Maximum likelihood estimation for random walks in random environments

F. Comets, M. Falconnet, A. Gloter, O. Loukianov, D. Loukianova, C. Mattias

SAPS 2017, Peterhof

Outline

Framework: We consider a one-dimensional random walk in i.i.d. random environment (RWRE) with a parametric distribution.

Result: Based on a single observation of the path, we provide a maximum likelihood estimation procedure for the law of the environment.

Definitions

Random environment on $\mathbb{Z}$

- $\omega = \{\omega_x\}_{x\in\mathbb{Z}}$ i.i.d. with $\omega_x \in]0,1[$ and $\omega_x \sim \mu$,
- $\mathbb{P} = \mu^{\otimes\mathbb{Z}}$ law on $]0,1[^{\mathbb{Z}}$ of ω and $\mathbb{E}$ expectation

Markov process conditional on the environment

For *fixed* ω, let $X = \{X_t\}_{t\in\mathbb{N}}$ be the Markov chain on $\mathbb{Z}$ starting at $X_0 = 0$ and with transitions

$$P_\omega(X_{t+1} = y | X_t = x) = \begin{cases} \omega_x & \text{if } y = x+1 \\ 1-\omega_x & \text{if } y = x-1 \end{cases}$$

P_ω is the measure on the path space of X given ω (*quenched* law) and E_ω corresponding expectation.

Definitions

Random environment on $\mathbb{Z}$

- $\omega = \{\omega_x\}_{x\in\mathbb{Z}}$ i.i.d. with $\omega_x \in]0,1[$ and $\omega_x \sim \mu$,
- $\mathbb{P} = \mu^{\otimes\mathbb{Z}}$ law on $]0,1[^{\mathbb{Z}}$ of ω and $\mathbb{E}$ expectation

Markov process conditional on the environment

For *fixed* ω, let $X = \{X_t\}_{t\in\mathbb{N}}$ be the Markov chain on $\mathbb{Z}$ starting at $X_0 = 0$ and with transitions

$$P_\omega(X_{t+1} = y | X_t = x) = \begin{cases} \omega_x & \text{if } y = x+1 \\ 1-\omega_x & \text{if } y = x-1 \end{cases}$$

P_ω is the measure on the path space of X given ω (*quenched* law) and E_ω corresponding expectation.

Random walk in random environment (RWRE)

The (unconditional) law of X is the *annealed* law

$$\mathbf{P}(\cdot) = \mathbb{E}(P_\omega(\cdot)) = \int P_\omega(\cdot)\mathrm{d}\mathbb{P}(\omega),$$

with $\mathbf{E}$ the corresponding expectation.

Note that X is *not* a Markov process under $\mathbf{P}$ in general.

Properties of RWRE

Consider the "left/right" ratio

$$\rho_x = \frac{1-\omega_x}{\omega_x}, \qquad x \in \mathbb{Z}$$

Solomon(1975) has proved the classification:

Recurrent case

If $\mathbb{E}(\log \rho_0) = 0$, then

$$-\infty = \liminf_{t\to\infty} X_t < \limsup_{t\to\infty} X_t = +\infty, \quad \mathbf{P}\text{-a.s.}$$

and X_t is null-recurrent.

Transient case

if $\mathbb{E}(\log \rho_0) < 0$, then

$$\lim_{t \to \infty} X_t = +\infty, \quad \mathbf{P}\text{-a.s.}$$

Moreover, if $T_n = \inf\{t \in \mathbb{N} : X_t = n\}$, then

- **Ballistic case:** if $\mathbb{E}(\rho_0) < 1$, then $T_n/n \to c$ $\mathbf{P}$-a.s. when $n \to \infty$.
- **Sub-ballistic case:** If $\mathbb{E}(\rho_0) \geq 1$ and $\mathbb{E}(\rho_0^\kappa) = 1$ for some $0 < \kappa \leq 1$ then (in general) $T_n \sim n^{1/\kappa}$.

