

Quantification of Margins and Uncertainties using an Active Subspace Method for Approximating Bounds

P. G. Constantine

*Department of Applied Mathematics and Statistics
Colorado School of Mines, Colorado, USA*

M. Emory, F. Palacios, N. Kseib, G. Iaccarino

*Department of Mechanical Engineering
Stanford University, California, USA.*

ABSTRACT: In many engineering application, bounds on the model output are sought to characterize safe operating conditions of a system. When a complex model contains many inputs, determining such bounds may be intractable. However, if one can determine that the model output varies with gross, global monotonic trends along a low-dimensional *active subspace* of the high-dimensional input space, then this information can be exploited to find corners of the input space that likely correspond to bounds on the output. We present a method for determining the active subspace, and we use it to determine the safe operating conditions of a hypersonic scramjet vehicle with 19 independent input parameters.

1 INTRODUCTION

Safety margins are routinely used in the design and operation of complex engineering systems. Although margins are typically estimated based on experience, the expansion of simulation capabilities and the cost of performing controlled tests provides fertile ground to apply advanced computational models. The determination of the safety margin requires the evaluation of bounds on performance metrics (model outputs). Realistic systems are characterized by a large number of components and their behavior is the result of complex interactions of many environmental and manufacturing factors and designers choices (model inputs) which are affected by uncertainties. Quantification of Margins and Uncertainties (QMU) is an engineering tool that assists designers in determining the operating envelope of a system balancing the need for a safety margin to operate away from a failure region with the presence of uncertainties.

We are primarily concerned with the estimation of the unstart limit of a supersonic combustion engine (scramjet)—a critical failure mode due to excessive increase of fuel injection in the combustor and subsequent thermal choking. The unstart process is rapid and leads to sudden appearance of strong shock waves, changes in the flow regime, and severe structural loads. The detection of unstart is typically based on the definition of early indicators based on wall-

mounted pressure sensors within the scramjet. We have employed QMU (Iaccarino, Pecnik, Glimm, & Sharp 2010) to assess the effect of uncertainties on the maximum pressure loading within the chamber using simple computational models and straightforward uncertainty quantification techniques. The use of detailed three-dimensional reactive flow simulations requires the introduction of efficient methods to determine the scramjet performance envelope under uncertainty.

In general, the computation of bounds on the output of a model over the space of model inputs is easily formulated as an optimization problem. However, this optimization problem is very difficult to solve when there are many model inputs—resulting in a high dimensional input space—and each evaluation of the model is computationally expensive. One approach to deal with these issues is to first find gross, global trends of the model output as a function of the inputs and exploit these trends to approximate the desired bounds. For example, if the gross trend of the model is monotonic with respect to some inputs, then the output evaluated at the bounds of the inputs may offer reasonable approximations to the bounds on the outputs.

In this paper we propose a method to uncover a gross trend of the model along a particular direction in the input space. The particular direction is the direction of primary variability of the function, which de-

finds an *active subspace* that contains the majority of the model output’s variability. Active subspace methods are used for such exploration in the 2010 Ph.D. thesis (Russi 2010), and the authors have developed a rigorous theoretical framework in (Constantine, Dow, & Wang 2013).

By approximating the model output within the dominant one-dimensional subspace of the input parameter space, one can plot the model output against the coordinate defining the subspace. Such a plot can reveal gross trends, e.g., monotonicity, of the function on the subspace. If the gross trends are determined to be monotonic—as is the case for many models in engineering—then the direction of primary variability can be used to determine the corners of the high dimensional input space where the bounds on the model output are likely to occur.

We apply this method to approximate the maximum exit pressure from a combustor on a hypersonic scramjet vehicle, where the model inputs are grouped into two sets: (i) seven parameters characterizing the inflow conditions of the scramjet combustor, and (ii) twelve reaction rates of a combustion model of fuel burning in the scramjet combustor.

