Multifidelity probability estimation via fusion of estimators

Boris Kramer * Alexandre Noll Marques* Benjamin Peherstorfer[†] Umberto Villa[‡] Karen Willcox[§]

April 22, 2019

This paper develops a multifidelity method that enables estimation of failure probabilities for expensive-to-evaluate models via information fusion and importance sampling. The presented general fusion method combines multiple probability estimators with the goal of variance reduction. We use low-fidelity models to derive biasing densities for importance sampling and then fuse the importance sampling estimators such that the fused multifidelity estimator is unbiased and has mean-squared error lower than or equal to that of any of the importance sampling estimators alone. By fusing all available estimators, the method circumvents the challenging problem of selecting the best biasing density and using only that density for sampling. A rigorous analysis shows that the fused estimator is optimal in the sense that it has minimal variance amongst all possible combinations of the estimators. The asymptotic behavior of the proposed method is demonstrated on a convectiondiffusion-reaction partial differential equation model for which 10^5 samples can be afforded. To illustrate the proposed method at scale, we consider a model of a free plane jet and quantify how uncertainties at the flow inlet propagate to a quantity of interest related to turbulent mixing. Compared to an importance sampling estimator that uses the high-fidelity model alone, our multifidelity estimator reduces the required CPU time by 65% while achieving a similar coefficient of variation.

Keywords: Multifidelity modeling, uncertainty quantification, information fusion, importance sampling, reduced-order modeling, failure probability estimation, PDEs, turbulent jet

^{*}Department of Aeronautics & Astronautics, Massachusetts Institute of Technology, Cambridge, MA bokramer@mit.edu, noll@mit.edu.

[†]Courant Institute of Mathematical Sciences, New York University, NY (pehersto@cims.nyu.edu).

[‡]Department of Electrical and Systems Engineering, Washington University in St. Louis, MO (uvilla@ wustl.edu).

[§]Oden Institute for Computational Engineering & Sciences, The University of Texas at Austin, TX (kwillcox@ices.utexas.edu)

1 Introduction

This paper considers multifidelity estimation of failure probabilities for large-scale applications with expensive-to-evaluate models. Failure probabilities are required in, e.g., reliable engineering design and risk analysis. Yet failure probability estimation with expensive-to-evaluate nonlinear models is computationally challenging due to the large number of Monte Carlo samples needed for low-variance estimates.

Efficient failure probability estimation methods aim to reduce the number of samples at which the expensive model is evaluated, e.g., by exploiting variance-reducing sampling strategies, multifidelity/multilevel estimation methods, or sequential sampling approaches. Variance reduction can be obtained through importance sampling [33], which allows for order-of-magnitude reductions in the number of samples needed to reliably estimate a small probability. However, importance sampling relies on having a good biasing distribution which in turn requires insight into the system. Surrogate models can provide such insight at much lower computational cost. Multifidelity approaches (see [38] for a review) that use surrogates for failure probability estimation via sampling have seen great interest recently [26, 8, 27, 13, 42, 35, 14], but require that the user selects a good importance sampling density. Multifidelity methods that avoid the selection of a single biasing density and instead use a suite of surrogate models to generate importance sampling densities were proposed in [36, 37, 32]. Nevertheless, this framework requires all knowledge about the small probability event to be available in the form of biasing densities, and is therefore only applicable to importance sampling estimators. Multilevel Monte Carlo [15, 2] methods use a hierarchy of approximations to the high-fidelity model in the sampling scheme. However, those model hierarchies have to satisfy certain error decay criteria, an assumption we do not make here. Subset simulation [3, 34] and line search [40, 11] can be used directly on the high-fidelity models, and therefore are of a black-box nature.

In this work, in addition to the computationally expensive model, we also have information about the system in form of surrogate models, analytical models, expert elicitation, and reduced models. In other settings where such a variety of information is available, information fusion has been used to combine multi-source probabilistic information into a single estimator, see [9, 31, 28]. Moreover, combining information from multiple models and sources via a weighted multifidelity method can lead to efficient data assimilation strategies [30].

Here, we propose a new approach to enable small probability estimation for large-scale, computationally expensive models that draws from prior work in information fusion, importance sampling, and multifidelity modeling. We use information fusion in combination with multifidelity importance-sampling-based failure probability estimators, where in addition to the variance reduction from importance sampling, we obtain further variance reduction through information fusion. The proposed multifidelity framework uses the available surrogates to compute multiple unbiased failure probability estimators. We then combine them *optimally* into a new unbiased estimator that has minimal variance amongst all possible linear combinations of those estimators. The method therefore

avoids the selection of the lowest variance biasing density to be used for sampling. Selecting the density that leads to the lowest variance in the failure probability estimator would require additional information, and not even error estimates on the surrogate model would suffice. Thus, we circumvent this step and optimally use all information available to us in form of probability estimators.

This paper is structured as follows: In Section 2 we illustrate the challenges in small failure probability computation and cover the necessary background material for multifidelity importance sampling used herein. Section 3 details our proposed approach of information fusion, importance sampling and multifidelity modeling. We then present in Section 4 a moderately expensive convection-diffusion-reaction test case, where we illustrate the asymptotic behavior of our approach. Section 5 discusses a turbulent jet model and demonstrates the computational efficiency of our proposed methods for this computationally expensive model. We close with conclusions in Section 6.

2 Small probability events and importance sampling estimators

We are interested in computing events with small probabilities, e.g., failure events, where the system fails to meet critical constraints. Section 2.1 describes small probability events, Section 2.2 introduces importance sampling and Section 2.3 briefly summarizes multifidelity importance sampling.

2.1 Small probability events

Let Ω be a sample space which, together with a sigma algebra and probability measure, defines a probability space. Define a *d*-dimensional random variable $Z : \Omega \mapsto \mathcal{D} \subseteq \mathbb{R}^d$ with probability density p, and let \mathbf{z} be a realization of Z. Let $f : \mathcal{D} \subseteq \mathbb{R}^d \mapsto \mathbb{R}^{d'}$ be an expensive-to-evaluate model of high fidelity with corresponding d'-dimensional quantity of interest $f(\mathbf{z}) \in \mathbb{R}^{d'}$. Let $g : \mathbb{R}^{d'} \mapsto \mathbb{R}$ denote a limit state function that defines failure of the system. If $g(f(\mathbf{z})) < 0$, then $\mathbf{z} \in \mathcal{D}$ is a configuration where the system fails. This defines a failure set

$$\mathcal{G} := \{ \mathbf{z} \in \mathcal{D} \mid g(f(\mathbf{z})) < 0 \}.$$

Define the indicator function $I_{\mathcal{G}}: \mathcal{D} \mapsto \{0, 1\}$ via

$$I_{\mathcal{G}}(\mathbf{z}) = \begin{cases} 1, & \mathbf{z} \in \mathcal{G}, \\ 0, & \text{otherwise} \end{cases}$$

The standard Monte Carlo estimator of the failure probability

$$P = \mathbb{E}_p[I_{\mathcal{G}}[Z]] = \int_{\mathcal{D}} I_{\mathcal{G}}(\mathbf{z}) p(\mathbf{z}) d\mathbf{z}$$

uses n realizations $\mathbf{z}_1, \ldots, \mathbf{z}_n$ of the random variable Z and estimates

$$P_n = \frac{1}{n} \sum_{i=1}^n I_{\mathcal{G}}(\mathbf{z}_i).$$
(1)

In the special case of small probabilities, standard Monte Carlo may be unfeasible due to the large number of samples needed to obtain good estimators. Since failure probabilities are generally small, most realizations \mathbf{z}_i will be outside the failure domain \mathcal{G} , and conversely, only a small fraction of the *n* samples lies in the failure region. The coefficient of variation (also called relative root-mean-squared error) of the estimator P_n is given by

$$e^{\text{CV}}(P_n) = \sqrt{\frac{\mathbb{V}[P_n]}{(\mathbb{E}[P_n])^2}} = \sqrt{\frac{P(1-P)}{nP^2}} = \sqrt{\frac{1-P}{nP}}.$$
 (2)

Thus, to obtain estimators with a small coefficient of variation, a large number of samples is necessary. For instance, if the small probability is $P = 10^{-4}$ and if we want $e^{CV} = 10^{-1}$ we would need $n = \mathcal{O}(10^6)$ samples via standard Monte Carlo approaches. This challenge is amplified by the presence of an expensive-to-evaluate model, such as the model of a free plane jet in Section 5.