Statistical problem

Environment law estimation

Estimate μ from a single observation $(X_t)_{0\leq t\leq T}$ of a RWRE path.

Assumptions

We suppose that $\mu = \mu_{\theta^\star} \in \{\mu_\theta\}_{\theta\in\Theta}$, where $\theta^\star \in \Theta$ is an unknown parameter, $\Theta \subset \mathbb{R}^d$ compact.

Example (finitely supported law)

$$\mu_{(\{a_i\},\{p_i\})} = \sum_{i=1}^{m} p_i \delta_{a_i}, \quad \mathbb{E}\rho_0 = \sum_{i=1}^{m} p_i \log \frac{1-a_i}{a_i}.$$

Statistical problem

Environment law estimation

Estimate μ from a single observation $(X_t)_{0\le t\le T}$ of a RWRE path.

Assumptions

We suppose that $\mu = \mu_{\theta^\star} \in \{\mu_\theta\}_{\theta\in\Theta}$, where $\theta^\star \in \Theta$ is an unknown parameter, $\Theta \subset \mathbb{R}^d$ compact.

Example (finitely supported law)

$$\mu_{(\{a_i\},\{p_i\})} = \sum_{i=1}^{m} p_i \delta_{a_i}, \quad \mathbb{E}\rho_0 = \sum_{i=1}^{m} p_i \log \frac{1-a_i}{a_i}.$$

We write $\mathbb{P}^\theta$, $\mathbf{P}^\theta$ and so on for RWRE generated by μ_θ, and $\mathbb{P}^\star$, $\mathbf{P}^\star$, ... for $\theta = \theta^\star$ (the true parameter to estimate).

Statistical problem

Environment law estimation

Estimate μ from a single observation $(X_t)_{0\leq t\leq T}$ of a RWRE path.

Assumptions

We suppose that $\mu = \mu_{\theta^\star} \in \{\mu_\theta\}_{\theta\in\Theta}$, where $\theta^\star \in \Theta$ is an unknown parameter, $\Theta \subset \mathbb{R}^d$ compact.

Example (finitely supported law)

$$\mu_{(\{a_i\},\{p_i\})} = \sum_{i=1}^{m} p_i \delta_{a_i}, \quad \mathbb{E}\rho_0 = \sum_{i=1}^{m} p_i \log \frac{1-a_i}{a_i}.$$

We write $\mathbb{P}^\theta$, $\mathbf{P}^\theta$ and so on for RWRE generated by μ_θ, and $\mathbb{P}^\star$, $\mathbf{P}^\star$, $\ldots$ for $\theta = \theta^\star$ (the true parameter to estimate).

Moments estimation

O. Adelman and N. Enriquez (2004), Random walks in random environment: what a single trajectory tells.

A nice family of estimators of *moments* of $\mu_{\theta^\star}$.

Example: first steps from each site $\Rightarrow$ first moment.

Drawback:

- Which moments to estimate to recover $\mu_{\theta^\star}$?
- Only some steps are used (loss of information).

Maximum likelihood estimator

Fix a time T, a trajectory $X_{[0,T]}$, and let $L_x = L_x(T)$ and $R_x = R_x(T)$ be the number of left and right steps from site x. Then,

$$P_\omega(X_{[0,T]}) = \prod_{x\in\mathbb{Z}} \omega_x^{R_x}(1-\omega_x)^{L_x}$$

and

$$\mathbf{P}^\theta(X_{[0,T]}) = \mathbb{E}^\theta \prod_{x\in\mathbb{Z}} \omega_x^{R_x}(1-\omega_x)^{L_x} =$$

$$\prod_{x\in\mathbb{Z}} \mathbb{E}^\theta \omega_x^{R_x}(1-\omega_x)^{L_x} = \prod_{x\in\mathbb{Z}} \int_0^1 a^{R_x}(1-a)^{L_x}\mathrm{d}\mu_\theta(a).$$