2 AN ACTIVE SUBSPACE METHOD

In this section we describe a method for approximating the bounds on a quantity of interest (QOI) by approximating it on a low dimensional subspace—i.e., the *active subspace*—of the input space. Further details of the active subspace method can be found in (Constantine, Dow, & Wang 2013). Denote the QOI by the scalar function

$$f = f(\mathbf{x}), \quad \mathbf{x} \in \mathcal{X} \subset \mathbb{R}^m. \quad (1)$$

where \mathbf{x} are the natural inputs of the system. The parameter space \mathcal{X} is the m -dimensional hypercube $[-1, 1]^m$, which we can also write as $-1 \leq \mathbf{x} \leq 1$. Let \mathbf{x}_j with $j = 1, \dots, M$ be a set of design sites from \mathcal{X} . Assume we have computed the function value $f_j = f(\mathbf{x}_j)$ and the gradient

$$\nabla f_j = \begin{bmatrix} \frac{\partial f}{\partial x_1}(\mathbf{x}_j) \\ \vdots \\ \frac{\partial f}{\partial x_m}(\mathbf{x}_j) \end{bmatrix} \quad (2)$$

at each design site.

For the models of interest, f_j typically involves computing a derived quantity from an approximate solution of a partial differential equation with inputs \mathbf{x}_j . The gradient may be approximated via an adjoint solution (Bryson & Ho 1975) or algorithmic differentiation (Griewank 2000), if such techniques are available. If not, one may use finite difference approximations, which use $m + 1$ evaluations of $f(\mathbf{x})$ per design site \mathbf{x}_j .

2.1 Computing the dominant direction of variability

Using the gradient evaluations ∇f_j , we will compute the direction in \mathcal{X} along which $f(\mathbf{x})$ varies the most on average. We accomplish this by forming the $m \times M$ matrix \mathbf{G} defined as

$$\mathbf{G} = [\nabla f_1 \quad \cdots \quad \nabla f_M]. \quad (3)$$

We compute the singular value decomposition (SVD)

$$\mathbf{G} = \mathbf{U}\Sigma\mathbf{V}^T, \quad \Sigma = \text{diag}(\sigma_1, \dots, \sigma_m), \quad (4)$$

where \mathbf{U} has dimension $m \times m$ and \mathbf{V} has dimension $M \times m$. The first column of \mathbf{U} , denoted \mathbf{u}_1 , defines a subspace of \mathcal{X} along which $f(\mathbf{x})$ varies the most on average. The subspace is defined precisely as

$$\mathcal{Y} = \{y : y = \mathbf{u}_1^T \mathbf{x}, \mathbf{x} \in \mathcal{X}\}. \quad (5)$$

The singular values give a relative measure of how much f varies along \mathbf{u}_1 ; see (Constantine, Dow, & Wang 2013) for a precise statement. Note that \mathbf{u}_1 is a global quantity. If f is monotonic with respect to each component of \mathbf{x} (i.e., the partial derivatives do not change sign), then \mathbf{u}_1 is closely related to the average gradient of f over the space \mathcal{X} .

In practice, one need not compute the full SVD; more efficient methods exist for computing the dominant singular vector. However, even extreme cases with thousands of input parameters (m) and thousands of samples of the gradient (M) will yield a matrix \mathbf{G} small enough to easily compute the SVD on a modern laptop.

2.2 Bootstrap for the dominant direction

The elements of \mathbf{u}_1 depend on the gradient samples ∇f_j used to construct \mathbf{G} . It can be shown that \mathbf{u}_1 has a limiting value as the number M of gradient samples increases. With finite samples, it is natural to ask if the gradient has been sufficiently sampled to uncover the true one-dimensional active subspace. Unfortunately, we typically have a limited budget of function evaluations, and checking convergence of \mathbf{u}_1 is not feasible. We propose to use the bootstrap technique (Efron & Tibshirani 1993) to estimate the variability in the computed components of \mathbf{u}_1 . Let

$$\boldsymbol{\pi}^k = [\pi_1^k \quad \cdots \quad \pi_M^k], \quad k = 1, \dots, N \quad (6)$$

be M -vectors of integers between 1 and M sampled with replacement. Define the matrices \mathbf{G}^k by

$$\mathbf{G}^k = \begin{bmatrix} \mathbf{G}_{1\pi_1^k} & \cdots & \mathbf{G}_{1\pi_M^k} \\ \vdots & \ddots & \vdots \\ \mathbf{G}_{m\pi_1^k} & \cdots & \mathbf{G}_{m\pi_M^k} \end{bmatrix}. \quad (7)$$