2.2 Importance sampling

Importance sampling achieves variance reduction by using realizations of a random variable $Z' : \Omega \mapsto \mathcal{D}$ with probability density q. This random variable Z' is chosen such that its probability density function q has higher mass (compared to the nominal density p) in the region of the event of interest. For a general introduction to importance sampling, see [33, Sec.9]. Define the support $\operatorname{supp}(p) = \{\mathbf{z} \in \mathcal{D} \mid p(\mathbf{z}) > 0\}$, and let $\operatorname{supp}(p) \subseteq \operatorname{supp}(q)$. Then

$$P = \int_{\mathcal{D}} I_{\mathcal{G}}(\mathbf{z}) p(\mathbf{z}) d\mathbf{z} = \int_{\mathcal{D}} I_{\mathcal{G}}(\mathbf{z}) \frac{p(\mathbf{z})}{q(\mathbf{z})} q(\mathbf{z}) d\mathbf{z}$$
(3)

is well defined, where $p(\mathbf{z})/q(\mathbf{z})$ is the *likelihood ratio*—in the context of importance sampling also called *importance weight*. The importance-sampling estimate of the failure probability P then draws n realizations $\mathbf{z}'_1, \ldots, \mathbf{z}'_n$ of the random variable Z' with density q and evaluates

$$P_n^{\rm IS} = \frac{1}{n} \sum_{i=1}^n I_{\mathcal{G}}(\mathbf{z}_i') \frac{p(\mathbf{z}_i')}{q(\mathbf{z}_i')}.$$
(4)

The variance of the importance sampling estimator is

$$\mathbb{V}[P_n^{\mathrm{IS}}] = \frac{\sigma_q^2}{n},\tag{5}$$

where

$$\sigma_q^2 = \int_{\mathcal{D}} \left(\frac{I_{\mathcal{G}}(\mathbf{z}')p(\mathbf{z}')}{q(\mathbf{z}')} - P \right)^2 q(\mathbf{z}') d\mathbf{z}'.$$
(6)

If $\operatorname{supp}(p) \subseteq \operatorname{supp}(q)$, and by using (3), one can show that the importance sampling estimator P_n^{IS} is an unbiased estimator of the failure probability, i.e.,

$$\mathbb{E}_q[P_n^{\mathrm{IS}}] = \mathbb{E}_p[I_{\mathcal{G}}(Z)] = P.$$

The importance sampling estimator P_n^{IS} has mean P and variance σ_q^2/n , and by the central limit theorem converges in distribution to the normal random variable $\mathcal{N}(P, \sigma_q^2/n)$. Constructing a good biasing density that leads to small σ_q^2 is challenging [33]. We next introduce low-fidelity surrogate models, which are then used to construct biasing densities.

2.3 Multifidelity Importance Sampling (MFIS)

Recall that by $f : \mathcal{D} \mapsto \mathbb{R}^{d'}$ we denote an expensive-to-evaluate model of high fidelity with corresponding quantity of interest $f(\mathbf{z}) \in \mathbb{R}^{d'}$. Let k surrogates

$$f^{(i)}: \mathcal{D} \mapsto \mathbb{R}^{d'}, \ i = 1, \dots, k$$

of lower fidelities be available, which are cheaper to evaluate than the high-fidelity model $f(\cdot)$. We do not assume any information about the accuracy of the $f^{(i)}(\cdot)$ with respect to the high-fidelity model $f(\cdot)$. Sections 4.2 and 5.4 detail the specific surrogate models used for the respective applications.

We use the MFIS method (see [35] for details) to obtain k estimators of the failure probability. First, MFIS evaluates the surrogate models $f^{(i)}$ at m_i samples to obtain a surrogate-model specific failure set $\mathcal{G}^{(i)}$. Second, MFIS computes a biasing density q_i by fitting a distribution in form of a Gaussian mixture model to the parameters in the failure set. If no failed samples are found by the surrogate model, i.e., if $\mathcal{G}^{(i)} = \emptyset$, then we set the biasing density to be the nominal density. This leads to k biasing densities q_1, \ldots, q_k from which we get importance sampling estimators

$$P_{n_i}^{\rm IS} = \frac{1}{n_i} \sum_{j=1}^{n_i} I_{\mathcal{G}}(\mathbf{z}_{i,j}) \frac{p(\mathbf{z}_{i,j})}{q_i(\mathbf{z}_{i,j})}, \qquad \mathbf{z}_{i,j} \sim q_i, \ j = 1, \dots, n_i, \tag{7}$$

for i = 1, ..., k. The variance of the importance sampling estimator is given by (5) with $n = n_i$ and $\sigma_q = \sigma_{q_i}$, with $\sigma_{q_i}^2$ being the asymptotic variance from (6) with $q = q_i$.

3 Fusion of multifidelity estimators

In many practical situations, a range of probability estimators are available, for instance in form of MFIS estimators derived from different biasing densities, in form of analytical models, or estimators derived from expert elicitation [31]. If one a priori knew which was the lowest variance estimator then a good strategy would be to sample only from that estimator. However, knowing *a priori* which estimator has the lowest variance is a formidable task, and one has to draw samples to assess which estimator has the lowest variance. In this section, we present our new approach that combines all available estimators in an optimal fashion by solving the following problem. **Problem 1.** Given k unbiased estimators, P_1, \ldots, P_k with expected value P, i.e. $\mathbb{E}[P_i] = P$, $i = 1, \ldots, k$, find an estimator with expected value P of the form

$$P_{\alpha} = \sum_{i=1}^{k} \alpha_i P_i,\tag{8}$$

such that it attains minimal variance amongst all estimators of the form (8). That is, find the optimal weights $\alpha_i \in \mathbb{R}$, i = 1, ..., k such that

$$\min_{\alpha} \mathbb{V}[P_{\alpha}] \quad \text{s.t.} \quad \mathbb{E}[P_{\alpha}] = P.$$
(9)

The fused estimator approach allows to still use information coming from the other (highvariance) estimators, whose samples would have otherwise gone to waste. Moreover, with the proposed method we can estimate small-probabilities for expensive-to-evaluate models by exploiting a variety of surrogates. We derive expressions for the mean and variance of the fused estimator in Section 3.1. In Section 3.2, we derive the optimal weights for the fused estimator. Section 3.3 then discusses the special case of uncorrelated estimators. Our proposed algorithm is discussed in Section 3.4, followed by a brief Section 3.5 that discusses measures of convergence of the estimators.

3.1 Mean and variance of fused estimator

We start with the observation that if the weights α_i of the fused estimator P_{α} sum to one, then the fused estimator is unbiased:

$$\sum_{i=1}^{k} \alpha_i = 1 \quad \Leftrightarrow \quad \mathbb{E}[P_{\alpha}] = \sum_{i=1}^{k} \alpha_i \mathbb{E}[P_i] = P \sum_{i=1}^{k} \alpha_i = P.$$

Let the estimators P_i have corresponding variances $0 < \sigma_i^2 < \infty$, $i = 1, \ldots k$. To compute the variance of the fused estimator P_{α} we have to consider covariances between the individual estimators. Define the Pearson product-moment correlation coefficient as

$$\rho_{i,j} = \frac{\mathbb{C}\mathrm{ov}(P_i, P_j)}{\sigma_i \sigma_j},\tag{10}$$

where $\mathbb{C}ov(P_i, P_j) = \mathbb{E}[(P_i - \mathbb{E}[P_i])(P_j - \mathbb{E}[P_j])] = \mathbb{E}[P_iP_j] - P^2$. We also define the symmetric, positive semi-definite *covariance matrix* $\Sigma_{ij} = \mathbb{C}ov(P_i, P_j)$ as:

$$\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_1^2 & \sigma_1 \sigma_2 \rho_{1,2} & \dots & \cdots & \sigma_1 \sigma_k \rho_{1,k} \\ \sigma_2 \sigma_1 \rho_{2,1} & \sigma_2^2 & \sigma_2 \sigma_3 \rho_{2,3} & \dots & \sigma_2 \sigma_k \rho_{2,k} \\ \vdots & \ddots & & & \\ \vdots & & \sigma_{k-1}^2 & \sigma_{k-1} \sigma_k \rho_{k-1,k} \\ \sigma_k \sigma_1 \rho_{k,1} & \sigma_k \sigma_2 \rho_{k,2} & \dots & \sigma_k \sigma_{k-1} \rho_{k,k-1} & \sigma_k^2 \end{bmatrix} .$$
(11)

It is worth noticing that if the estimators P_1, \ldots, P_k are independent, then Σ is diagonal. The variance of the fused estimator from (8) is

$$\mathbb{V}[P_{\alpha}] = \mathbb{V}\left[\sum_{i=1}^{k} \alpha_{i} P_{i}\right] = \sum_{i=1}^{k} \alpha_{i}^{2} \mathbb{V}[P_{i}] + 2\mathbb{C}\operatorname{ov}\left(\sum_{i=1}^{k} \alpha_{i} P_{i}, \sum_{j=1}^{k} \alpha_{j} P_{j}\right)$$
$$= \sum_{i=1}^{k} \alpha_{i}^{2} \sigma_{i}^{2} + 2\sum_{i=1}^{k} \sum_{j>i}^{k} \alpha_{i} \alpha_{j} \sigma_{i} \sigma_{j} \rho_{i,j},$$

which can be written in vector form as

$$\mathbb{V}[P_{\alpha}] = \alpha^T \Sigma \alpha. \tag{12}$$

In the following section, we provide an explicit formula to find the optimal weights α for the general case of (possibly)-correlated estimators P_1, \ldots, P_k ; while in Section 3.3 we discuss the case of independent estimators, such as those constructed with the MFIS method.

3.2 Optimizing the weights for minimum-variance estimate

Problem (9) seeks the optimal α such that the variance in (12) is minimized and P_{α} remains unbiased. In this section, we show that such weights exist, are unique, and present a closed-form solution, provided that the covariance matrix Σ is invertible. This is summarized in the following result.