Let ϕ_θ be the function from $\mathbb{N}^2$ to $\mathbb{R}$ given by

$$\phi_\theta(x, y) = \log \int_0^1 a^x (1-a)^y \mathrm{d}\mu_\theta(a).$$

The criterion function $\theta \mapsto \ell_T(\theta)$ is defined as

$$\ell_T(\theta) = \log \mathbf{P}^\theta(X_{[0,T]}) = \sum_{x \in \mathbb{Z}} \phi_\theta(R_x, L_x),$$

and our estimator is

$$\hat{\theta}_T \in \operatorname*{Argmax}_{\theta \in \Theta} \ell_T(\theta).$$

Aim

Study the convergence of $\hat{\theta}_T$ to $\theta^\star$.

Method: Show that $\ell_T(\theta)$ converges (after appropriate normalisation) to some $\ell(\theta)$ with

$$\operatorname*{Argmax}_{\theta\in\Theta} \ell(\theta) = \theta^\star$$

and apply classical M-estimation theory.

Question: where $\ell(\theta)$ comes from?

Transient case

We take $n \in \mathbb{N}$ and $T = T_n = \inf\{t \in \mathbb{N} : X_t = n\}$.

Note that

- Only the visited sites contribute to $\ell_{T_n}(\theta)$.
- The number of visited sites $x < 0$ is bounded (since X is transient to the right).
- Moreover, $R_x = L_{x+1} + 1$ for $x = 0, 1, \ldots, n-1$.

Hence

$$\ell_{T_n}(\theta) \approx \sum_{x=0}^{n-1} \phi_\theta(L_{x+1} + 1, L_x).$$

Underlying BPIRE I

Under $\mathbf{P}^\theta$, the sequence $L_n, L_{n-1}, \ldots, L_0$ has the same distribution as a BPI denoted $Z_0, \ldots, Z_n$, and defined by

$$Z_0 = 0 \quad \text{and} \quad Z_{k+1} = \sum_{i=0}^{Z_k} \xi_{k+1,i} \text{ for } k \geq 0,$$

with $\{\xi_{k,i}\}$ independent and

$$\forall m \in \mathbb{N}, \quad P_\omega(\xi_{k,i} = m) = (1 - \omega_k)^m \omega_k,$$

Underlying BPIRE II

Under $\mathbf{P}^\theta$, $\{Z_n\}$ is an irreducible positive recurrent homogeneous Markov chain with the transition kernel

$$Q_\theta(x,y) = \binom{x+y}{x} \int_0^1 a^{x+1}(1-a)^y \mathrm{d}\mu_\theta(a).$$

Consequence

$$\frac{1}{n}\ell_{T_n}(\theta) \sim \frac{1}{n}\sum_{k=0}^{n-1}\phi_\theta(Z_k+1, Z_{k+1}) \quad \text{under } \mathbf{P}^\star$$

and the right-hand side is (up to constants) the likelihood of a Markov process.

It follows that $\phi_\theta(Z_k+1, Z_{k+1})$ satisfies a law of large numbers.

We deduce that $\ell_{T_n}(\theta)/n$ converges in $\mathbf{P}^{\star}$-probability to a deterministic limit $\ell(\theta)$:

$$\ell(\theta) = \mathbf{E}^{\star}\phi_\theta(\tilde{Z}_0 + 1, \tilde{Z}_1)$$

Ballistic case

$\tilde{Z}_k$ has a finite first order moment.

Sub-ballistic case

We fix $\theta_0 \in \Theta$ and replace $\ell_{T_n}(\theta)$ with

$$\ell_{T_n}(\theta) - \ell_{T_n}(\theta_0) \sim \sum_{k=0}^{n-1} \left(\phi_\theta(Z_k + 1, Z_{k+1}) - \phi_{\theta_0}(Z_k + 1, Z_{k+1})\right)$$

and assume that $\phi_\theta - \phi_{\theta_0}$ is uniformly integrable (true in most cases).