In words, \mathbf{G}^k is an $m \times M$ matrix whose columns are sampled with replacement from the columns of

G . Compute \mathbf{u}_1^k for each G^k and build bootstrap histograms for the components of \mathbf{u}_1 . One must be careful during this procedure because \mathbf{u}_1^k is unique up to its sign. Therefore, each \mathbf{u}_1^k should be normalized, e.g., so that its first component is positive.

2.3 Approximation and monotonicity

Given the vector \mathbf{u}_1 defining the direction of primary variability, we can project the design sites \mathbf{x}_j to the one-dimensional active subspace and plot the corresponding function evaluations. We have found that this visualization often provides great insight into the function. For example, one can often tell if the function is roughly monotonic along the active subspace of the inputs; we make this notion more precise below.

To project the design sites, compute

$$y_j = \mathbf{u}_1^T \mathbf{x}_j, \quad j = 1, \dots, M. \quad (8)$$

Then plot the pairs $\{y_j, f_j\}$ to visualize the function on the subspace. In many cases, gross trends become immediately apparent. Additionally, defining the function values on a low-dimensional subspace enables a host of analyses that are not feasible in an m -dimensional space, such as response surface construction or derivative approximation.¹ In particular, we can test the function values to see if the gross trends are monotonic on the subspace. In practice, the visualization may be sufficient to determine monotonicity of gross trends. Although, more rigorous statistical approaches are also available (Bowman, Jones, & Gijbels 1998).

2.4 Approximating bounds

If the gross trends are indeed monotonic on the one-dimensional subspace, and if the dominant subspace is well separated as measured by the singular values, then one may justifiably evaluate f at the corners of the full input space \mathcal{X} corresponding to the bounds of the active subspace and treat these function values as approximate bounds on the range of f . More specifically, define $\underline{\mathbf{x}}$ by

$$\underline{\mathbf{x}} = \underset{-1 \leq \mathbf{x} \leq 1}{\operatorname{argmin}} \mathbf{u}_1^T \mathbf{x}, \quad (9)$$

And define $\bar{\mathbf{x}} = -\underline{\mathbf{x}}$. Let $\bar{f} = f(\bar{\mathbf{x}})$ and $\underline{f} = f(\underline{\mathbf{x}})$. Then the range of $f(\mathbf{x})$ can be approximated by $[\bar{f}, \underline{f}]$ or $[\underline{f}, \bar{f}]$, depending on whether $\bar{f} < \underline{f}$ or vice versa. If the dominant subspace is not well separated, then $\bar{\mathbf{x}}$ and $\underline{\mathbf{x}}$ may still be excellent initial guesses for local optimization procedures, if available.

¹Note that integration on the subspace is substantially more complicated because the projection of the measure on \mathcal{X} requires a $m - 1$ -dimensional integral for every y .

3 APPLICATION TO HYSHOT II

The problem of interest is the study of the operability limits associated with the propulsion system of vehicles flying at high Mach numbers. In these circumstances, the engine operates under supersonic conditions (scramjet mode) and a key *performance metric* is the amount of heat released in the combustion chamber, as this is directly connected to the generation of thrust.

