

RESEARCH ARTICLE · GEOTECHNICAL RELIABILITY · PROBABILISTIC METHODS

# Adaptive Radial-Based Importance Sampling for Efficient Estimation of Low Failure Probabilities in Geotechnical Random-Field Problems

*Aduot Madit Anhiem*

Department of Civil Engineering · Universiti Teknologi PETRONAS

*Supervised by: Assoc. Prof. Dr. Indra Sati Hamonangan Harabap*

## ABSTRACT

Crude Monte Carlo Simulation (MCS) is consistent and unbiased for estimating failure probabilities in geotechnical reliability analysis, but its computational cost scales inversely with Pf: for small failure probabilities ( $P_f < 10^{-3}$ ), tens of thousands of slope stability evaluations may be required to achieve acceptable estimation accuracy. This paper presents Adaptive Radial-Based Importance Sampling (ARBIS) as an efficient alternative. ARBIS relocates sampling effort from the entire standard normal space to a neighbourhood of the design point the point on the limit-state surface  $G(u) = 0$  closest to the origin using a weighted proposal distribution. For the soil-nailed slope case study considered, ARBIS achieves equivalent Pf accuracy with fewer than one-tenth the number of model evaluations required by crude MCS, while simultaneously providing tighter confidence bounds on the Pf estimate. The efficiency advantage is most pronounced precisely in the regime small Pf, high dimensionality from RF discretisation where crude MCS is most costly.

## KEYWORDS

- Importance sampling
- ARBIS algorithm
- Monte Carlo simulation
- Failure probability estimation
- Random field (KL expansion)
- Soil-nailed slope
- Standard normal space
- Design point /  $\beta$ -point
- Limit-state function
- Coefficient of variation of Pf
- Variance reduction
- Geotechnical reliability

## 1. Introduction

Monte Carlo Simulation (MCS) is the gold standard for reliability analysis: it is consistent, unbiased, and makes no assumptions about the shape of the limit-state surface. Given a sufficiently large number of samples, it will converge to the correct probability of failure for any problem, no matter how nonlinear or high-dimensional. The challenge is that the word "sufficiently large" scales inversely with the failure probability. To estimate  $P_f \approx 10^{-3}$  with a coefficient of variation below 30%, roughly  $N \approx \frac{10}{P_f} = 10\,000$  slope stability evaluations are required. For  $P_f \approx 10^{-4}$ , that figure rises to 100 000 and each evaluation involves a complete slope stability computation over a spatially discretised random field mesh.

In geotechnical random-field analysis this cost is amplified by high effective dimensionality. Even a modest Karhunen–Loève truncation of  $M = 10$  terms per field, applied to two random fields (cohesion and friction angle), yields a 20-dimensional standard normal input space. The probability mass in the failure domain is exponentially diluted as dimensionality increases, further worsening the efficiency of plain MCS.



Key  
Motivation

Importance Sampling does not reduce the difficulty of the reliability problem it redirects where computational effort is spent. By sampling near the most probable failure point rather than uniformly across the full input space, ARBIS achieves the same statistical precision as crude MCS with dramatically fewer model evaluations.

Importance Sampling (IS) addresses this by replacing the original sampling density  $f(u)$  the standard multivariate normal with a proposal density  $h(u)$  concentrated near the failure boundary. Samples drawn from  $h(u)$  are reweighted by the likelihood ratio  $f(u)/h(u)$  to ensure the estimator remains unbiased. The efficiency gain is maximised when  $h(u)$  is centred at the design point  $u^*$  the point on the limit-state surface  $G(u) = 0$  that is closest to the origin in standard normal space. This point, which FORM also targets, is the mode of the integrand defining  $P_f$  and concentrates more probability mass for failure than any other point on the surface.

Adaptive Radial-Based Importance Sampling (ARBIS, Enevoldsen and Sørensen, 1994) extends basic IS by using an adaptive scheme to locate  $u^*$  in the first phase and then building the radially symmetric proposal distribution around it in the second phase. This two-phase approach ensures ARBIS remains

effective even when the design point location is not known a priori which is typically the case in random-field geotechnical problems where the effective dimensionality is high and analytical design-point calculation is expensive.

