On the control landscape topology

Eurika Kaiser [*] **Ruiying Li** [**] **Bernd R. Noack** [***]

[*] *University of Washington, Seattle, WA 98195 USA (e-mail: eurika@uw.edu)*
[**] *Institut PPRIME, CNRS – Université de Poitiers – ISAE-ENSMA, 86962 Futuroscope Chasseneuil, France (e-mail: ruiying.li@ensma.fr)*
[***] *LIMSI-CNRS, UPR 3251, 91405 Orsay cedex, France Technische Universitat Braunschweig, 38108 Braunschweig, Germany (e-mail: noack@limsi.fr)*

Abstract: Evolutionary algorithms are powerful tools to optimize parameters and structure of control laws. However, these approaches are often very costly, or even prohibitive, for expensive experiments due to long evaluation times and large population sizes. Reducing the learning time, e.g. by decreasing the number of function evaluations, is a challenging problem as it often requires additional knowledge on the objective function and assumptions. We address the need to analyze these algorithms and guide their acceleration through examination of the search space topology and the exploratory and exploitative nature of the genetic operators. We show how this gives insights on the convergence and performance behavior of Genetic Programming Control for the drag reduction of a car model (Li et al., 2016). Profiling machine learning algorithms, that are very powerful but also more complex to analyze, aids the goal to increase their performance and making them eventually feasible for a wide range of applications.

Keywords: Statistical data analysis, Evolutionary algorithms in control and identification, Knowledge discover (data mining), Information processing and decision support

1. INTRODUCTION

Feedback turbulence control is still a challenging endeavor due to the high-dimensionality, strong nonlinearities, and time delays of the fluid flows. Most successful control strategies employ adaptive or model-free approaches and target the optimization of open-loop controllers, not exploiting advantages of in-time flow response like disturbance mitigation. The recently introduced Genetic Programming Control (GPC) (Gautier et al., 2015; Parezanović et al., 2016; Duriez et al., 2016), inspired by machine learning techniques to learn optimal functions with respect to an objective function using evolutionary algorithms (Koza, 1992), has been shown to systematically learn sensor-based feedback control laws in an unsupervised manner by optimizing the structure of the control law. The reader is referred to (Brunton and Noack, 2015) for a recent review on closed-loop turbulence control. In this work, we address the need to tune GPC to decrease the learning time, which is particularly important for expensive experiments, where the large number of control laws to be evaluated and their evaluation time makes its application very costly, even prohibitive for certain applications.

The point of departure of GPC is a set of candidate control laws, referred to as *individuals*, which are evaluated to determine their performance, called *fitness*, and then evolved using genetic operations like mutation, replication, cross-over, and elitism (Koza, 1992). While this approach has been shown to be powerful in finding globally optimal control laws exploiting strong nonlinearities, this is accompanied by long learning times. The lower bound of the evaluation time of an individual is limited by the time duration needed to compute a confident measure of the average quantity of interest like the average drag or mean recirculation area. Thus, there is a strong focus on reducing the number of function evaluations, e.g., by discarding or approximating the fitness of individuals. Examples are based on a meta-model of the fitness function (Emmerich et al., 2002; Ziegler and Banzhaf, 2003; Jin et al., 2002) to interpolate between or classify individuals, a proxy fitness for un-evaluated functions using ancestry (Sastry et al., 2001) or neighborhood information (Kim, 2001), or employ other statistical and information-theoretic methods (Giacobini et al., 2002).

In this work, we analyze the performance of GPC through examination of the search space topology and the exploratory and exploitative nature of the genetic operators. We employ Multidimensional Scaling (MDS) (Mardia et al., 1979; Cox and Cox, 2000) to determine a low-dimensional subspace in which the distances between control laws are preserved. This requires a suitable definition of a (dis)similarity metric between control laws based on input-output data and their fitness. Delaunay triangulation is then used to approximate the fitness topology on which the evolution of the control laws is examined. Then, ancestry information is analyzed to determine the exploratory and exploitative nature of the genetic operators in the search space. In addition, we propose a performance estimator for untested control laws that is assessed in a post-processing analysis. We analyze data from applying GPC, particularly the recently introduced Linear Genetic

Programming Control (LGPC), to a turbulent bluff body flow to reduce drag (Li et al., 2016). This analysis gives insights into the degree of exploration due to the different genetic operators and may guide future improvements in the learning time.

In Sec. 2, we present a brief background on the tools employed. The main results and conclusions are provided in Sec. 3 .

2. FEATURE-BASED ANALYSIS AND COST ESTIMATOR

2.1 Similarity of control law functions

Let be a control law denoted by K. The time-dependent actuation is given by $b(t) = K(s(t))$ where $s(t)$ is the time-dependent sensor reading. In GPC, we have an ensemble of control laws $\{K^i\}_{i=1}^N$ with $N = N_I \times N_G$ where N_I is the number of individuals in a generation and N_G is the number of generations. Each control law K^i is evaluated in the experiment and a fitness J^i is assigned to it. The (dis)similarity between different control laws K^i and K^j is measured based on the time series information and their difference in performance. The latter incorporates that very similar control laws may yield very different performance. The elements of the square of the cross-generational distance matrix, $\mathbf{D} = (D_{ij})$, are defined by

$$D_{ij}^2 = \langle K^i(s), K^j(s)\rangle_{i,j} + \alpha |J^i - J^j| \tag{1}$$

where $\langle \bullet \rangle$ denotes a suitable time or ensemble average of the differences between control laws K^i and K^j. The first term in (1) is given by

$$\begin{aligned}\langle K^i(s), K^j(s)\rangle_{i,j} = & \frac{1}{2M}\sum_{m=1}^{M}\left[|K^i(s^i(t_m)) - K^j(s^i(t_m))|^2\right. ,\\ + & \left.|K^i(s^j(t_m)) - K^j(s^j(t_m))|^2\right]\end{aligned} \tag{2}$$

where $s^i(t_m)$ are the sensor readings collected at discrete time t_m, $m = 1, \ldots, M$, when K^i was applied.