Proposition 1. Let $\mathbf{P} = [P_1, \ldots, P_k]^T$ be the vector of probability estimators and assume that Σ is not singular. Define $\mathbf{1}_k = [1, \ldots, 1]^T$ as a column-vector of length k. The optimization problem (9) has the unique solution

$$oldsymbol{lpha} = rac{oldsymbol{\Sigma}^{-1} oldsymbol{1}_k}{oldsymbol{1}_k^T oldsymbol{\Sigma}^{-1} oldsymbol{1}_k}.$$

That is, the minimal variance unbiased estimator P_{α} is such that

$$P_{\boldsymbol{\alpha}} = \frac{\mathbf{1}_k^T \boldsymbol{\Sigma}^{-1} \boldsymbol{P}}{\mathbf{1}_k^T \boldsymbol{\Sigma}^{-1} \mathbf{1}_k}, \qquad \mathbb{V}[P_{\boldsymbol{\alpha}}] = \frac{1}{\mathbf{1}_k^T \boldsymbol{\Sigma}^{-1} \mathbf{1}_k}.$$

Proof. We have seen above that $\sum_{i=1}^{k} \alpha_i = 1$ if and only if $\mathbb{E}[P_{\alpha}] = P$. Define the cost function $J(\alpha) := \mathbb{V}[P_{\alpha}] = \alpha^T \Sigma \alpha$ by using equation (12). Therefore, the optimization problem (9) can be written as the quadratic program

$$\min_{\alpha} J(\alpha) = \alpha^T \Sigma \alpha, \quad \text{s.t.} \quad \alpha^T \mathbf{1} = 1.$$
(13)

Letting $\mathcal{L}(\boldsymbol{\alpha}, \lambda) := \boldsymbol{\alpha}^T \boldsymbol{\Sigma} \boldsymbol{\alpha} + \lambda(\boldsymbol{\alpha}^T \mathbf{1} - 1)$ denote the Lagrangian cost function associated to (13), the optimality conditions are $\nabla_{\boldsymbol{\alpha}} \mathcal{L}(\boldsymbol{\alpha}, \lambda) = \mathbf{0}$ and $\frac{d\mathcal{L}}{d\lambda}(\boldsymbol{\alpha}, \lambda) = 0$. This optimality system is written as

$$\begin{bmatrix} \boldsymbol{\Sigma} & \mathbf{1}_k \\ \mathbf{1}_k^T & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha} \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{0}_k \\ 1 \end{bmatrix}.$$
 (14)

For invertible Σ , the unique weights to this quadratic program are then obtained by

$$\boldsymbol{\alpha} = \frac{\boldsymbol{\Sigma}^{-1} \mathbf{1}}{\mathbf{1}_k^T \boldsymbol{\Sigma}^{-1} \mathbf{1}_k},\tag{15}$$

and the expression for the variance follows by inserting these weights into (13). The estimator is obtained by inserting the weights into (8). \Box

The weights can be expressed explicitly in terms of the components of the covariance matrix as

$$\alpha_i = \frac{1}{\sigma_i^2} \left[\frac{1}{\sum_{l=1}^k \frac{1}{\sigma_l^2}} \left(1 + \sum_{l=1}^k \frac{1}{\sigma_l^2} \sum_{j>l}^k \alpha_j \sigma_l \sigma_j \rho_{l,j} \right) - \sum_{j>i}^k \alpha_j \sigma_i \sigma_j \rho_{i,j} \right].$$
(16)

Note, that the weights are inversely proportional to the variance of the individual estimators and the weight α_i depends on the covariance between the estimators P_i and P_j . Also, note that if P_i are correlated some weights may be negative, while for a diagonal Σ all weights α_i are *strictly* positive. In the next section, we have a closer look at the uncorrelated case.

3.3 The special case of uncorrelated estimators

In the situation where all estimators are uncorrelated, we recover the classical result of the *inverse variance-weighted mean* [29]. As a corollary from Proposition 1 we get the following result.

Corollary 1. Consider the setting from Proposition 1, and let $\Sigma = \text{diag}(\sigma_1^2, \ldots, \sigma_k^2)$ be diagonal. Then the unique solution to the optimization problem (9) is given by

$$\alpha_{i} = \frac{1}{\sigma_{i}^{2} \sum_{i=1}^{k} \frac{1}{\sigma_{i}^{2}}}, \qquad \mathbb{V}[P_{\alpha}] = \frac{1}{\sum_{i=1}^{k} \frac{1}{\sigma_{i}^{2}}}.$$
(17)

A few observations about this special case are in order:

- 1. The optimal coefficients α_i of the combined estimator P_{α} are inversely proportional to the asymptotic variance σ_i of the corresponding estimator P_i . To reduce the variance via a weighted combination of estimators, smaller weights are assigned to estimators with larger variance.
- 2. If one variance is small compared to all other ones, say $\sigma_1^2 \ll \sigma_i^2$, i = 2, ..., k, then $\sum_{i=1}^k \frac{1}{\sigma_i^2} \approx \frac{1}{\sigma_1^2}$ so that $\mathbb{V}[P_{\alpha}] \approx \sigma_1^2$. The estimators with large variance cannot reduce the variance of the fused estimator much more.
- 3. If all estimators have equal variance, $\sigma_1^2 = \ldots = \sigma_k^2$, then $\sum_{i=1}^k \frac{1}{\sigma_i^2} = \frac{k}{\sigma_1}$ so that $\mathbb{V}[P_{\alpha}] = \frac{\sigma_1^2}{k}$. Hence, combining the estimators reduces the variance by a factor of k.

4. Since $0 < \alpha_i < 1, \forall i$, it follows from both equations in (17) that

$$\mathbb{V}[P_{\alpha}] = \frac{1}{\sum_{i=1}^{k} \frac{1}{\sigma_i^2}} = \sigma_i^2 \alpha_i < \sigma_i^2, \forall i \quad \Rightarrow \quad \mathbb{V}[P_{\alpha}] < \min_{i=1,\dots,k} \sigma_i^2.$$
(18)

Consequently, we are guaranteed to reduce the variance in P_{α} by combining all estimators in the optimal way described above.

3.4 Fused multifidelity importance sampling: Algorithm and Analysis

We now use the general fusion framework to obtain a failure probability estimator. Thus, we solve Problem 1 in the context of importance-sampling-based failure probability estimators so that $P_i = P_{n_i}^{\text{IS}}$. Our proposed method optimally fuses the k MFIS estimators from (7), such that

$$P_{\alpha} = \sum_{i=1}^{k} \alpha_i P_{n_i}^{\rm IS},\tag{19}$$

with the optimal weights chosen as in Proposition 1 and $\sum_{i=1}^{k} n_i = n$. Since estimator $P_{n_i}^{\text{IS}}$ is computed from n_i samples, P_{α} uses $n = \sum_{i=1}^{k} n_i$ samples. We now discuss how P_{α} compares to a single importance sampling estimator with n

We now discuss how P_{α} compares to a single importance sampling estimator with n samples. Consider the estimator $P_{j'}^{\text{IS}}$ that uses n samples drawn from a single biasing density $q_{j'}$ for $j' \in \{1, \ldots, k\}$. This estimator would require selection of the lowest biasing density *a priori*, a formidable task. The next results compares P_{α} and $P_{j'}^{\text{IS}}$, and gives a criterion for which the former has lower variance than the latter.

Proposition 2. Let k estimators $P_{n_i}^{IS}$ with $n_1 = n_2 = \ldots = n_k$ samples be given. Let $j' \in \{1, \ldots, k\}$, and $q_{j'}$ be a biasing density that is used to derive an IS estimator $P_{j'}^{IS}$ with $n = kn_1$ samples. If

$$\sigma_{q'_j}^2 > \frac{k}{\sum_{i=1}^k \frac{1}{\sigma_i^2}}$$

then the variance of the fused estimator P_{α} in (19) with n samples is smaller than the variance of the estimator with biasing density $q_{j'}$ with n samples, i.e.,

$$\mathbb{V}[P_{\alpha}] < \mathbb{V}[P_{j'}]$$

Proof. Set $n_i = n/k$, i = 1, ..., k, so that all estimators use the same number of samples. According to equation (17),

$$\mathbb{V}[P_{\alpha}] = \frac{1}{\sum_{i=1}^{k} \frac{n_i}{\sigma_{q_i}^2}} = \frac{k}{n \sum_{i=1}^{k} \frac{1}{\sigma_{q_i}^2}}$$

as well as $\mathbb{V}[P_{j'}^{\mathrm{IS}}] = \frac{\sigma_{q_{j'}}^2}{n}$, so that

$$\mathbb{V}[P_{\alpha}] = \frac{k}{n\sum_{i=1}^{k} \frac{1}{\sigma_{q_i}^2}} < \frac{\sigma_{q_{j'}}^2}{n} = \mathbb{V}[P_{j'}^{\mathrm{IS}}] \qquad \Leftrightarrow \qquad \frac{k}{\sum_{i=1}^{k} \frac{1}{\sigma_{q_i}^2}} < \sigma_{q_{j'}}^2$$

The importance sampling estimate (7) requires evaluating the high-fidelity model at n_i samples from the biasing density. While not required, we use $n_i = n/k$, $i = 1, \ldots, k$ to distribute the computational load evenly. Extension of Proposition 2 is straightforward to the case with different number of samples n_i for each estimator P_i

The computational procedure is summarized in Algorithm 1. Here, we denote samplingbased estimates as $\hat{P}_{n_i}^{\text{IS}}$, which are realizations of the estimator $P_{n_i}^{\text{IS}}$.

Algorithm 1 Computing failure probability estimate \hat{P}_{α} via fused importance sampling **Input:** Nominal distribution p, biasing distributions $\{q_i\}_{i=1}^k$, # of evaluations $\{n_i\}_{i=1}^k$, limit state function $q(\cdot)$.