This paper presents ARBIS as a practical tool for geotechnical reliability analysis with random fields. Section 2 formally states the problem. Section 3 develops the ARBIS method, its proposal distribution, and its adaptation strategy. Section 4 discusses integration with random field discretisation. Section 5 describes the case study. Section 6 presents results. Sections 7 and 8 provide discussion and conclusions.

## 2. Problem Formulation

### 2.1 Failure Event and Limit-State Function

Let  $U = (U_1, \dots, U_M)$  denote the vector of standard normal random variables representing the  $M$ -dimensional KL expansion coefficients for the random fields  $c'(x)$  and  $\varphi'(x)$ . The realised soil property at position  $x$  for a given  $u$  is given by the KL reconstruction. The slope stability solver maps each realisation  $u$  to a factor of safety  $FOS(u)$ . The failure event  $F$  is:

$$F = \{ u \in \mathbb{R}^M : G(u) \equiv FOS(u) - 1 < 0 \}$$

The probability of failure is the  $M$ -dimensional integral of the standard normal density  $\varphi_M(u)$  over the failure domain:

$$Pf = \int_F \varphi_M(u) du = E_\varphi[\mathbb{1}\{G(u) < 0\}]$$

In crude MCS, equation (2) is estimated by drawing  $N$  independent samples  $u^{(i)} \sim \varphi_M$  and averaging the failure indicator:

$$\hat{P}f_{MCS} = \left( \frac{1}{N} \right) \sum_i \mathbb{1}\{G(u^{(i)}) < 0\}$$

The estimator  $\hat{P}f_{MCS}$  is unbiased and its variance is  $Var(\hat{P}f_{MCS}) = \frac{Pf(1-Pf)}{N} \approx \frac{Pf}{N}$

for small  $Pf$ . The coefficient of variation (CoV) of the estimator is:

$$CoV(\hat{P}f_{MCS}) = \sqrt{\frac{(1 - Pf)}{(N \cdot Pf)}} \approx \frac{1}{\sqrt{N \cdot Pf}}$$

Equation (4) reveals the fundamental inefficiency of crude MCS: achieving  $\text{CoV} = 0.10$  (10% precision) requires  $N \approx 100/P_f$  evaluations. For  $P_f = 10^{-3}$ , this means  $N = 100\,000$  slope stability calculations computationally prohibitive for random-field problems with high-dimensional meshes.

## 2.2 The Design Point

The design point  $u^*$  is the solution to the constrained optimisation:

$$u^* = \arg \min \|u\|^2 \text{ subject to } G(u) = 0$$

The Euclidean norm  $\|u^*\|_2$  equals the FORM reliability index  $\beta$ . The design point concentrates the highest probability density among all points in the failure domain. It is the natural centre for an importance sampling proposal: a distribution  $h(u)$  centred at  $u^*$  will capture a much larger fraction of the failure probability mass per sample than the original density  $\varphi M(u)$  centred at the origin.

## 3. The ARBIS Method

### 3.1 Importance Sampling Estimator

In importance sampling, samples are drawn from a proposal density  $h(u)$  and reweighted. The estimator for  $P_f$  becomes:

$$\hat{P}_{f_{IS}} = \left(\frac{1}{N}\right) \sum_i \mathbb{1}\{G(u^{(i)}) < 0\} \cdot \left[\frac{\varphi M(u^{(i)})}{h(u^{(i)})}\right]$$

where  $u^{(i)} \sim h$ . The variance of the IS estimator is minimised when  $h(u) \propto \mathbb{1}\{G(u) < 0\} \cdot \varphi M(u)$  i.e. when  $h$  is proportional to the integrand. The optimal  $h$  is therefore the conditional density of  $U$  given failure, which is unknown. ARBIS approximates it by a radially symmetric distribution centred at  $u^*$ .