The engine injection system is carefully designed to mix the fuel and the incoming supersonic air stream and to produce auto-ignition of the mixture. This process is complicated by the presence of a shock wave system within the chamber (shock train) interacting with the viscous boundary layers formed on the walls. The vehicle thrust is directly related to the amount of heat release occurring in the chamber, but there is an upper limit corresponding to thermal choking: a normal shock and a consequent subsonic flow region is established in the combustor. In addition to a reduction in the thrust, this can lead to increased structural and thermal loads and eventually to failure. The normal shock can also propagate upstream in the combustion chamber creating extensive regions of boundary layer separation; moreover, the shock motion can lead to engine *unstart* with the entire shock train moving upstream. Under these conditions the vehicle performance is compromised and extreme actions have to be taken to *restart* the engine. Although the heat release is the critical performance metric, this is not easy to measure directly in ground laboratories or in flight. Instead, pressure sensors are embedded in the combustor walls and used to detect the occurrence of subsonic flow within the chamber as an early indicator to unstart. The objective of the present computations is to evaluate the highest pressure in the system—typically towards the exit of the combustor—and how it is affected by variability in the operating conditions and in the physical model inputs.

3.1 Uncertainties

We are concerned with the study of the Hyshot II engine which was analyzed experimentally in Germany at DLR (Schramm, Karl, Hannemann, & Steelant 2008). The engine was mounted in a high-enthalpy shock tube and the system performance were investigated for a variety of conditions, characterized by different mass fuel injection (or equivalence ratios). In this paper only the nominal design scenario is considered. Although the shock tube operation is closely controlled, the extreme thermodynamic conditions and the high associated cost do not allow to perform repeated tests thus introducing uncertainties and imprecisions.

We have considered seven parameters that describe the uncertainties associated with the operating conditions within the shock tube: the geometrical posi-

tion of the model within the facility (angle of attack, AoA), the thermodynamic conditions of the incoming airflow (stagnation pressure and enthalpy), the turbulence level (turbulence intensity and length scale) and finally the occurrence of transition from laminar to turbulent flow on the vehicle nose (ramp) and within the combustor (cowl). The characterization of the input ranges for each of these parameters is based on available information provided by DLR.

In addition to the uncertainty described above and hereafter referred to as *inflow uncertainties*, we have also considered the parameters required as inputs to the combustion model used in the computations, e.g. the reaction rates. The heat release is dominated by the exothermic reactions of the hydrogen with oxygen but the overall kinetic mechanism involves twelve reactions and associated Arrhenius rates. These rates are measured indirectly in dedicated facilities, but remain affected by uncertainties whose ranges are reported in the literature.

In summary, for the present study we have included 17 uncertainties divided in two sets of parameters: inflow ($m = 7$) and reaction rates ($m = 12$).

3.2 Applying the ASM

Next we apply the active subspace method separately for each set of uncertain input parameters, inflow parameters and chemical reaction rates. When applied to the inflow parameters, we use a nominal set of reaction rates; when applied to the reaction rates, we will use a nominal inflow condition. By treating these sets of inputs independently, we can analyze the exit pressure variability attributable to inflow profile versus the reaction rates.

3.2.1 Inflow parameters

There are $m = 7$ input parameters characterizing the inflow profile. We first shift and scale the parameter domain to the reference hypercube $[-1, 1]^7$. We then sample a finite difference approximation of the gradient at $M = 14$ locations in the hypercube to construct the matrix \mathbf{G} from (3). This required 14 runs at nominal design sites plus 7 perturbations per design site for a total of 112 simulations. Given that the simulation uses an iterative method to compute the solution and the derived exit pressure, we anticipate noise in the finite difference approximation due to finite iterations of the solver; we do not attempt to address this issue.

Given \mathbf{G} , we follow the recipe described in Section 2 to determine the dominant subspace, including the bootstrap estimates of density functions. The singular values of \mathbf{G} are shown in Figure 1(a), and the exit pressures from the nominal runs plotted against the reduced coordinate are shown in Figure 1(b). Figure 2 plots the bootstrap histograms for the components of the vector defining the dominant subspace. While there is significant variability, we are confident of the signs of the components, which will determine

the corners of the hypercube chosen to get approximate bounds.