Output: Failure probability estimate \hat{P}_{α} and variance estimate $\mathbb{V}[\hat{P}_{\alpha}]$

1: for j = 1 : k do {Loop over all surrogates}

2: Draw $\mathbf{z}_{j,1}, \ldots, \mathbf{z}_{j,n_j}$ independent realizations from Z_j with density q_j and compute

$$\hat{P}_{n_j}^{\rm IS} = \frac{1}{n_j} \sum_{i=1}^{n_j} I_{\mathcal{G}}(\mathbf{z}_{j,i}) \frac{p(\mathbf{z}_{j,i})}{q_j(\mathbf{z}_{j,i})}$$
(20)

Compute the sample variances 3:

$$\hat{\sigma}_{q_j}^2 = \frac{1}{n_j - 1} \sum_{i=1}^{n_j} \left(I_{\mathcal{G}}(\mathbf{z}_{j,i}) \frac{p(\mathbf{z}_{j,i})}{q_j(\mathbf{z}_{j,i})} - \hat{P}_{n_j}^{\mathrm{IS}} \right)^2 \tag{21}$$

4: end for

- 5: Define the vector $\boldsymbol{P} = [\hat{P}_{n_1}^{\text{IS}}, \dots, \hat{P}_{n_k}^{\text{IS}}]^T$ 6: Let $\hat{\boldsymbol{\Sigma}} = \text{diag}(\hat{\sigma}_{q_1}^2/n_1, \dots, \hat{\sigma}_{q_k}^2/n_k)$
- 7: Compute the fused estimate as in (13):

$$\hat{P}_{\alpha} = \frac{\mathbf{1}_{k}^{T} \hat{\boldsymbol{\Sigma}}^{-1} \boldsymbol{P}}{\mathbf{1}_{k}^{T} \hat{\boldsymbol{\Sigma}}^{-1} \mathbf{1}_{k}}, \qquad \mathbb{V}[\hat{P}_{\alpha}] = \frac{1}{\mathbf{1}_{k}^{T} \hat{\boldsymbol{\Sigma}}^{-1} \mathbf{1}_{k}}$$
(22)

3.5 Error measures and practical computation

The failure probability estimate $\hat{P}_{n_i}^{\mathrm{IS}}$ is computed as in (20) and the sample variance $\hat{\sigma}_{q_i}^2$ as in (21). The root-mean-squared-error (RMSE) of the estimate P_{n_i} is

$$e^{\text{RMSE}}(\hat{P}_{n_i}) = \sqrt{\frac{\hat{\sigma}_{q_i}^2}{n_i}},\tag{23}$$

and the relative mean-squared-error, or coefficient of variation is computed as

$$e^{\text{CV}}(\hat{P}_{n_i}) = \sqrt{\frac{\hat{\sigma}_{q_i}^2}{n_i(\hat{P}_{n_i}^{\text{IS}})^2}}.$$
 (24)

4 Test case: Convection-diffusion-reaction

We first consider a PDE model whose solution can be numerically evaluated with moderate computational cost. With this model, we demonstrate the asymptotic behavior of our method because we can afford to sample the high-fidelity model $n = 10^5$ times, which will be too costly for the model in Section 5. The test problem is the convectiondiffusion-reaction PDE introduced in Section 4.1. Its discretizations and reduced-order models are described in Section 4.2. Numerical results are presented in Section 4.3.

4.1 Convection-diffusion-reaction PDE model

We consider a simplified model of a premixed combustion flame at constant and uniform pressure, and follow the notation and set-up in [7, Sec.3]. The model includes a one-step reaction of the species

$$2H_2 + O_2 \rightarrow 2H_2O$$

in the presence of an additional non-reactive species, nitrogen. The physical combustor domain is 18mm in length (x-direction), and 9mm in height (y-direction), as shown in Figure 1.



Figure 1: Set-up of combustor, with details of the boundary conditions in Table 1.

The velocity field U is set to be constant in the positive x direction, and divergence free. The molecular diffusivity κ is modeled as constant, equal and uniform for all species and temperature. The PDE model is given by

$$0 = \kappa \Delta s - U \nabla s + \mathcal{F}(s, \mathbf{z}) \quad \in \Omega \tag{25}$$

where the state is comprised of the components $s = [T, Y_{H_2}, Y_{O_2}, Y_{H_2O}]$, with the Y_i being the mass fractions of the species (fuel, oxidizer, product), and T denoting the temperature. Referring to Figure 1, we have that $\Gamma_D = \Gamma_{D,i} \cup \Gamma_{D,0}$ is the Dirichlet part of the boundary and Γ_N combines the top, bottom and right boundary, where Neumann

Table 1: Boundary conditions for the combustion model from [7]. Boundary Temperature Species

o and any	romboracaro	Species
$\Gamma_{D,i}$	$T = 950 \mathrm{K}$	$Y_{H_2} = 0.0282, \ Y_{O_2} = 0.2259, \ Y_{H_2O} = 0$
$\Gamma_{D,0}$	$T = 300 \mathrm{K}$	$Y_{H_2} = 0, \ Y_{O_2} = 0, \ Y_{H_2O} = 0$
Γ_N	$\nabla T \cdot \boldsymbol{n} = 0$	$\nabla Y_i \cdot \boldsymbol{n} = 0$

Table 2: Parameters for the combustion model from [7].

quantity	physical meaning	assumptions	value
κ	molecular diffusivity	const., equal, uniform $\forall i$	$2.0\frac{\mathrm{cm}^2}{\mathrm{s}}$
U	velocity	const.	$50\frac{\text{cm}}{\text{s}}$
W_{H_2}	molecular weight	const.	$2.016 \frac{\text{g}}{\text{mol}}$
W_{O_2}	molecular weight	const.	$31.9\frac{\text{g}}{\text{mol}}$
W_{H_2O}	molecular weight	const.	$18\frac{\text{g}}{\text{mol}}$
ho	density of mixture	const.	$1.39 \times 10^{-3} \frac{\text{g}}{\text{cm}^3}$
R	univ. gas constant	const.	$8.314472 \frac{J^{-11}}{\text{mol K}}$
Q	heat of reaction	const.	9800K
$ u_{H_2}$	stochiometric coefficient	const.	2
ν_{O_2}	stochiometric coefficient	const.	1
$ u_{H_2O}$	stochiometric coefficient	const.	2

conditions are prescribed. In sum, $\partial \tilde{\Omega} = \Gamma_D \cup \Gamma_N$; the boundary conditions are imposed as given in Table 1. The nonlinear reaction term $\mathcal{F}(s, \mathbf{z}) = [\mathcal{F}_T, \mathcal{F}_{H_2}, \mathcal{F}_{O_2}, \mathcal{F}_{H_2O}](s, \mathbf{z})$ is of Arrhenius type [10], and modeled as

$$\mathcal{F}_{i}(s,\mathbf{z}) = -\nu_{i} \left(\frac{W_{i}}{\rho}\right) \left(\frac{\rho Y_{F}}{W_{F}}\right)^{\nu_{F}} \left(\frac{\rho Y_{O}}{W_{O}}\right)^{\nu_{O}} A \exp\left(-\frac{E}{RT}\right), \quad i = H_{2}, \ O_{2}, \ H_{2}O \quad (26)$$
$$\mathcal{F}_{T}(s,\mathbf{z}) = Q \ \mathcal{F}_{H_{2}O}(s,\mathbf{z}). \tag{27}$$

The parameters of the model are defined in Table 2. The uncertain parameters are the pre-exponential factor A and the activation energy E of the Arrhenius model. The domain for these parameters is denoted as \mathcal{D} . In particular, we have that

$$\mathbf{z} = [A, E] \in \mathcal{D} = [5.5 \times 10^{11}, \ 1.5 \times 10^{13}] \times [1.5 \times 10^3, \ 9.5 \times 10^3].$$

4.2 Discretization and reduced-order models

The model is discretized using a finite difference approximation in two spatial dimensions, with 72 nodes in x direction, and 36 nodes in y direction, leading to 10,804 unknowns in the model. The nonlinear system is solved with Newton's method. Let $\mathbf{T}(\mathbf{z})$ be the vector with components corresponding to the approximations of the temperature $T(x, y; \mathbf{z})$ at the grid points. The high-fidelity model (HFM) is $f : \mathcal{D} \to \mathbb{R}$ and the quantity of interest is the maximum temperature over all grid points:

$$f(\mathbf{z}) = \max \mathbf{T}(\mathbf{z})$$

Reduced-order models provide a powerful framework to obtain surrogates for expensiveto-evaluate models. In the case of nonlinear systems, reduced-order models can be obtained via reduced-basis methods [19], dynamic mode decomposition [24], proper orthogonal decomposition [5], and many others; for a survey, see [4]. Here, we compute reduced-order models $f^{(i)}$ for our multifidelity approach via Proper Orthogonal Decomposition and the Discrete Empirical Interpolation Method (DEIM) for an efficient evaluation of the nonlinear term. The training snapshots are generated from solutions to the high-fidelity model on a parameter grid of 50×50 equally spaced values $\mathbf{z} \in \mathcal{D}$. The three surrogate models are built from 2, 10, 15 POD basis vectors, and accordingly 2, 5, 10 DEIM interpolation points. The corresponding models are denoted as ROM1, ROM2, ROM3, respectively. We denote by $\mathbf{T}_r^{(i)}(\mathbf{z})$ the approximation to the temperature $T(x, y; \mathbf{z})$ via the *i*th ROM. The surrogate models $f^{(i)}$ are the mappings $f^{(i)} : \mathcal{D} \mapsto \mathbb{R}$ with corresponding quantity of interest denoted as

$$f^{(i)}(\mathbf{z}) = \max \mathbf{T}_r^{(i)}(\mathbf{z}), \quad i = 1, \dots, k.$$

We refer the reader to [7] for more details on the discretization and ROM construction for this convection-diffusion-reaction model.