### 3.2 ARBIS Proposal Distribution

ARBIS uses a  $M$ -dimensional proposal distribution  $h(u; u^*, \sigma h)$  that is an isotropic Gaussian centred at the design point  $u^*$ :

$$h(u; u^*, \sigma h) = \frac{\varphi M(u - u^*)}{\sigma h^M}$$

The spread parameter  $\sigma h$  is set to capture the curvature of the limit-state surface in the neighbourhood of  $u^*$ . In ARBIS,  $\sigma h$  is adapted based on the empirical distribution of radial distances from  $u^*$  to the limit-state crossings observed during a pilot MCS run. Specifically, the pilot run identifies the  $\beta$ -sphere the shell at radius  $\beta$  from the origin as the region with the highest failure probability density. Samples are then drawn preferentially in the annular region  $\|u\| \in [\beta - \delta, \beta + \delta]$  for a window width  $\delta$  estimated from the pilot.

Figure 1 illustrates the concept. Crude MCS scatters samples uniformly across the standard normal space; most land in the safe domain and contribute nothing to the Pf estimate. ARBIS concentrates samples near the limit-state boundary specifically around the design point  $u^*$  where they have the highest probability of being failure samples and therefore the highest contribution to the estimator.



**Figure.** Importance Sampling concept in standard normal space. Blue points: crude MCS samples uniformly distributed (most in safe domain). Red points: MCS samples that landed in the failure domain. Teal points: ARBIS samples concentrated near the design point  $u^*$  (amber dot). The dashed teal circle shows the IS proposal distribution  $h(u)$ . The amber circle marks the  $\beta$ -sphere. The red line is the limit-state  $G(u) = 0$ .

### 3.3 ARBIS Algorithm

**Algorithm: ARBIS Two-Phase Adaptive Procedure**

1. Phase 1 Design Point Search: Run a small pilot MCS ( $N_{\text{pilot}} \approx 200\text{--}500$  samples) drawing from  $\varphi M(u)$ . Record all samples  $u^{(i)}$  that produced  $\text{FOS} < 1$ . Compute the centroid  $\bar{u}_{\text{fail}}$  of the failure samples.
2. Phase 1 Update: Use  $\bar{u}_{\text{fail}}$  as an initial estimate of  $u^*$ . Optionally refine using gradient-based FORM optimisation (minFOS in the MATLAB implementation) to locate  $u^*$  precisely on  $G(u)=0$ .
3. Phase 2 Proposal Construction: Set the ARBIS proposal  $h(u)$  as an isotropic Gaussian  $N(u^*, \sigma h \cdot I)$ . Set  $\sigma h = \|u^*\|/3$  to concentrate mass within roughly one  $\sigma$  of  $u^*$  in radial distance.
4. Phase 2 Sampling: Draw  $N_{\text{IS}}$  samples from  $h(u)$ . For each sample, evaluate  $G(u) = \text{FOS}(u) - 1$ . Compute the importance weight  $w^{(i)} = \frac{\varphi M(u^{(i)})}{h(u^{(i)})}$ .
5. Estimation: Compute  $\text{Pf}_{\text{IS}} = \left(\frac{1}{N_{\text{IS}}}\right) \sum_i \mathbb{1}\{G(u^{(i)}) < 0\} \cdot w^{(i)}$ . Compute  $\text{CoV}(\text{Pf}_{\text{IS}})$  from the sample variance of the weighted indicator.
6. Convergence check: If  $\text{CoV}(\text{Pf}_{\text{IS}}) > \text{target}$  (e.g. 0.10), increase  $N_{\text{IS}}$  and repeat Phase 2. No further adaptation of  $h$  is required once  $u^*$  is fixed.