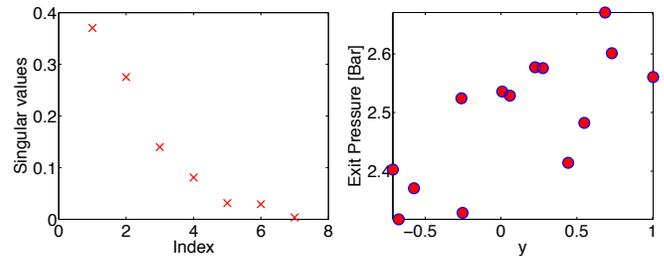


Figure 1: (Left) Singular values of the finite difference approximation to \mathbf{G} from (3) for the inflow parameters. (Right) Exit pressure plotted versus the reduced coordinate to show the dominant trend.

3.2.2 Reaction rates

We follow the same procedure for the reaction rates using a finite difference approximation to the matrix \mathbf{G} from (3). In this case, there are $m = 12$ parameters, and we sample the approximate gradients $M = 24$ times. This results in a total of 312 full simulations.

Figures 3(a) and 3(b) show, respectively, the singular values of \mathbf{G} and the exit pressures at the design sites plotted with respect to the reduced coordinate defined by \mathbf{u}_1 . The bootstrap densities for the components of \mathbf{u}_1 are shown in Figure 4.

3.2.3 An upper bound on exit pressure

Since we are primarily concerned with the safe operating conditions of the hypersonic vehicle, we seek only the location in the parameter space that corresponds to the approximate maximum of the exit pressure. The sign of the components of the dominant singular vector \mathbf{u}_1 determine the corner of the hypercube that we will use as inputs for the simulation to find an approximate maximum. At this point the separate analyses of each set of input parameters are combined to find an approximate maximum exit pressure over the total space of inputs.

The exit pressure that we find at the chosen corner of the hypercube is 2.6883 bar. As a simple check, the maximum exit pressure over the runs used to compute \mathbf{u}_1 is 2.6707 bar. Thus, we have found a particular corner of the 19-dimensional hypercube input space that yields a larger exit pressure than the initial set of designs used to determine the dominant one-dimensional subspace.

4 CONCLUSIONS

We have described and demonstrated a method for approximating bounds on a quantity of interest that depends on many input parameters. The method relies on determining the one-dimensional *active subspace* of the high dimensional space of input parameters. By visualizing the function on this one-dimensional