Using the almost linear nature of ϕ_θ, we prove that $\ell(\theta)$ is finite, with a maximum at $\theta^{\star}$.

Results: consistency, asymptotic normality and efficiency

The standard M-estimators theory then applies. Under appropriate (classical) assumptions, in the *transient* case, we establish that MLE satisfies

- **Consistency**: $\lim_{n\to+\infty} \hat{\theta}_{T_n} = \theta^\star, \quad \mathbf{P}^\star$-a.s.
- **Asymptotic normality**: $\sqrt{n}(\hat{\theta}_{T_n} - \theta^\star) \rightsquigarrow^{\mathbf{P}^\star - dist.} \mathcal{N}(0, \Sigma_{\theta^\star}^{-1})$.
- **Efficiency**: $\Sigma_{\theta^\star}$ is the Fisher information matrix.

Hence the rate of convergence is of the order $\sqrt{T}$ in the ballistic case, and $T^{\kappa/2}$ in the sub-ballistic case $(\kappa \leq 1)$.

Recurrent case

We consider the distributions of the form

$$\mu_{(\mathbf{a},\mathbf{p})} = \sum_{i=1}^{d} p_i \delta_{a_i}$$

and assume that the true parameter $\theta^\star = (\mathbf{a}^\star, \mathbf{p}^\star)$ belongs to a compact $\Theta \subset (0,1)^{2d}$ satisfying

Assumption (Identifiability)

For any $\theta = (\mathbf{a}, \mathbf{p})$ in Θ,

$$0 < a_1 < a_2 < \ldots < a_d < 1$$

Assumption (Recurrent environment)

$$\mathbb{E}^\star \rho_0 = \sum_{i=1}^{d} p_i^\star \log \frac{1 - a_i^\star}{a_i^\star} = 0.$$

Example (Temkin)

Let

$$\mu_\theta = \frac{1}{2}\delta_a + \frac{1}{2}\delta_{1-a}.$$

Here, the unknown parameter is $a \in \Theta \subset (0, 1/2)$.

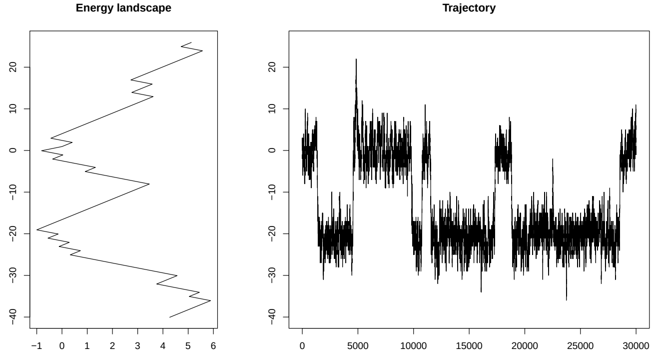

Properties of recurrent RWRE

An important property of recurrent RWRE is *localization*: the RE creates *traps* where the walk spends much time.

A useful trap visualization is the potential landscape V where $V = \{V(x) : x \in \mathbb{Z}\}$ is defined by

$$V(x) = \begin{cases} \sum_{y=0}^{x} \log \rho_y - \log \rho_0 & \text{if } x \geq 0 \\ -\sum_{y=x+1}^{0} \log \rho_y & \text{if } x < 0 \end{cases}$$