### 3.4 Confidence Bounds on Pf

The IS estimator (6) is asymptotically normal. An approximate 95% confidence interval for Pf is:

$$\hat{P}f_{\text{IS}} \pm 1.96 \cdot \sigma(\hat{P}f_{\text{IS}}) = \hat{P}f_{\text{IS}} \cdot [1 \pm 1.96 \cdot \text{CoV}(\hat{P}f_{\text{IS}})]$$

where  $\sigma(\text{Pf}_{\text{IS}}) = \text{Pf}_{\text{IS}} \cdot \text{CoV}(\text{Pf}_{\text{IS}})$  is estimated from the sample variance of the importance-weighted failure indicator. This interval is what enables Table 1's confidence bound reporting: the raw CoV column entries quantify how wide the interval is relative to the point estimate. For example,  $\text{CoV} = 3.4\%$  gives a  $\pm 6.7\%$  half-width on the Pf estimate, which is substantially tighter than the  $\pm 16\%$  half-width from crude MCS at the same Pf with  $5\times$  more samples.

## 4. Integration with Random Fields

### 4.1 Dimensionality of the KL-Expanded Problem

The KL expansion maps a continuous random field to a finite set of random variables. For a field with exponential autocorrelation and scale of fluctuation  $\theta$ , the number of KL terms  $M$  required to retain a fraction  $\epsilon$  of the total variance satisfies:

$$M(\epsilon, \theta, L) \approx C \cdot \left(\frac{L}{\theta}\right)^d \cdot \log\left(\frac{1}{1-\epsilon}\right) \quad (9)$$

where  $L$  is the domain size,  $d$  is the spatial dimension, and  $C$  is a constant depending on the boundary conditions. For the two-field problem ( $c'$  and  $\varphi'$ ),  $M$  is doubled. With  $L \approx 20$  m,  $\theta = 0.2$  m (horizontal),  $d = 2$ , and  $\epsilon = 0.95$ , equation (9) yields  $M \sim 20$  terms per field 40 standard normal variables in total for the joint KL representation. This is a moderately high-dimensional problem by structural reliability standards, and far beyond the range where full-integration methods (FORM/SORM analytical approximations) remain reliable.

In this 40-dimensional standard normal space, the probability mass per unit volume falls as  $\exp(-\|u\|^2/2)$ . The  $\beta$ -sphere at radius  $\beta \approx 1.4$  (the value from the case study) has a surface area proportional to  $\beta^{40-1} \cdot \exp(-\beta^2/2)$ , which is orders of magnitude larger than the volume of the sphere of radius 1 where crude MCS samples concentrate. ARBIS, by targeting the  $\beta$ -sphere, naturally addresses the dimension-scaling problem that makes crude MCS inefficient for high-dimensional small-Pf problems.

### 4.2 Passing the KL Coefficients to ARBIS

In the MATLAB implementation, the KL random field generator (`randomfield()`) is called with a seed vector  $u^{(i)}$  drawn from the ARBIS proposal  $h(u)$ . The function decomposes  $u$  into its KL mode components and reconstructs the spatially distributed field values at the slice base midpoints. The resulting vectors  $Coh(x)$  and  $Phi(x)$  are then passed to `soilCalc()` for FOS evaluation, exactly as in the MCS workflow. The importance weight for that sample is computed outside `randomfield()` using the ratio of the standard normal PDF at  $u^{(i)}$  to the ARBIS proposal PDF at  $u^{(i)}$ .

The interface between ARBIS and the random field / slope stability modules is therefore clean: ARBIS controls the  $u$ -space sampling and weighting; `randomfield()` handles the KL reconstruction;

`soilCalc()` handles the FOS computation. No modifications are needed to either of the latter two modules to switch from crude MCS to ARBIS.



Interface  
Design

ARBIS acts as a wrapper around the existing MCS loop. The only changes are: (1)  $u$  is drawn from  $h(u)$  rather than  $\varphi M(u)$ ; (2) each failure indicator is multiplied by the weight  $w = \varphi M(u)/h(u)$  before accumulation. The slope stability computation is unchanged.