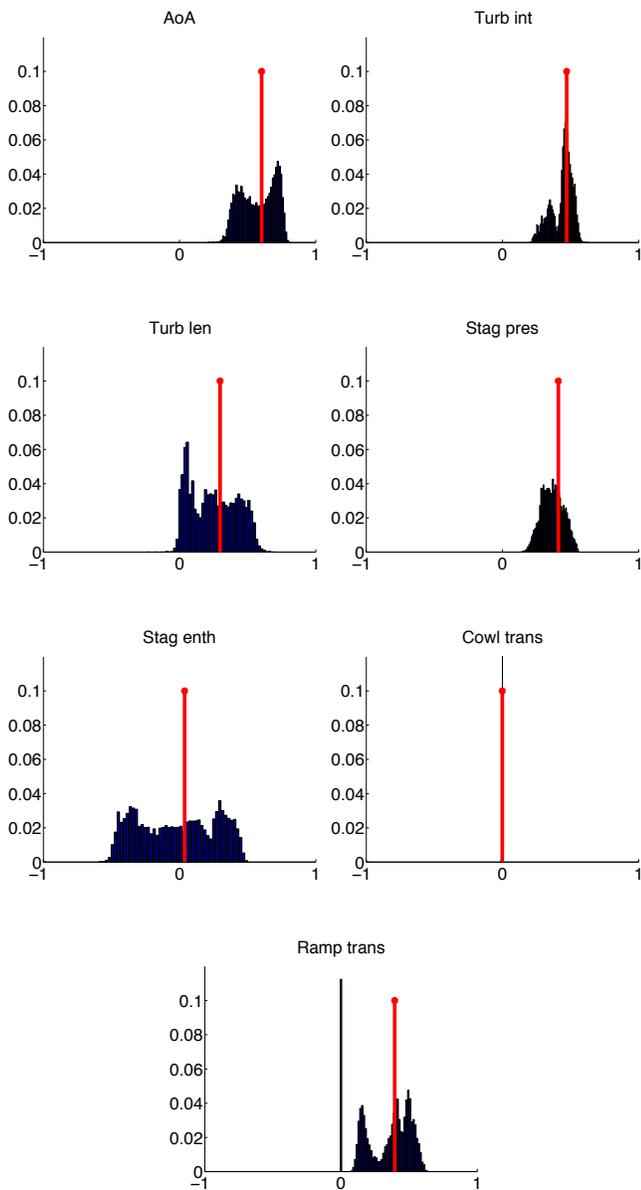


Figure 2: Bootstrap histograms of the components of the dominant direction \mathbf{u}_1 corresponding to the inflow parameters. The red stems are the computed components of \mathbf{u}_1 . The titles of each figure tell the specific inflow parameter.

subspace, one can uncover gross, global trends such as monotonicity. If the function appears to be monotonic on the subspace, and if the subspace is well separated, then one can use the direction defining the subspace to determine the extreme points on the parameter space (i.e., the corners of the hypercube) that yield approximations to the maximum and minimum of the quantity of interest over the parameter space. We have successfully applied this method to the hypersonic scramjet vehicle with two sets of input parameters to determine the safe operating conditions.

REFERENCES

- Bowman, A., M. Jones, & I. Gijbels (1998). Testing monotonicity of regression. *Journal of Computational and Graphical Statistics* 7(4), 489–500.
- Bryson, A. E. & Y.-C. Ho (1975). *Applied Optimal Control: Optimization, Estimation, and Control*. Hemisphere Publishing

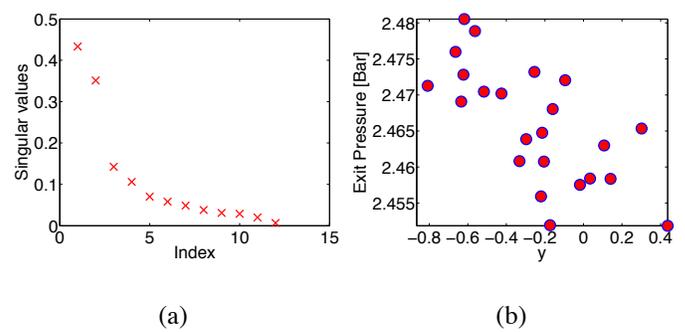


Figure 3: (Left) Singular values of the finite difference approximation to \mathbf{G} from (3) for the chemical reaction rates. (Right) Exit pressure plotted versus the reduced coordinate to show the dominant trend.

Corportation.

- Constantine, P. G., E. Dow, & Q. Wang (2013). Active subspace methods in theory and practice: applications to kriging surfaces. *arXiv preprint arXiv:1304.2070*.
- Efron, B. & R. Tibshirani (1993). *An introduction to the bootstrap*, Volume 57. Chapman & Hall/CRC.
- Griewank, A. (2000). *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. SIAM.
- Iaccarino, G., R. Pecnik, J. Glimm, & D. Sharp (2010). A qmu approach for characterizing the operability limits of air-breathing hypersonic vehicles. *Reliability Engineering & System Safety* 96(9), 1150–1160.
- Russi, T. M. (2010). *Uncertainty Quantification with Experimental Data and Complex System Models*. Ph. D. thesis, UC Berkeley.
- Schramm, J. M., S. Karl, K. Hannemann, & J. Steelant (2008). Ground testing of the hyshot ii scramjet configuration in heg. *AIAA Paper 2008 2547*.