4.3 Results for multifidelity fusion of failure probabilities

We define a failure of the system when the maximum temperature in the combustor exceeds 2430K, so that the limit state function is

$$g(f(\mathbf{z})) = 2430 - f(\mathbf{z}),$$
 (28)

and likewise for the reduced-order models $g(f^{(i)}(\mathbf{z})) = 2430 - f^{(i)}(\mathbf{z})$.

To compute the biasing densities, we draw $\hat{m} = 20,000$ samples from the uniform distribution on \mathcal{D} , compute surrogate-based solutions, and evaluate the limit state function for those solutions. If the limit state function indicates failure of the system for a solution obtained from the *i*th surrogate model, the corresponding parameter is added to $\mathcal{G}^{(i)}$, the failure set computed from the *i*th surrogate model. We compute the biasing densities q_1, q_2, q_3 via MFIS (see Section 2.3) as Gaussian mixture distributions with a single component. Table 3 shows the computational cost in CPU time of computing the biasing distributions from the various ROMs and the HFM. Computing a biasing density using the high-fidelity model with $\hat{m} = 20,000$ samples costs approximately 2.1 CPUhours. Constructing the biasing density via the low-fidelity models ROM2 and ROM3 reduces the computational time by a factor of 66 and 58, respectively. Note, that ROM1 is the reduced-order model that is cheapest to execute per model evaluation, but it is also the least accurate. In our case, ROM1 did not produce any samples in the failure region, even after $\hat{m} = 10^5$ samples. It is not unexpected that ROM1 is so inaccurate, since only two POD modes are not enough to resolve the important character of this problem. ROM1 is included to demonstrate how the fusion approach can be effective even in the presence of highly inaccurate surrogate models.

Table 3: CPU time to generate the biasing densities, and the number of samples in the failure domain.

	ROM1	ROM2	ROM3	HFM
# of samples drawn	10^{5}	2×10^4	2×10^4	2×10^4
# of samples in failure domain $\mathcal{G}^{(i)}$	0	13	17	17
time needed	N.A.	11.2[s]	12.7[s]	2.1[h]

In Figure 2 we show the quantity of interest, i.e., the maximum temperature. The plots are obtained by generating $m = 10^5$ samples from the nominal distribution (left) and the respective biasing distributions (right), and evaluating the HFM at those samples. Figure 2, left, shows that the typical range of the quantity of interest is between approximately 1200K and 2440K. However, only the events where the quantity of interest is above 2430K are relevant for the failure probability computation. By using the biasing distributions in Figure 2, right, a large portion of the outputs leads to a failure of the system. This indicates that the biasing distributions are successful in generating samples at the failure region of the high-fidelity model.



Figure 2: Quantity of interest $f(\mathbf{z})$ in [K] of HFM ordered by magnitude versus # of samples \mathbf{z} , for $m = 10^5$ samples. Left: Samples are from the nominal (uniform) distribution. Right: The parameter samples are drawn from different biasing distributions (biased towards failure above 2430K). This demonstrates that the biasing distributions are good since the outputs are largely above the failure threshold. Here, ROM1 did not have any parameters in the failure domain, and hence defaulted to being the nominal distribution and is therefore not plotted.

Next, we show results for the fused multifidelity estimator P_{α} with n samples and compare it with importance sampling estimators $\hat{P}_{n_i}^{\text{IS}}$ that only use a single biasing density and also n samples. The fused estimator is obtained via Algorithm 1 with $n_i = \lfloor n/3 \rfloor, i = 1, 2, 3$, samples by fusing the three surrogate-model-based importance sampling estimators. For reference purposes, a biasing density is constructed as described above using the HFM with $\hat{m} = 20,000$ samples. Based on this density, we compute an importance sampling estimate of the failure probability with $n = 10^5$ samples, resulting in $\hat{P}_{105}^{\text{IS}} = 8.42 \times 10^{-4}$.

To assess the quality of the fused estimator P_{α} , we consider the error measures introduced in Section 3.5. In Figure 3, left, we show the root mean-squared error of the importance sampling estimators $\hat{P}_{n_i}^{\text{IS}}$ as well as the combined estimator \hat{P}_{α} . Figure 3, right, shows the coefficient of variation defined in (24) for the estimators. The fused estimator is competitive in RMSE and coefficient of variation with the estimator using the high-fidelity biasing density, but comes at a much lower computational cost.

Note, that the fused estimator does not use any of the high-fidelity information. We only plotted the high-fidelity estimator for comparison reasons, but the high-fidelity density is not used in our algorithm. Heuristically, we could expect the fused estimator to perform better than the MFIS estimator with high-fidelity-derived biasing density in the following situation. Let the HFM be so expensive that the HF biasing density is built only from a few failure samples, and assume the low-fidelity models are good surrogates, hence able to cheaply explore the failure region. Then the low-fidelity biasing density could be better than the high-fidelity biasing density.



Figure 3: Left: Root mean-squared error from (23); Right: Coefficient of variation as defined in (24) for the convection-diffusion-reaction simulation.

In Table 4 we show the weights for the fused estimator \hat{P}_{α} . The fused estimator assigns only a small weight α_1 to the estimator $\hat{P}_{n_1}^{\text{IS}}$ which uses biasing density q_1 . This was expected, as the estimator has large variance due to the fact that biasing density q_1 is actually the nominal density, see Table 3 as the ROM1 evaluation did not yield any

samples in the failure domain.

	α and α and α and α and α and α						
	$n = 10^2$	$n = 10^{3}$	$n = 10^4$	$n=2\times 10^4$	$n=4\times 10^4$	$n = 10^5$	
α_1	0	0	0.005	0.001	0.002	0.005	
α_2	0.587	0.471	0.331	0.294	0.415	0.742	
α_3	0.413	0.529	0.664	0.705	0.583	0.253	

				~~		
Table 4: Weights	of the	fused	estimator	P_{α}	with n	i samples

4.4 Comparison to subset simulation methods

To demonstrate the efficiency of our proposed multifidelity method compared to stateof-the-art existing methods in failure probability estimation, we compare our results to subset simulation [3], a widely used method for reliability analysis and failure probability estimation. The method defines intermediate failure events

$$\mathcal{G}_j := \{ \mathbf{z} \in \mathcal{D} \mid g(f(\mathbf{z})) < b_j \}, \ j = 1, \dots, L,$$

for a sequence of threshold levels $b_1 > b_2 > \ldots > b_L = 0$ and L being the final level. This ensures that the intermediate failure events are nested as $\mathcal{G}_1 \supset \mathcal{G}_2 \supset \ldots \supset \mathcal{G}_L = \mathcal{G}$. The failure probability can then be expressed as

$$P = P(I_{\mathcal{G}}) = \mathbb{E}(I_{\mathcal{G}_1}) \prod_{j=2}^{L} P(I_{\mathcal{G}_j} | I_{\mathcal{G}_{j-1}}).$$

Thus, this method requires sampling from the conditional events $\mathcal{G}_j|\mathcal{G}_{j-1}$, and the efficiency of this sampling is pivotal to the success of subset simulation. Markov Chain Monte Carlo (MCMC) methods provide efficient solutions to this problem [34]. Note, that the b_j cannot be determined in advance, but are found adaptively by specifying an intermediate failure probability $p_0 = P(\mathcal{G}_j|\mathcal{G}_{j-1})$. A typical choice is $p_0 = 0.1$ which yields efficient subset simulation results, see [3].

Here, we compare our fused importance sampling approach for failure probability estimation to a direct application of subset simulation to the full model. We follow the recent MCMC implementation for subset simulation of [34]. Table 5 lists the computational results that include the number of levels L that subset simulation needed to arrive at the failure probability estimate, the samples at each level (user defined), the failure probability estimate, and the overall number of samples needed (not known beforehand). All results were averaged over ten independent runs. We also give an approximate coefficient of variation, although we caution that this is *not* the same coefficient of variation defined in (2), since at the intermediate levels, subset simulation produces correlated samples. Thus, we used an approximated coefficient of variation as suggested in [45, Eq. (19)]. For a thorough discussion of the coefficient of variation estimation within subset simulation we refer to [6, Sec.5.3]. We observe from Table 5 that the coefficient of variation monotonically decreases as more samples are added. To compare our proposed multifidelity fusion method with subset simulation, we first note that the estimate from subset simulation \hat{P}_f is biased for finite N (see [3, Sec.6.3]), whereas our fused estimator \hat{P}_{α} is unbiased. Moreover, the numerical results in Table 5 show that the estimated coefficients of variation are about one order of magnitude larger than the coefficients of variation we reported in Figure 3, right. From a computational cost perspective, the estimator with 20,000 samples in subset simulation produces an approximated coefficient of variation of 1.18×10^{-1} whereas our fused estimator \hat{P}_{α} produces a coefficient of variation of 1.34×10^{-2} for the same number of high-fidelity model evaluations. Thus, the fused estimator outperforms subset simulation in this particular example. In sum, our method can successfully take advantage of cheaper low-fidelity methods to get accurate estimators, while the subset simulation method works directly with the full model and therefore does not have access to cheaper surrogate model information.