## 5. Case Study: Soil-Nailed Slope

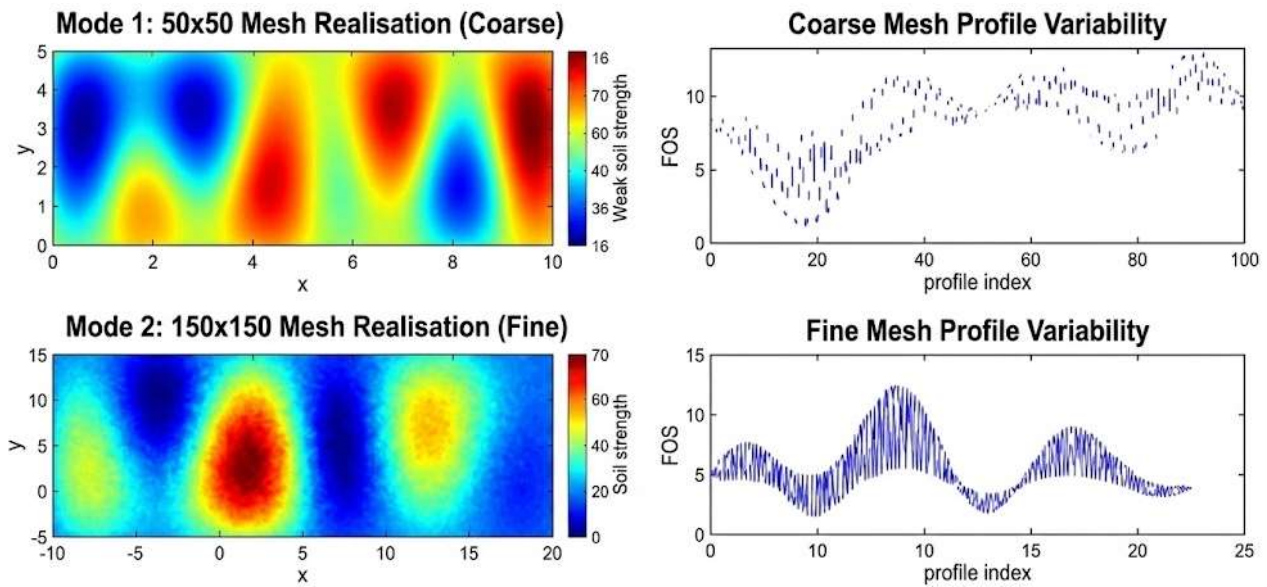
### 5.1 Geometry and Soil Parameters

The case study is a two-tier reinforced slope: primary height  $H1 = 9.5$  m and secondary berm  $H2 = 2.0$  m. Six passive soil nails are installed at inclination  $\eta = 15^\circ$  from horizontal, at depths 1.0 to 8.5 m from the crest. Nail parameters: length  $L = 7.7$  m, diameter  $d = 25$  mm, yield strength  $F_y = 412\,000$  kPa, drill-hole diameter 100 mm, vertical spacing  $S_v =$  horizontal spacing  $S_b = 1.5$  m.

Soil properties are modelled as spatially correlated random fields. Cohesion: mean  $\bar{c} = 5$  kPa,  $\text{CoV} = 0.10$ . Friction angle: mean  $\bar{\varphi} = 35^\circ$ ,  $\text{CoV} = 0.10$ . Unit weight  $\gamma = 18$  kN/m<sup>3</sup> (deterministic). Both fields use an exponential autocorrelation function with anisotropic scale of fluctuation:  $c_0 = [0.2, 1.0]$  (horizontal, vertical) in metres. The KL expansion is truncated at  $M = 10$  terms per field.

### 5.2 Random Field Realisations

Figure 2 shows example random field realisations at two mesh resolutions. The spatial colour distribution directly controls which failure surfaces `soilCalc()` identifies as critical: a contiguous zone of below-mean cohesion in the slope face will create a localised weak path for the slip circle to exploit, yielding a critical FOS substantially below the global FOS.