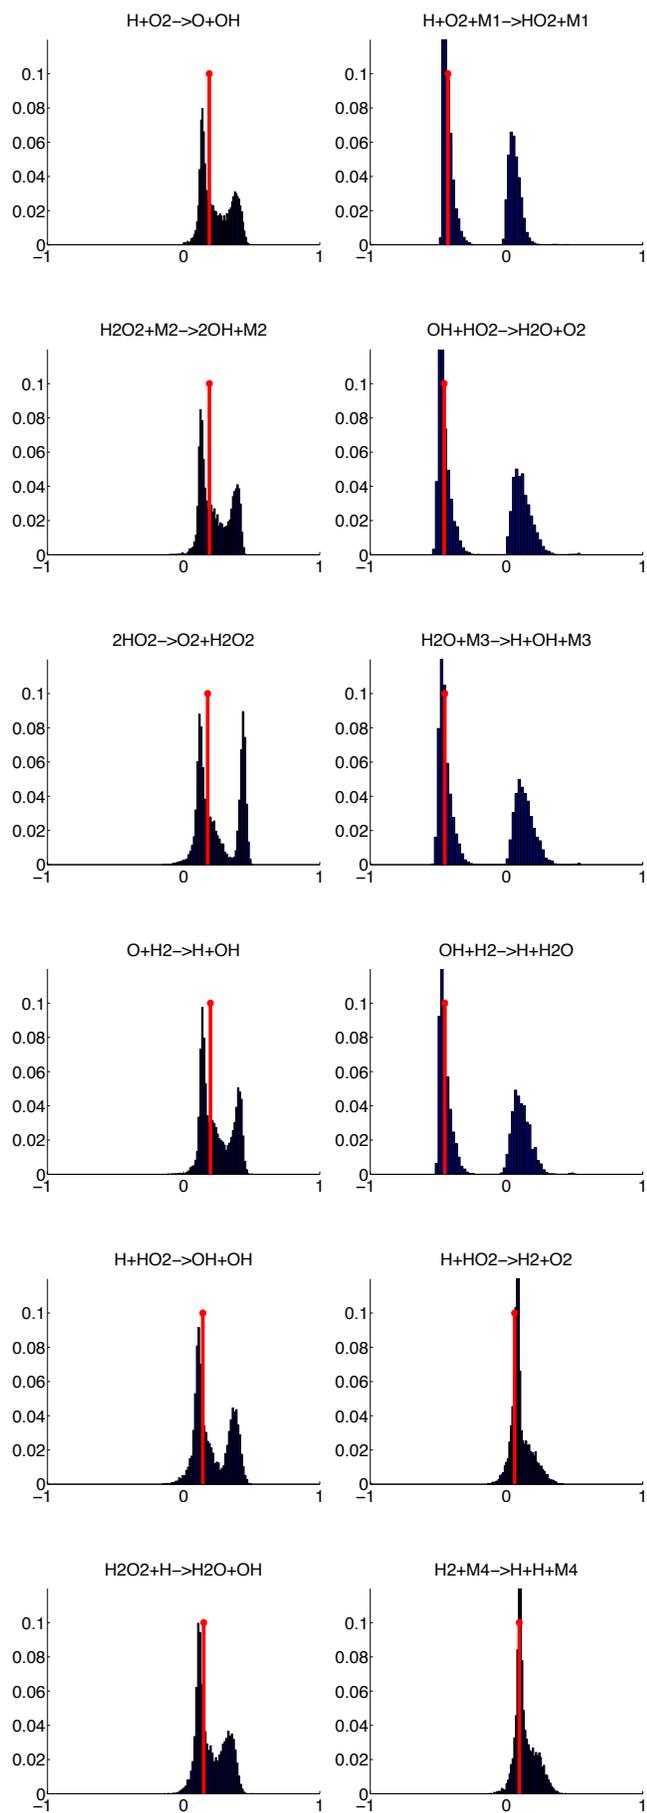


Figure 4: Bootstrap histograms of the components of the dominant direction \mathbf{u}_1 corresponding to the reaction rates. The red stems are the computed components of \mathbf{u}_1 . The titles of each figure tell the specific reaction.