Table 5: Results for subset simulation to compute failure probabilities for the convectiondiffusion-reaction problem.

l C.o.V.
10^{-1}
10^{-1}
10^{-1}
10^{-1}
10^{-1}
10^{-1}
1

5 Failure probability estimation related to a free plane jet

We apply the proposed fusion of estimators to quantify the influence of uncertain parameters on the amount of turbulent mixing produced by a free plane jet.

This is a challenging problem, since it involves an expensive-to-evaluate model for which the naive computation of low probabilities requires thousands of hours of computation. We reduce this number significantly with our multifidelity importance sampling framework via fusion of estimators.

The remainder of this section is organized as follows. Section 5.1 introduces the free plane jet, followed by details of the model and its governing equations in Section 5.2. The uncertain parameters and quantity of interest are defined in Section 5.3. The low-fidelity surrogate models used in this investigation are discussed in Section 5.4. Finally, the results for multifidelity fusion of small probability estimators are presented in Section 5.5.

5.1 Large-scale application: Free plane jet

Free turbulent jets are prototypical flows believed to represent the dynamics in many engineering applications, such as combustion and propulsion. As such, free jet flows are the subject of several experimental [18, 17, 23] and numerical investigations [44, 39, 41, 21, 22] and constitute an important benchmark for turbulent flows.

Our expensive-to-evaluate model of a free plane jet is based on the two-dimensional incompressible Reynolds-averaged Navier-Stokes (RANS) equations, complemented by the $k - \epsilon$ turbulence model. Although a RANS model does not resolve all relevant turbulent features of the flow, it represents a challenging large-scale application for the computation of small probabilities. We use this model to investigate the influence of five uncertain parameters on the amount of turbulent mixing produced by the jet. We quantify turbulent mixing using a relatively simple metric: the integral jet width. One of the uncertain parameters is the Reynolds number at the inlet of the jet, which is assumed to vary from 5,000 to 15,000. The other four uncertain parameters correspond to coefficients of the $k - \epsilon$ turbulence model and its boundary condition, as detailed in Section 5.3. Figure 4 shows a flow field typical of the cases considered here.



Figure 4: Flow field of a two-dimensional plane jet at Reynolds number 10,000, computed with standard coefficients of $k - \epsilon$ turbulence model.

5.2 Modeling and governing equations

We consider a free plane jet in conditions similar to the ones reported in [21, 22]. Namely, the flow exits a rectangular nozzle into quiescent surroundings with a prescribed top-hat velocity profile and turbulence intensity. The nozzle has width D, and is infinite along the span-wise direction. The main difference between the free plane jet we considered here and the one described in [21, 22] is the Reynolds number at the exit nozzle. Here the Reynolds number varies between 5,000 and 15,000.

Our simulation model computes the flow in a rectangular domain Ω located at a distance 5D downstream from the exit of the jet nozzle, as illustrated in Figure 5. By doing so, modeling the conditions at the exit plane of the jet nozzle is avoided. Instead, direct numerical simulation data are used to define inlet conditions at the surface $\Gamma_{\rm in}$. The dynamics are modeled with the incompressible Reynolds-averaged Navier-Stokes



Figure 5: Illustration of the free plane jet setup. The diameter of the nozzle is denoted by D. The simulation domain Ω is composed of a $30D \times 10D$ box situated at a distance 5D downstream to the nozzle exit.

equations, complemented by the $k - \epsilon$ turbulence model [25]:

$$(\boldsymbol{v}\cdot\nabla)\boldsymbol{v} + \frac{1}{\rho}\nabla p - \nabla\cdot((\nu+\nu_t)\bar{\bar{S}}(\boldsymbol{v})) = 0, \qquad (29)$$

$$\nabla \cdot \boldsymbol{v} = 0, \qquad (30)$$

$$\boldsymbol{v} \cdot \nabla k - 2\nu_t(\bar{\bar{S}}(\boldsymbol{v}) : \bar{\bar{S}}(\boldsymbol{v})) + \epsilon - \nabla \cdot \left(\left(\nu + \frac{\nu_t}{\sigma_k} \right) \nabla k \right) = 0, \quad (31)$$

$$\boldsymbol{v} \cdot \nabla \epsilon - 2C_{1\epsilon} \frac{\epsilon \nu_t}{k} (\bar{\bar{S}}(\boldsymbol{v}) : \bar{\bar{S}}(\boldsymbol{v})) + C_{2\epsilon} \frac{\epsilon^2}{k} - \nabla \cdot \left(\left(\nu + \frac{\nu_t}{\sigma_\epsilon} \right) \nabla \epsilon \right) = 0, \quad (32)$$

where $\boldsymbol{v} = [v_x, v_y]$ denotes the velocity vector, p denotes pressure, ρ is the density, ν is the kinematic viscosity, and $\overline{\bar{S}}$ is the strain rate tensor given by

$$\bar{\bar{S}}(\boldsymbol{v}) = \frac{1}{2} (\nabla \boldsymbol{v} + (\nabla \boldsymbol{v})^T).$$

In the $k - \epsilon$ turbulence model, k denotes the turbulent kinetic energy, ϵ denotes the turbulent dissipation, and ν_t denotes the turbulent kinematic viscosity, defined as

$$\nu_t = C_\mu \frac{k^2}{\epsilon}.\tag{33}$$

The coefficients C_{μ} , $C_{1\epsilon}$, $C_{2\epsilon}$, σ_k , σ_{ϵ} in (31)–(33) are either considered as uncertain parameters, or are functions of uncertain parameters, as detailed in Section 5.3.

At the inlet surface $\Gamma_{\rm in}$ Dirichlet boundary conditions are imposed. Data obtained by the direct numerical simulation described in [22] (Reynolds number 10,000) are used to determine reference inlet profiles for velocity, $v_{\rm ref}$, and for turbulent kinetic energy, $k_{\rm ref}$. Inlet conditions are allowed to vary by defining a velocity intensity (U) scale, which

¹We use σ_k and σ_{ϵ} here as model coefficients, which is typical notation in fluids community. These are only used in this section, and throughout the paper σ 's are variances.

is applied to the reference profiles. Turbulent dissipation at the inlet is estimated by assuming a mixing length model. Thus, the boundary conditions at the inlet surface are given by

$$oldsymbol{v}|_{\Gamma_{
m in}} = Uoldsymbol{v}_{
m ref}, \qquad \qquad k|_{\Gamma_{
m in}} = U^2 k_{
m ref}, \qquad \qquad \epsilon|_{\Gamma_{
m in}} = C_\mu rac{k^{3/2}}{\ell_m},$$

where ℓ_m denotes the mixing length parameter.

At the symmetry axis surface, Γ_{sym} , no-flux boundary conditions are imposed through a combination of Dirichlet and Neumann conditions of the form

$$v_y|_{\Gamma_{\text{sym}}} = 0, \qquad \left. \frac{\partial v_x}{\partial n} \right|_{\Gamma_{\text{sym}}} = 0, \qquad \left. \frac{\partial k}{\partial n} \right|_{\Gamma_{\text{sym}}} = 0, \qquad \left. \frac{\partial \epsilon}{\partial n} \right|_{\Gamma_{\text{sym}}} = 0.$$

Finally, at the surface $\Gamma_{\rm ff}$ "far-field" conditions that allow the entrainment of air around the jet are imposed through weak Dirichlet conditions, as detailed in [43].

The complete model includes additional features that make it more amenable to numerical discretization. The most delicate issue in the solution of the RANS model is the possible loss of positivity of the turbulence variables. To avoid this issue, we introduce an appropriately mollified (and thus smoothly differentiable) max function to ensure positivity of k and ε . In addition, if inflow is detected at any point on the far-field boundary, the boundary condition is switched from Neumann to Dirichlet by means of a suitably mollified indicator of the inflow region. Finally, we stabilize the discrete equations using a strongly consistent stabilization technique (Galerkin Least Squares, GLS, stabilization) to address the convection-dominated nature of the RANS equations. The complete formulation is shown in [43].

The model equations described above are solved numerically using a finite element discretization. The discretization is implemented in FEniCS [1] by specifying the weak form of the residual, including the GLS stabilization and mollified versions of the positivity constraints on k and ϵ and the switching boundary condition on the outflow boundary. To solve the nonlinear system of equations that arise from the finite element discretization, we employ a damped Newton method. The bilinear form of the state Jacobian operator is computed using FEniCS's symbolic differentiation capabilities. Finally, we use pseudo-time continuation to guarantee global convergence of the Newton method to a physically stable solution (if such solution exists) [20]. The finite element solver is detailed in [43].

5.3 Uncertain parameters and quantity of interest

In this investigation five uncertain parameters are considered: velocity intensity at inlet² (U), mixing length at inlet (ℓ_m), and the $k - \epsilon$ turbulence model coefficients C_{μ} , $C_{2\epsilon}$, and σ_k :

$$\mathbf{z} = [U, \ell_m, C_\mu, C_{2\epsilon}, \sigma_k].$$

²Since we keep other physical parameters constant, by varying the velocity intensity we effectively change the Reynolds number.

The parameter domain is $\mathbf{z} \in \mathcal{D} = [0.5, 1.5] \times [0.05, 0.15] \times [0.01, 0.15] \times [1.1, 2.5] \times [0.5, 2.5]$, and the nominal distribution of parameters is uniform in \mathcal{D} .