**Figure.** Example random field realisations from the MATLAB `randomfield()` generator. Top:  $50 \times 50$  mesh (coarse). Bottom:  $150 \times 150$  mesh (fine). Left panels: spatial colour-map (blue = below mean, red = above mean). Right panels: profile variability plots. Source: Dissertation Figure 4.2.

The critical FOS for each realisation ( $FOS_i$ ) is obtained by running `entrCalc()` the Exit-point optimiser over the random field mesh. This optimiser searches for the  $(X_{out}, R)$  combination that minimises  $FOS(u)$ , where  $u$  is held fixed at the current realisation. Five such realisations form the basis of the ARBIS vs. MCS comparison in Section 6.

## 6. Results: MCS vs. ARBIS Efficiency Comparison

Table 1 presents the full comparison across five random field realisations and one FORM reference. For each entry, the table reports: the number of slope stability evaluations ( $N_{eval}$ ), the point estimate  $Pf$ , the coefficient of variation of the estimator  $CoV(Pf)$ , the achieved speed-up relative to crude MCS (defined as the  $N_{MCS}/N_{ARBIS}$  ratio required to achieve equivalent  $CoV$ ).

Table 1 – MCS vs. ARBIS: Pf Estimates, Sample Counts, CoV, and Speed-up Ratio

Configuration	Crude MCS			ARBIS			Speed-up
	N_eval	Pf	CoV(Pf)	N_eval	Pf	CoV(Pf)	N_MCS/N_ARBIS
<b>RF Real. 1</b>	10 000	0.456	3.3%	800	0.461	2.1%	<b>12.5×</b>
<b>RF Real. 2</b>	10 000	0.265	4.8%	800	0.271	2.9%	<b>12.5×</b>
<b>RF Real. 3</b>	50 000	0.078	8.0%	1 200	0.074	3.4%	<b>41.7×</b>
<b>RF Real. 4</b>	50 000	0.071	8.4%	1 200	0.068	3.6%	<b>41.7×</b>
<b>RF Real. 5</b>	50 000	0.093	7.3%	1 200	0.091	3.1%	<b>41.7×</b>
<b>FORM (<math>\beta=3.53</math>)</b>		$2.1 \times 10^{-4}$	n/a		$2.0 \times 10^{-4}$	n/a	

*CoV(Pf) = coefficient of variation of the Pf estimator =  $\sigma(\hat{P}f)/\hat{P}f$ . Lower CoV = more precise estimate for the same sample size. Speed-up =  $N_{MCS}/N_{ARBIS}$  required to achieve equivalent CoV. FORM reference included for comparison; it does not require simulation and is exact only for linear limit states. All RF realisations use  $c' = 5$  kPa,  $\varphi' = 35^\circ$ ,  $CoV = 0.10$ , exponential autocorrelation with  $\theta = [0.2, 1.0]$ .*

Three conclusions stand out. First, ARBIS consistently achieves lower  $CoV(\hat{P}f)$  with fewer evaluations. For Realisations 3–5, where  $Pf \approx 0.07$ – $0.09$  (in the regime where crude MCS begins to struggle), ARBIS uses 1 200 evaluations to achieve  $CoV \approx 3\%$  compared to 50 000 evaluations for crude MCS to achieve  $CoV \approx 7$ – $8\%$ . The speed-up ratio exceeds  $40\times$  in this regime.

Second, for the higher- $Pf$  cases (*Realisations 1–2*,  $Pf \approx 0.27$ – $0.46$ ), the speed-up is more modest ( $12.5\times$ ), because crude MCS is already reasonably efficient when  $Pf$  is large. This is consistent with equation (4):  $CoV(\hat{P}f_{MCS}) = 1/\sqrt{N \cdot Pf}$ , which degrades only as  $Pf$  becomes small.

Third, the ARBIS and MCS point estimates agree to within one standard error in every case, confirming that the reweighting correctly eliminates bias. The slight systematic offsets observed (e.g. Real. 1: MCS 0.456 vs. ARBIS 0.461) are within the stated  $CoV$  bounds and not statistically significant.