The environment $\{\omega_x\}$ can be recovered from its potential:

$$\omega_x = \frac{\exp(-V(x))}{\exp(-V(x)) + \exp(-V(x-1))}.$$

Main valleys

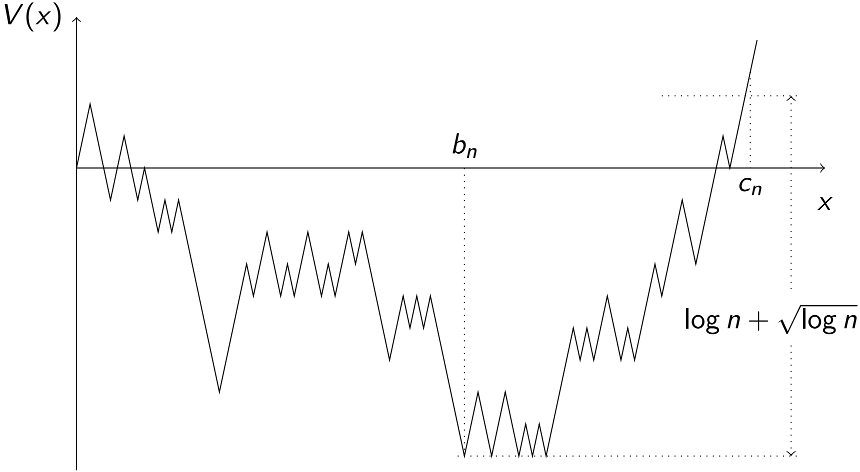

Figure :
$c_n = \min\left\{x \geq 0 \,:\, V(x) - \min_{0 \leq y \leq x} V(y) \geq \log n + (\log n)^{1/2}\right\}$,
$b_n = \min\{x \geq 0 \,:\, V(x) = \min_{0 \leq y \leq c_n} V(y)\}$.

Classical localization results

Basic localization properties of the RW are known since the works of Sinai (1982), Golosov (1984) and others. Namely,

- with an overwhelming probability, the (reflected) walk $X_{[0,n]}$ stays between 0 and c_n (Arrhenius law);
- $b_n/\log^2 n$ and $c_n/\log^2 n$ converge in law to some non-degenerate random variables;
- moreover, $(X_n - b_n)/\log^2 n$ converges to 0 in probability.

For MLE, we need to describe the distribution of local times of $X_n - b_n$.

Infinite valley

Let $\widetilde{V} = \{\widetilde{V}(x) : x \in \mathbb{Z}\}$ be a collection of random variables distributed as V "conditioned" to stay positive. For each $\widetilde{V}$, let $\widetilde{\omega}$ be the corresponding environment on $\mathbb{Z}$.

Let $\nu(x) = \nu^+(x) + \nu^-(x)$ be the invariant measure of the corresponding (ergodic) Markov chain $\widetilde{X}_n$ on $\mathbb{Z}$, where

$$\nu^+(x) = \frac{e^{-\widetilde{V}(x)}}{2\sum_{z\in\mathbb{Z}} e^{-\widetilde{V}(z)}} \quad \text{and} \quad \nu^-(x) = \frac{e^{-\widetilde{V}(x-1)}}{2\sum_{z\in\mathbb{Z}} e^{-\widetilde{V}(z)}}.$$

Then $\widetilde{\omega}_x = \nu^+(x)/\nu(x)$.

Remark

The possible values of $\widetilde{\omega}$ are those of ω, though their distributions are different (not i.i.d.)

Gantert-Peres-Shi theorem

Put $\nu_n^+(x) = R_x/n$, $\nu_n^-(x) = L_x/n$.

Theorem (Gantert-Peres-Shi, 2010)

The distributions of

$$\{(\nu_n^+(x+b_n), \nu_n^-(x+b_n)) \,:\, x \in \mathbb{Z}\}$$

converge weakly to the distribution of

$$\{(\nu^+(x), \nu^-(x)) \,:\, x \in \mathbb{Z}\}.$$

As a consequence, for each strongly continuous functional f which is translation invariant, we have

$$f\left(\{(\nu_n^+, \nu_n^-)\}\right) \xrightarrow[n\to\infty]{law} f\left(\{(\nu^+, \nu^-)\}\right).$$

MLE

The annealed log-likelihood $\ell_n(\theta) = \log \mathbf{P}^\theta(X_{[0,n]})$ in our case is given by

$$\ell_n(\theta) = \sum_{x \in \mathbb{Z}} \log \left[\sum_{i=1}^{d} a_i^{R_x} (1 - a_i)^{L_x} p_i \right]$$

Recurrence imply $R_x, L_x \to \infty$ as $n \to \infty$, so

- the branching explodes;
- but we can apply Laplace methods.