The other two coefficients of the $k - \epsilon$ turbulence model, $C_{1\epsilon}$ and σ_{ϵ} , are also uncertain but do not vary independently. According to Dunn et al. [12], empirical evidence suggests that $C_{1\epsilon}$ is related to $C_{2\epsilon}$ by

$$C_{1\epsilon} = \frac{C_{2\epsilon} - 0.8}{1.8}$$

In addition, as noted in [12, 16], the log-law implies that σ_{ϵ} must follow from

$$\sigma_{\epsilon} = \frac{\kappa^2}{\sqrt{C_{\mu}}(C_{2\epsilon} - C_{1\epsilon})}$$

where $\kappa = 0.41$ is the von Kárman constant.

The quantity of interest is the integral jet width measured at x = 27.5D:

$$w(\boldsymbol{v}; \mathbf{z}) = \frac{1}{v_{x_0} D} \int_0^{10D} v_x(x = 27.5D, y; \mathbf{z}) \, dy, \tag{34}$$

where $v_{x_0} = v_x(x = 27.5D, y = 0; \mathbf{z})$. Figure 6 illustrates a typical solution behavior for this turbulent jet by plotting contours of the turbulent kinetic energy for selected samples in \mathcal{D} .

5.4 Simplified-physics surrogate models

We consider four surrogate models to represent the dynamics of the free plane jet flow. The models are based on two distinct computational grids (fine and coarse), and on two representations of turbulence effects. The fine computational grid contains 10,000 elements and 5,151 nodes, while the coarse grid contains 2,500 elements and 1,326 nodes. Furthermore, the models are based either on the complete $k - \epsilon$ turbulence model described in the previous section, or on a prescribed turbulent viscosity field.

In the latter case, the turbulent viscosity field is estimated by a linear interpolation based on 243 conditions that span the input parameter space \mathcal{D} on a uniform grid (3 points along each of the 5 dimensions). At each of these 243 conditions, the turbulent viscosity field is computed with the $k - \epsilon$ turbulence model and the fine computational grid.

The following four low-fidelity models are increasingly complex in terms of either modeled physics or grid resolution:

- LFM1–CI: Coarse, interpolated; combines the interpolated turbulence viscosity field with the coarse computational grid (3,978 degrees of freedom); average computational time 25s;
- LFM2–FI: Fine, interpolated; combines the interpolated turbulence viscosity field with the fine computational grid (15,453 degrees of freedom); average computational time 72s;



- Figure 6: Samples of the flow solution computed at different points of the input parameter space \mathcal{D} . The plots show contours of turbulent kinetic energy and velocity streamlines. The white bars denote the integral jet width w associated with each case.
 - LFM3–CKE: Coarse $k \epsilon$; combines the $k \epsilon$ turbulence model with the coarse computational grid (6,630 degrees of freedom); average computational time 109s;
 - **HFM**: High-fidelity model; combines the $k \epsilon$ turbulence model with the fine computational grid (25,755 degrees of freedom); average computational time 590s.

Note that the models based on an interpolated turbulent viscosity field run four to eight times faster than the corresponding models based on the $k - \epsilon$ turbulence model.

This speedup results from eliminating (31)–(32) from the governing equations, which leads to a reduction in the total number of degrees of freedom (elimination of variables k and ϵ) and simplifications in the numerical discretization.

Let v_i , i = HFM, LFM1, LFM2, LFM3, denote the velocity field computed with the models above. The high-fidelity model is the mapping from the inputs to the quantity of interest (jet width from (34)) for a velocity field computed with the most complex representation of the flow dynamics, v_{HFM} :

$$f: \mathcal{D} \mapsto \mathbb{R}, \qquad f(\mathbf{z}) = w(\boldsymbol{v}_{\text{HFM}}; \mathbf{z}).$$

The surrogate models are defined in a similar fashion as

$$f^{(i)}: \mathcal{D} \mapsto \mathbb{R}, \qquad f^{(i)}(\mathbf{z}) = w(\boldsymbol{v}_i; \mathbf{z}), \quad i = \text{LFM1}, \text{LFM2}, \text{LFM3}$$

5.5 Results for multifidelity fusion of small probability estimators

We define a design failure when the jet width is below the value 0.98. Hence, the limit state function is given by

$$g(f(\mathbf{z})) = f(\mathbf{z}) - 0.98.$$
 (35)

We compute the biasing distributions q_i for i = LFM1, LFM2, LFM3 from the three low-fidelity surrogate models via MFIS (see Section 2.3). For each surrogate, we draw $\hat{m} = 20,000$ parameter samples from the uniform distribution on \mathcal{D} and evaluate the limit state function applied to the resulting quantity of interest. If the limit state function indicates failure of the system for a solution obtained from the *i*th surrogate model, the corresponding parameter is added to $\mathcal{G}^{(i)}$, the failure set computed from the *i*th surrogate model. We then fit a multivariate Gaussian to the samples in $\mathcal{G}^{(i)}$, resulting in the biasing densities $q_{\text{LFM1}}, q_{\text{LFM2}}, q_{\text{LFM3}}$.

Evaluation of the limit state function with the threshold value of 0.98 resulted in few samples in the failure region, so we increased it to 1.12 to obtain more samples to compute the biasing density from. For the three surrogate models and the high-fidelity model, the $\hat{m} = 20,000$ evaluations yield 21, 21, 30 and 76 samples, respectively, where the QoI falls below that increased threshold. This strategy yields an efficient biasing density as we see below. As reference, we repeat the same process with the high-fidelity model, resulting in the biasing distribution $q_{\rm HFM}$.

First, we investigate the quality of the biasing distributions. For reference, Figure 7, left, shows the result of 10^3 uniform sample evaluations with the four computational models. Note that hardly any samples are below the failure threshold. In contrast, the quantity of interest computed from samples of the four biasing distributions is shown in Figure 7, right. The biasing distributions give between 10%-50% of the 1000 drawn samples in the failure domain. Note, that the y-axis scaling of both figures is different, which also shows that the biased samples result in a tighter range of QoI values than the unbiased samples. Thus, the biasing distributions are indeed biased towards the failure region, and therefore the multifidelity strategy provides a viable way of saving computational time to inform a biasing distribution.



Figure 7: Quantity of interest, the width of the jet at x = 27.5D, for $n = 10^3$ samples. Left: The input parameters are drawn from the nominal distribution. Right: The input parameters are drawn from different biasing distributions. Note the large portion of samples falling below the failure threshold, and the different scaling of the *y*-axis.

The reference failure probability is computed via importance sampling with $n = 10^4$ samples drawn from the HFM biasing distribution and is $\hat{P}_{7,500,q_{\rm HFM}}^{\rm IS} = 7.25 \times 10^{-4}$. We compute the estimators $P_{n_i}^{\rm IS}$, i = 1, 2, 3 with n_i samples using the biasing densities $q_{\rm LFM1}$, $q_{\rm LFM2}$, $q_{\rm LFM3}$ derived from the three surrogate models. We obtain the fused multifidelity estimator P_{α} as described above in Algorithm 1 with $n_i = \lfloor n/3 \rfloor$, i = 1, 2, 3, samples by fusing the three surrogate-model-based importance sampling estimators. The fused estimator thus uses a total of n samples. We compare these estimators with an estimator that uses n samples from the HFM biasing density $q_{\rm HFM}$. The estimators and the error measures are averaged over three independent runs.

The coefficient of variation (24) is shown in Figure 8, left. The biasing density derived from the high-fidelity model yields the best estimator among all the models, as expected. The fused estimator yields a better coefficient of variation than LFM2 and LFM3, shows almost identical convergence as the estimator using $q_{\rm LFM1}$. Table 6 shows the three weights for the fused estimator \hat{P}_{α} as given in Proposition 1, according to which, the estimates with the lowest variance get assigned the largest weights.

	n = 300	n = 600	n = 900	n = 1200
α_1	0.500	0.502	0.610	0.900
α_2	0.448	0.044	0.055	0.057
α_3	0.052	0.453	0.335	0.043

Table 6: Weights of the fused estimator \hat{P}_{α} with *n* samples.

The CPU-hours to compute the biasing densities via this approach are shown in Fig-

ure 8, right. Since MC methods are embarrassingly parallel, any practical implementation can take advantage of this. Our numerical experiments were parallelized on a computing cluster with 55 nodes. Each node is a quad-core Intel Xenon E5-1620 with 3.6 GHz and 10MB Cache. The nodes have either 32GB or 64GB RAM. To put the CPU-hours savings achieved by using the high-fidelity model versus the lower-fidelity models to construct biasing densities into perspective, we see that using LFM1 reduces the computational cost by 96%, LFM2 by 88% and LFM3 by 81.5%. If we are using a fused estimator of all three models, we still save more than 65% computational effort compared to using the HFM, see Figure 8, right. This significant time difference can have important implications for engineering practice, as it translates into faster evaluation time and savings in CPU-hours .



Figure 8: Free plane jet application with five uncertain parameters. Left: Coefficient of variation for the different estimators. Right: CPU-hours for construction of the biasing densities used in the estimators in the left plot.

