limiting covariance  $\Gamma^{-1}S\Gamma^{-1} = I^{-1}$ ; the rate  $\gamma_n = c/n$  needs  $cI > \frac{1}{2}$  [6, 18].

*Non-asymptotic MSE (Proposition 1).* With  $\pi_{n,j} = \prod_{k=j}^{n-1} (1 - cI/k) = (j/n)^{cI} (1 + O(1/j))$ , the variation-of-constants form of the linearised recursion has a deterministic transient contributing  $O(n^{-cI})$ , a stochastic term  $1/(In) + o(1/n)$ , and a nonlinearity remainder  $O(n^{-cI})$ ; squaring and summing gives (3), with a  $(\log n)^2$  factor at  $cI = \frac{1}{2}$ .

*Moment bound (Lemma 1).* With  $V_n = d_n^2$ ,  $\mathbb{E}[V_{n+1} \mid \mathcal{F}_n] \leq (1 - 2\lambda c/n)V_n + Cc^2/n^2(1 + V_n)$ ; for  $2\lambda c > 1$  the substitution  $Z_n = nV_n$  is contractive in expectation and Robbins-Siegmund gives  $\sup_n \mathbb{E}Z_n < \infty$ , the  $L^4$  bound analogously [18, Ch. 4].

**Competing interests.** The author has no conflicts of interest relevant to the content of this article.

### REFERENCES

1. Anscombe, F.J.: Large-sample theory of sequential estimation. *Proc. Camb. Philos. Soc.* **48**, 600-607 (1952)
2. Bickel, P.J., Klaassen, C.A.J., Ritov, Y., Wellner, J.A.: *Efficient and Adaptive Estimation for Semiparametric Models*. Johns Hopkins University Press, Baltimore (1993)
3. Chen, X., Lee, J.D., Tong, X.T., Zhang, Y.: Statistical inference for model parameters in stochastic gradient descent. *Ann. Statist.* **48**(1), 251-273 (2020)
4. Chow, Y.S., Robbins, H.: On the asymptotic theory of fixed-width sequential confidence intervals for the mean. *Ann. Math. Statist.* **36**, 457-462 (1965)
5. Durmus, A., Moulines, E., Naumov, A., Samsonov, S.: Finite-time high-probability bounds for Polyak-Ruppert averaged iterates of linear stochastic approximation. *arXiv:2207.04475* (2022)
6. Fabian, V.: On asymptotic normality in stochastic approximation. *Ann. Math. Statist.* **39**, 1327-1332 (1968)
7. Fang, Y., Xu, J., Yang, L.: Online bootstrap confidence intervals for the stochastic gradient descent estimator. *J. Mach. Learn. Res.* **19**(78), 1-21 (2018)
8. Gut, A.: *Stopped Random Walks: Limit Theorems and Applications*, 2nd edn. Springer, New York (2009)
9. Hájek, J.: A characterization of limiting distributions of regular estimates. *Z. Wahrscheinlichkeitstheorie Verw. Gebiete* **14**, 323-330 (1970)
10. Kushner, H.J., Yin, G.G.: *Stochastic Approximation and Recursive Algorithms and Applications*, 2nd edn. Springer, New York (2003)
11. Le Cam, L.: *Asymptotic Methods in Statistical Decision Theory*. Springer, New York (1986)
12. Le Cam, L., Yang, G.L.: *Asymptotics in Statistics: Some Basic Concepts*, 2nd edn. Springer, New York (2000)
13. Lee, S., Liao, Y., Seo, M.H., Shin, Y.: Fast and robust online inference with stochastic gradient descent via random scaling. In: *Proc. AAAI Conf. Artif. Intell.*, vol. 36, pp. 7381-7389 (2022)
14. Louis, T.A.: Finding the observed information matrix when using the EM algorithm. *J. R. Stat. Soc. B* **44**, 226-233 (1982)
15. Luo, L., Song, P.X.-K.: Renewable estimation and incremental inference in generalized linear models with streaming data sets. *J. R. Stat. Soc. B* **82**(1), 69-97 (2020)
16. Mou, W., Li, C.J., Wainwright, M.J., Bartlett, P.L., Jordan, M.I.: On linear stochastic approximation: fine-grained Polyak-Ruppert and non-asymptotic concentration. In: *Proc. COLT, PMLR* **125**, 2947-2997 (2020)

17. Moulines, E., Bach, F.: Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In: Advances in Neural Information Processing Systems, vol. 24 (2011)
18. Nevelson, M.B., Hasminskii, R.Z.: Stochastic Approximation and Recursive Estimation. American Mathematical Society, Providence (1973)
19. Orchard, T., Woodbury, M.A.: A missing information principle: theory and applications. In: Proc. 6th Berkeley Symp. Math. Statist. Probab., vol. 1, pp. 697–715 (1972)
20. Pelletier, M.: On the almost sure asymptotic behaviour of stochastic algorithms. Stochastic Process. Appl. **78**, 217-244 (1998)
21. Polyak, B.T.: New method of stochastic approximation type. Autom. Remote Control **51**, 937-946 (1990)
22. Polyak, B.T., Juditsky, A.B.: Acceleration of stochastic approximation by averaging. SIAM J. Control Optim. **30**, 838-855 (1992)
23. Robbins, H., Monro, S.: A stochastic approximation method. Ann. Math. Statist. **22**, 400-407 (1951)
24. Robbins, H., Siegmund, D.: A convergence theorem for non-negative almost supermartingales. In: Optimizing Methods in Statistics, pp. 233-257. Academic Press, New York (1971)
25. Ruppert, D.: Efficient estimation from a slowly convergent Robbins-Monro process. Tech. Rep. 781, Cornell University (1988)
26. Toulis, P., Airoldi, E.M.: Asymptotic and finite-sample properties of estimators based on stochastic gradients. Ann. Statist. **45**(4), 1694-1727 (2017)
27. Van der Vaart, A.W.: Asymptotic Statistics. Cambridge University Press, Cambridge (1998)
28. Woodroffe, M.: Nonlinear Renewal Theory in Sequential Analysis. SIAM, Philadelphia (1982)
29. Xu, J., et al.: Scalable estimation and inference with large-scale or online survival data. arXiv:2001.01434 (2020)

[Cite this article](#)

Boltaeva N. *Averaged Stochastic Approximation for Sequential Fixed-Width Confidence Intervals under Right Censoring*, **Acta NUUZ**, 2026, 2/3, pp. 59-67.