Denote by $\mathcal{R}_n$ the range of the walk:

$$\mathcal{R}_n = \{x \ : \ \exists t \le n, \ X_t = x\}$$

Recall that $|\mathcal{R}_n| = O_{\mathbb{P}^\star}(\log^2 n)$.

Log-likelihood decomposition

For any x in $\mathcal{R}_n$, define the random integer

$$\hat{\imath} = \hat{\imath}(\mathbf{a}, n, x) = \operatorname*{Argmax}_i \left\{ a_i^{R_x} (1 - a_i)^{L_x} \right\}$$

Then

$$\ell_n(\theta) = \sum_{x \in \mathbb{Z}} \left(R_x \log a_{\hat{\imath}} + L_x \log(1 - a_{\hat{\imath}}) \right) + \sum_{x \in \mathcal{R}_n} \log p_{\hat{\imath}}$$
$$+ \sum_{x \in \mathcal{R}_n} \log \left(1 + \sum_{i \neq \hat{\imath}} \left(\frac{a_i}{a_{\hat{\imath}}} \right)^{R_x} \left(\frac{1 - a_i}{1 - a_{\hat{\imath}}} \right)^{L_x} \frac{p_i}{p_{\hat{\imath}}} \right)$$
$$= M_n + K_n + r_n = O(n) + O_{\mathbb{P}^\star}(\log^2 n) + o_{\mathbb{P}^\star}(\log^2 n)$$

Estimators

We define a MLE as

$$\widehat{\theta}_n = (\widehat{\mathbf{a}}_n, \widehat{\mathbf{p}}_n) = \operatorname*{Argmax}_{(\mathbf{a},\mathbf{p})\in\Theta} \ell_n(\theta),$$

and a (pseudo) MLE $(\overline{\mathbf{a}_n}, \overline{\mathbf{p}_n})$ as

$$\begin{cases} \overline{\mathbf{a}_n} & = & \operatorname*{Argmax}_{\mathbf{a}\in\Theta_a} M_n(\mathbf{a}), \\ \overline{\mathbf{p}_n} & = & \operatorname*{Argmax}_{\mathbf{p}} K_n(\overline{\mathbf{a}_n}, \mathbf{p}). \end{cases}$$

Theorem

Both the the MPL estimator $(\overline{\mathbf{a}_n}, \overline{\mathbf{p}_n})$ and ML estimator $(\widehat{\mathbf{a}}_n, \widehat{\mathbf{p}}_n)$ converge in $\mathbf{P}^\star$-probability to the true parameter value $\theta^\star$.

Estimators

We define a MLE as

$$\widehat{\theta}_n = (\widehat{\mathbf{a}}_n, \widehat{\mathbf{p}}_n) = \operatorname*{Argmax}_{(\mathbf{a},\mathbf{p})\in\Theta} \ell_n(\theta),$$

and a (pseudo) MLE $(\overline{\mathbf{a}_n}, \overline{\mathbf{p}_n})$ as

$$\begin{cases} \overline{\mathbf{a}_n} & = & \operatorname*{Argmax}_{\mathbf{a}\in\Theta_a} \; M_n(\mathbf{a}), \\ \overline{\mathbf{p}_n} & = & \operatorname*{Argmax}_{\mathbf{p}} K_n(\overline{\mathbf{a}_n}, \mathbf{p}). \end{cases}$$

Theorem

Both the the MPL estimator $(\overline{\mathbf{a}_n}, \overline{\mathbf{p}_n})$ and ML estimator $(\widehat{\mathbf{a}}_n, \widehat{\mathbf{p}}_n)$ converge in $\mathbf{P}^\star$-probability to the true parameter value $\theta^\star$.