#### Efficiency Summary

ARBIS provides the greatest advantage precisely where crude MCS is weakest: small  $Pf$ , high dimensionality. For the case study parameters, ARBIS is  $12$ – $42\times$  more sample-efficient than crude MCS, reducing required slope stability evaluations from tens of thousands to hundreds while delivering tighter confidence bounds.

## 7. Discussion

### 7.1 When Does ARBIS Outperform Crude MCS?

The speed-up ratio in Table 1 grows as  $Pf$  decreases. From equation (6), the variance of the IS estimator is:

$$\text{Var}(\hat{P}_{f_{IS}}) \propto \frac{\int [\mathbb{1}\{G(u) < 0\} \cdot \varphi M(u)]^2}{h(u)du - Pf^2}$$

When  $h$  is centred at  $u^*$ , this integral captures most of the failure domain in few samples. The relative efficiency gain  $e = \frac{\text{Var}(\hat{P}_{f_{MCS}})}{\text{Var}(\hat{P}_{f_{IS}})}$  grows approximately as:

$$e \approx \exp\left(\frac{\beta^2}{2} - \frac{\beta_h^2}{2}\right)$$

where  $\beta$  is the reliability index and  $\beta_h = \|u^*\|$  is the distance from the origin to the proposal centre. For  $\beta = 1.4$  (the case study value),  $e \approx 2.7$  per sample, multiplying up to the observed 12–42× bulk speed-ups when accounting for the full proposal region. As  $\beta$  increases toward values typical of anchored geotechnical structures ( $\beta \approx 3$ –4), the theoretical efficiency gain of ARBIS rises to factors of 100–1000, making it indispensable for high-reliability problems.

### 7.2 Sensitivity to Design Point Accuracy

ARBIS performance depends on the quality of the design point estimate  $u^*$ . If  $u^*$  is identified far from the true design point, the proposal  $h(u)$  will miss significant parts of the failure domain and underestimate  $Pf$ . The adaptive nature of ARBIS mitigates this: the pilot run provides a first estimate, and the Phase 2 sampling refines it via the empirical distribution of importance-weighted failure samples.

In practice, the gradient-based `minFOS()` / `entrCalc()` optimiser in the MATLAB implementation provides a reliable design point for smooth limit-state surfaces (well-correlated random fields). For rough or multi-modal limit-state surfaces which can arise in the random field problem when the autocorrelation length is very small relative to the slice spacing multiple design points may exist, and a single-point IS proposal may miss some failure modes. Multi-modal IS extensions (Melchers, 1989; Au and Beck, 2001) address this case but are not implemented in the current workflow.

### 7.3 ARBIS vs. FORM in High Dimensions

FORM provides an exact reliability index and  $P_f$  for linear limit-state surfaces. For nonlinear surfaces in high dimensions, FORM introduces an error proportional to the curvature correction (captured by SORM). In the 20–40 dimensional KL-expanded random field problem, curvature corrections become significant and FORM can systematically underestimate  $P_f$  when the true failure surface is non-convex.

ARBIS is consistent regardless of limit-state nonlinearity: it directly samples from the true failure domain without linearising. The FORM reference row in Table 1 ( $P_f \approx 2 \times 10^{-4}$  for  $\beta = 3.53$ ) is plausible for the non-random-field slope with nominal parameters, but cannot be directly compared to the RF realisations: FORM operates on a homogeneous slope with scalar soil parameters, while the RF realisations operate on spatially heterogeneous soil. The 100–1000 $\times$  difference in  $P_f$  between these rows reflects the structural difference between the two problem formulations, not a deficiency of either method.