6 Conclusions

We enabled the estimation of small probabilities for expensive-to-evaluate models via a new approach drawing from importance sampling, multifidelity modeling and information fusion. The effectiveness of our proposed approach is demonstrated on a convectiondiffusion-reaction PDE, where asymptotic numerical results could be obtained. The strength of the proposed framework is then shown on the target application of the turbulent jet, a challenging problem for small-probability computation due to its high computational cost. The proposed framework was illustrated for the special case of importance-sampling based estimators, but applies to a much broader class of estimators, as long as the estimators are unbiased. An investigation of correlated estimators and the effect of correlation for variance reduction would be an interesting future direction. By fusing different estimators, we avoid the difficult biasing density selection problem. We also showed that this strategy always outperforms sampling from the worst biasing density. The numerical results suggest that the fused estimator is often comparable to an estimator that samples from the best biasing density only.

Acknowledgements

The authors thank Prof. M. Klein for sharing the DNS data in [22] with us. This work was supported by the Defense Advanced Research Projects Agency [EQUiPS program, award W911NF-15-2-0121, Program Manager F. Fahroo]; the Air Force [Center of Excellence on Multi-Fidelity Modeling of Rocket Combustor Dynamics, award FA9550-17-1-0195]; and the US Department of Energy, Office of Advanced Scientific Computing Research (ASCR) [Applied Mathematics Program, awards DE-FG02-08ER2585 and DE-SC0009297, as part of the DiaMonD Multifaceted Mathematics Integrated Capability Center].

References

- M. S. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells. The FEniCS project version 1.5. Archive of Numerical Software, 3(100), 2015.
- [2] L. J. Aslett, T. Nagapetyan, and S. J. Vollmer. Multilevel Monte Carlo for reliability theory. *Reliability Engineering & System Safety*, 165:188–196, 2017.
- [3] S.-K. Au and J. L. Beck. Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics*, 16(4):263–277, 2001.
- [4] P. Benner, M. Ohlberger, A. Cohen, and K. Willcox. Model Reduction and Approximation: Theory and Algorithms. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017.
- [5] G. Berkooz, P. Holmes, and J. L. Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics*, 25(1):539–575, 1993.
- [6] W. Betz. Bayesian inference of engineering models. PhD thesis, Technische Universität München, 2017.
- [7] M. Buffoni and K. Willcox. Projection-based model reduction for reacting flows. In 40th Fluid Dynamics Conference and Exhibit, page 5008, 2010.
- [8] P. Chen and A. Quarteroni. Accurate and efficient evaluation of failure probability for partial different equations with random input data. *Computer Methods in Applied Mechanics and Engineering*, 267:233–260, 2013.
- [9] R. T. Clemen and R. L. Winkler. Combining probability distributions from experts in risk analysis. *Risk analysis*, 19(2):187–203, 1999.

- [10] B. Cuenot and T. Poinsot. Asymptotic and numerical study of diffusion flames with variable Lewis number and finite rate chemistry. *Combustion and Flame*, 104(1):111–137, 1996.
- [11] M. de Angelis, E. Patelli, and M. Beer. Advanced line sampling for efficient robust reliability analysis. *Structural Safety*, 52:170–182, 2015.
- [12] M. C. Dunn, B. Shotorban, and A. Frendi. Uncertainty quantification of turbulence model coefficients via latin hypercube sampling method. *Journal of Fluids Engineering*, 133(4):041402–041402–7, 2011.
- [13] D. Elfverson, F. Hellman, and A. Mlqvist. A multilevel Monte Carlo method for computing failure probabilities. SIAM/ASA Journal on Uncertainty Quantification, 4(1):312–330, 2016.
- [14] F. Fagerlund, F. Hellman, A. Mlqvist, and A. Niemi. Multilevel Monte Carlo methods for computing failure probability of porous media flow systems. Advances in Water Resources, 94:498 – 509, 2016.
- [15] M. B. Giles. Multilevel Monte Carlo path simulation. Operations Research, 56(3):607–617, 2008.
- [16] S. Guillas, N. Glover, and L. Malki-Epshtein. Bayesian calibration of the constants of the $k\epsilon$ turbulence model for a CFD model of street canyon flow. Computer Methods in Applied Mechanics and Engineering, 279:536 553, 2014.
- [17] E. Gutmark, M. Wolfshtein, and I. Wygnanski. The plane turbulent impinging jet. Journal of Fluid Mechanics, 88(4):737756, 1978.
- [18] E. Gutmark and I. Wygnanski. The planar turbulent jet. Journal of Fluid Mechanics, 73(3):465495, 1976.
- [19] J. S. Jan S Hesthaven, G. Rozza, and B. Stamm. Certified Reduced Basis Methods for Parametrized Partial Differential Equations. Springer, 2016.
- [20] C. T. Kelley and D. E. Keyes. Convergence analysis of pseudo-transient continuation. SIAM Journal on Numerical Analysis, 35(2):508–523, 1998.
- [21] M. Klein, A. Sadiki, and J. Janicka. Investigation of the influence of the Reynolds number on a plane jet using direct numerical simulation. *International Journal of Heat and Fluid Flow*, 24(6):785 – 794, 2003.
- [22] M. Klein, C. Wolff, and E. Tangermann. A-priori analysis of a les model for scalar flux based on interscale energy transfer. In 9th International Symposium on Turbulence and Shear Flow, 2015.
- [23] A. Krothapalli, D. Baganoff, and K. Karamcheti. On the mixing of a rectangular jet. Journal of Fluid Mechanics, 107:201220, 1981.

- [24] J. Kutz, S. Brunton, B. Brunton, and J. Proctor. Dynamic Mode Decomposition. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2016.
- [25] B. E. Launder and D. B. Spalding. The numerical computation of turbulent flows. Computer Methods in Applied Mechanics and Engineering, 3(2):269 – 289, 1974.
- [26] J. Li, J. Li, and D. Xiu. An efficient surrogate-based method for computing rare failure probability. *Journal of Computational Physics*, 230(24):8683–8697, 2011.
- [27] J. Li and D. Xiu. Surrogate based method for evaluation of failure probability under multiple constraints. SIAM Journal on Scientific Computing, 36(2):A828– A845, 2014.
- [28] F. Marín-Martínez and J. Sánchez-Meca. Weighting by inverse variance or by sample size in random-effects meta-analysis. *Educational and Psychological Measurement*, 70(1):56–73, 2010.
- [29] P. Meier. Variance of a weighted mean. *Biometrics*, 9(1):59–73, 1953.
- [30] A. Narayan, Y. Marzouk, and D. Xiu. Sequential data assimilation with multiple models. *Journal of Computational Physics*, 231(19):6401 – 6418, 2012.
- [31] A. O'Hagan, C. E. Buck, A. Daneshkhah, J. R. Eiser, P. H. Garthwaite, D. J. Jenkinson, J. E. Oakley, and T. Rakow. Multiple experts. In *Uncertain Judgements: Eliciting Experts' Probabilities*, pages 179–192. John Wiley and Sons, Ltd, 2006.
- [32] A. Owen and Y. Zhou. Safe and effective importance sampling. Journal of the American Statistical Association, 95(449):135–143, 2000.
- [33] A. B. Owen. Monte Carlo theory, methods and examples. 2013.
- [34] I. Papaioannou, W. Betz, K. Zwirglmaier, and D. Straub. MCMC algorithms for subset simulation. *Probabilistic Engineering Mechanics*, 41:89–103, 2015.
- [35] B. Peherstorfer, T. Cui, Y. Marzouk, and K. Willcox. Multifidelity importance sampling. Computer Methods in Applied Mechanics and Engineering, 300:490 – 509, 2016.
- [36] B. Peherstorfer, B. Kramer, and K. Willcox. Combining multiple surrogate models to accelerate failure probability estimation with expensive high-fidelity models. *Journal of Computational Physics*, 341:61–75, 2017.
- [37] B. Peherstorfer, B. Kramer, and K. Willcox. Multifidelity preconditioning of the cross-entropy method for rare event simulation and failure probability estimation. *SIAM/ASA Journal of Uncertainty Quantification*, 6(2):737–761, 2018.
- [38] B. Peherstorfer, K. Willcox, and M. Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. SIAM Review, 60(3):550–591, 2018.

- [39] C. L. Ribault, S. Sarkar, and S. A. Stanley. Large eddy simulation of a plane jet. *Physics of Fluids*, 11(10):3069–3083, 1999.
- [40] G. Schuëller, H. Pradlwarter, and P. Koutsourelakis. A critical appraisal of reliability estimation procedures for high dimensions. *Probabilistic Engineering Mechanics*, 19(4):463–474, 2004.
- [41] S. A. Stanley, S. Sarkar, and J. P. Mellado. A study of the flow-field evolution and mixing in a planar turbulent jet using direct numerical simulation. *Journal of Fluid Mechanics*, 450:377407, 2002.
- [42] E. Ullmann and I. Papaioannou. Multilevel estimation of rare events. SIAM/ASA Journal on Uncertainty Quantification, 3(1):922–953, 2015.
- [43] U. Villa and A. N. Marques. An UQ-ready finite element solver for a twodimensional RANS model of free plane jets, 2017.
- [44] X. Zhou, Z. Sun, F. Durst, and G. Brenner. Numerical simulation of turbulent jet flow and combustion. *Computers & Mathematics with Applications*, 38(9):179– 191, 1999.
- [45] K. M. Zuev, J. L. Beck, S.-K. Au, and L. S. Katafygiotis. Bayesian post-processor and other enhancements of subset simulation for estimating failure probabilities in high dimensions. *Computers & Structures*, 92:283–296, 2012.