Cross-entropy

Recall the properties of the *cross-entropy* $H(\mathbf{p}, \mathbf{q})$ of two (finitely supported) probability measures:

$$H(\mathbf{p}, \mathbf{q}) = -E_{\mathbf{p}} \log \mathbf{q} = -\sum_i p_i \log q_i$$

$$H(\mathbf{p}) = H(\mathbf{p}, \mathbf{p}) < H(\mathbf{p}, \mathbf{q}) \quad \text{if} \quad \mathbf{p} \neq \mathbf{q}$$

In particular, for $0 < p, q < 1$

$$\max_q \{p \log q + (1-p) \log(1-q)\}$$
$$= p \log p + (1-p) \log(1-p)$$

Support estimation

Recall that $\nu_n^+(x) = R_x/n$, $\nu_n^- = L_x/n$ and $\widetilde{\omega} = \nu^+/\nu$. So

$$M_n = n \sum_{x \in \mathbb{Z}} \max_i \left\{ \nu_n^+(x) \log a_i + \nu_n^-(x) \log(1 - a_i) \right\}$$

and GPS theorem yields

$$\frac{M_n}{n} \xrightarrow[n\to\infty]{\text{law}} M(\mathbf{a}, \nu^+, \nu^-)$$
$$= \sum_{x \in \mathbb{Z}} \nu(x) \max_i \{\widetilde{\omega}_x \log a_i + (1 - \widetilde{\omega}_x) \log(1 - a_i)\}$$
$$= -\sum_{x \in \mathbb{Z}} \nu(x) \max_i H\left(\widetilde{\omega}_x, a_i\right)$$

Using $\widetilde{\omega}_x \in \{a_i^\star\}$, it's easily seen that for $\mathbf{a} \neq \mathbf{a}^\star$,

$$M(\mathbf{a}) < M(\mathbf{a}^\star) = -\sum_{x\in\mathbb{Z}} \nu(x) H\left(\widetilde{\omega}_x\right).$$

Finally,

$$\frac{M_n(\mathbf{a}^\star) - M_n(\mathbf{a})}{n} \xrightarrow[n\to\infty]{\text{law}} M(\mathbf{a}^\star) - M(\mathbf{a}) > 0$$

whence it can be deduced that

$$\operatorname*{Argmax}_{\mathbf{a}\in\Theta} \; M_n(\mathbf{a}) = \overline{\mathbf{a}_n} \to \mathbf{a}^\star.$$

Probability estimation

Recall the definition of $\overline{\mathbf{p}_n}$:

$$\overline{\mathbf{p}_n} = \operatorname*{Argmax}_{\mathbf{p}} \sum_{x \in \mathcal{R}_n} \log p_{\hat{\imath}} = \operatorname*{Argmax}_{\mathbf{p}} \sum_{i=1}^{d} \frac{|\mathcal{R}_n(i)|}{|\mathcal{R}_n|} \log p_i,$$

where

$$\hat{\imath} = \operatorname*{Argmax}_{i} \{R_x \log(\overline{\mathbf{a}_n})_i + L_x \log(1 - (\overline{\mathbf{a}_n})_i)\}$$

and $\mathcal{R}_n(i) = \{x \in \mathcal{R}_n \ : \ \hat{\imath} = i\}$.

By the law of large numbers, $R_x/(R_x + L_x) \to \omega_x$, hence $\omega_x = a^\star_{\hat{\imath}}$ for $\mathbf{a} \approx \mathbf{a}^\star$ if n is large enough.