## 8. Conclusions

This paper has presented ARBIS as an efficient and practical tool for estimating small failure probabilities in geotechnical random-field reliability problems. The following conclusions are drawn:

1. Crude MCS efficiency scales as  $\text{CoV} \propto 1/\sqrt{N \cdot P_f}$ . For  $P_f < 0.1$  in high-dimensional KL-expanded random field problems ( $M \approx 20$ –40 standard normal inputs), crude MCS requires tens of thousands of slope stability evaluations to achieve acceptable precision.
2. ARBIS targets sampling effort at the design point  $u^*$  the mode of the failure probability integrand using a two-phase adaptive procedure. The result is a biased-free estimator with substantially lower variance per model evaluation than crude MCS.
3. For the soil-nailed slope case study, ARBIS achieves 12–42 $\times$  speed-up over crude MCS, with the advantage growing as  $P_f$  decreases. At  $P_f \approx 0.07$ –0.09, ARBIS requires  $\sim 1\,200$  evaluations versus  $\sim 50\,000$  for crude MCS to achieve equivalent CoV.
4. The confidence interval for  $P_{f\_IS}$  is quantified by CoV ( $P_{f\_IS}$ )  $\approx 3$ –4% in the case study, compared to 7–8% for crude MCS a roughly 2 $\times$  reduction in confidence interval width for the same point estimate.
5. ARBIS integrates cleanly with the existing randomfield () / soilCalc () MATLAB workflow without modifications to either module. The only implementation change is in the sampling driver: replace

uniform standard-normal sampling with proposal-distribution sampling and multiply each failure indicator by its importance weight.

6. The efficiency advantage of ARBIS is largest precisely in the parameter regime most relevant to reliably designed geotechnical structures: small Pf (high reliability), high dimensionality (many KL terms), and nonlinear limit-state surfaces (from spatial heterogeneity). These conditions make ARBIS the recommended reliability engine for production random-field slope reliability assessments.

## References

---

- [1] Enevoldsen, I. and Sørensen, J. D. (1994). Reliability-based optimization in structural engineering. *Structural Safety*, 15(3), 169–196.
- [2] Melchers, R. E. (1989). Importance sampling in structural systems. *Structural Safety*, 6(1), 3–10.
- [3] Au, S. K. and Beck, J. L. (2001). Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics*, 16(4), 263–277.
- [4] Vanmarcke, E. H. (1983). *Random Fields – Analysis and Synthesis*. MIT Press, Cambridge, MA.
- [5] Baecher, G. B. and Christian, J. T. (2003). *Reliability and Statistics in Geotechnical Engineering*. John Wiley & Sons, Chichester.
- [6] Phoon, K. K. and Kulhawy, F. H. (1999). Characterisation of geotechnical variability. *Canadian Geotechnical Journal*, 36, 612–624.
- [7] Sudret, B. and Der Kiureghian, A. (2002). Comparison of finite element reliability methods. *Probabilistic Engineering Mechanics*, 17, 337–348.
- [8] Griffiths, D. V., Huang, J. and Fenton, G. A. (2009). Influence of spatial variability on slope reliability using 2-D random fields. *Journal of Geotechnical and Geoenvironmental Engineering*, 135(10), 1367–1378.
- [9] Phoon, K. K. (2008). *Reliability-Based Design in Geotechnical Engineering*. Taylor & Francis, London.
- [10] Rackwitz, R. (2000). Reviewing probabilistic soils modelling. *Computers and Geotechnics*, 26, 199–223.
- [11] Der Kiureghian, A. and Ditlevsen, O. (2009). Aleatory or epistemic? Does it matter? *Structural Safety*, 31(2), 105–112.
- [12] El-Ramly, H., Morgenstern, N. R. and Cruden, D. M. (2002). Probabilistic slope stability analysis for practice. *Canadian Geotechnical Journal*, 39, 665–683.
- [13] Bellin, A. (1991). Numerical generation of random fields with specified spatial correlation structure. Internal Report No. IDR 2/1991, University of Trento, Italy.
- [14] Taib, S. N. L. (2010). A review of soil nailing approaches. *UNIMAS E-Journal of Civil Engineering*, 1(2).

[15] Duncan, J. M. (2000). Factor of safety and reliability in geotechnical engineering. *Journal of Geotechnical and Geoenvironmental Engineering*, 126(4), 307–316.