Probability estimation

Recall the definition of $\overline{\mathbf{p}_n}$:

$$\overline{\mathbf{p}_n} = \operatorname*{Argmax}_{\mathbf{p}} \sum_{x \in \mathcal{R}_n} \log p_{\hat{\imath}} = \operatorname*{Argmax}_{\mathbf{p}} \sum_{i=1}^{d} \frac{|\mathcal{R}_n(i)|}{|\mathcal{R}_n|} \log p_i,$$

where

$$\hat{\imath} = \operatorname*{Argmax}_{i} \{R_x \log(\overline{\mathbf{a}_n})_i + L_x \log(1 - (\overline{\mathbf{a}_n})_i)\}$$

and $\mathcal{R}_n(i) = \{x \in \mathcal{R}_n \ : \ \hat{\imath} = i\}$.

By the law of large numbers, $R_x/(R_x + L_x) \to \omega_x$, hence $\omega_x = a^\star_{\hat{\imath}}$ for $\mathbf{a} \approx \mathbf{a}^\star$ if n is large enough.

Since $\overline{\mathbf{a}_n} \to \mathbf{a}^\star$, we get $|\mathcal{R}_n(i)| \to \#\{x \in \mathcal{R}_n \ : \ \omega_x = a^\star_i\}$, whence $|\mathcal{R}_n(\cdot)|/|\mathcal{R}_n| = \overline{\mathbf{p}_n} \to \mathbf{p}^\star$.

Probability estimation

Recall the definition of $\overline{\mathbf{p}_n}$:

$$\overline{\mathbf{p}_n} = \operatorname*{Argmax}_{\mathbf{p}} \sum_{x \in \mathcal{R}_n} \log p_{\hat{\imath}} = \operatorname*{Argmax}_{\mathbf{p}} \sum_{i=1}^{d} \frac{|\mathcal{R}_n(i)|}{|\mathcal{R}_n|} \log p_i ,$$

where

$$\hat{\imath} = \operatorname*{Argmax}_{i} \left\{ R_x \log(\overline{\mathbf{a}_n})_i + L_x \log(1 - (\overline{\mathbf{a}_n})_i) \right\}$$

and $\mathcal{R}_n(i) = \{x \in \mathcal{R}_n \ : \ \hat{\imath} = i\}$.

By the law of large numbers, $R_x/(R_x + L_x) \to \omega_x$, hence $\omega_x = a^\star_{\hat{\imath}}$ for $\mathbf{a} \approx \mathbf{a}^\star$ if n is large enough.

Since $\overline{\mathbf{a}_n} \to \mathbf{a}^\star$, we get $|\mathcal{R}_n(i)| \to \#\{x \in \mathcal{R}_n \ : \ \omega_x = a^\star_i\}$, whence $|\mathcal{R}_n(\cdot)|/|\mathcal{R}_n| = \overline{\mathbf{p}_n} \to \mathbf{p}^\star$.

Simulations

Example (Temkin)

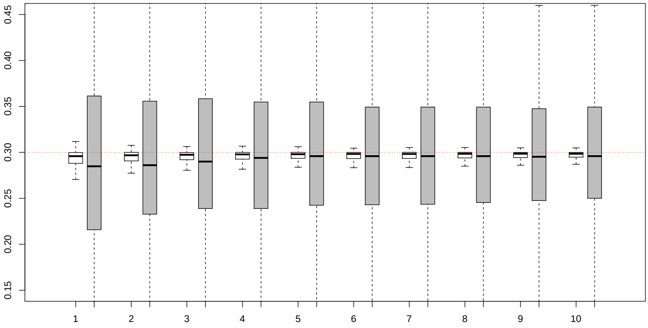

Figure : Boxplots of our estimator (white) and Adelman and Enriquez estimator (grey). The true value of $\theta^\star$ is 0.3.

Extensions

- P. Andreoletti, D. Loukianova, C. Matias (2015): Hidden Markov model for parameter estimation of a random walk in a Markov environment.
- R. Diel, M. Lerasle (2016): Non parametric estimation for random walks in random environment.