Equation (2) represents the difference between control laws K^i and K^j in an average sense evaluated in the relevant sensor space and considering both forced attractors with equal probability.

2.2 Multidimensional Scaling

Multidimensional Scaling (MDS) comprises a collection of algorithms to determine a meaningful low-dimensional embedding from a given distance matrix, in which the similarity between objects is depicted by their mutual distances. These geometry-preserving algorithms are particularly useful for visualization purposes of the relative distances or similarity of high-dimensional data objects. We employ particularly Classical Multidimensional Scaling (CMDS) which originated from the works of Young and Householder (1938) and Schoenberg (1935). The aim of CMDS is to find a centered representation of points $\mathbf{\Gamma} = [\boldsymbol{\gamma}_1 \quad \ldots \boldsymbol{\gamma}_N]$ with $\boldsymbol{\gamma}_i \in \mathbb{R}^r$ where $r = 2$ here for visualization purposes, such that the pairwise distances of the objects approximate the true distances, i.e. $||\boldsymbol{\gamma}^i - \boldsymbol{\gamma}^j||_2 \approx D_{ij}$.

2.3 Determining the search space topology

The search space topology shall be approximated by the evaluated control laws, particularly, by fitting a surface of the form $J = J(\gamma_1, \gamma_2)$ to the scattered data points given by $(\gamma_1^i, \gamma_2^i, J^i)$ for all control laws K^i, $i = 1, \ldots, N$. This is based on a Delaunay triangulation of the point ensemble $\mathbf{\Gamma}$. A triangulation-based cubic interpolation (interpolating surface is $C2$ continuous) is employed for approximation which yields for any point (γ_1, γ_2) within the convex hull of the dataset $\mathbf{\Gamma}$ a unique value for J.

2.4 Analysis of genetic operators

The exploratory and exploitative nature of the genetic operators is examined by a statistical analysis of the ancestry of control laws. For this purpose, a coarse grid of N_f square fields on the $(\gamma_1, \gamma_2)-$space is defined, which elements are denoted by f_m, $m = 1, \ldots, N_f$. The conditional probabilities from parent to offspring generation under different operators is given by

$$P_{mn}^{\dagger} = \text{Prob}\,(\text{offspring in } f_m|\text{parent in } f_n, \dagger)\,, \tag{3}$$

with $m, n = 1, \ldots, N_f$. The symbol $\dagger$ stands for one of the four genetic operations and can assume letters E for elitism, R for replication, C for cross-over, or M for mutation. The probabilities are computed as relative frequencies that a transition occurs and are column-stochastic, i.e. $\sum_{m=1}^{N_f} P_{mn}^{\dagger} = 1$ for all n and operators $\dagger$.

In particular, we are interested in the probabilities that a parent lies in any of the tiles f_k,

$$P_{mn}^{\dagger,\text{from}} = \text{Prob}\,(\text{parent in } f_m|\dagger)\,, \tag{4}$$

and where the offspring will be located

$$P_{mn}^{\dagger,\text{to}} = \text{Prob}\,(\text{offspring in } f_m|\dagger)\,, \tag{5}$$

under the action of the genetic operators. A more uniform distribution is an indicator for the exploratory nature of the operator, while a peak probability in a certain region demonstrates its exploitative action.

2.5 Performance estimator for untested control laws

The increasing amount of information collected by testing generations of individuals can be leveraged to develop a performance estimator for newly bred, untested individuals to discard potentially similarly performing control laws and to specifically select exploratory or exploitative control laws before their evaluation. Given the proximity map obtained from applying CMDS to those generations which are already evaluated, the estimator shall determine the location $\hat{\boldsymbol{\gamma}}^k$, where the symbol '$\hat{}$' refers to estimated values, of the untested individual. Selection criteria based on the proximity to other control laws can then be applied to pick control laws for the next generation to be evaluated. For this purpose, we employ Sparse Landmark Multidimensional Scaling (sLMDS) (de Silva, 2004), which computes embeddings for large distance matrices for which MDS is generally very expensive, by first finding the subspace from the distances of a subset of objects and then estimating the coefficients of the remaining objects in this subspace. This requires a suitable distance measure between untested individuals and evaluated individuals. An overview of the estimator is displayed in Fig. 1.

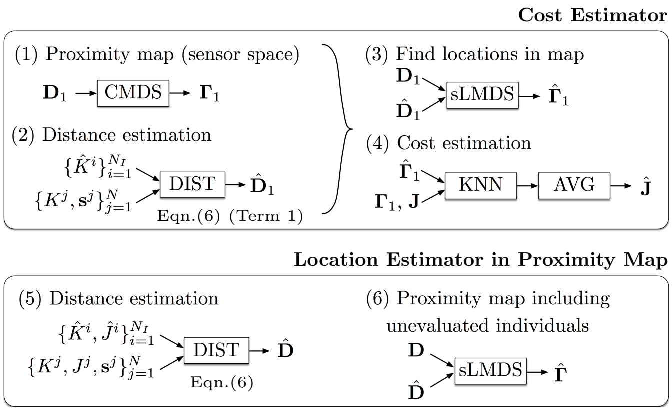

Fig. 1. Schematic of the performance estimator for untested control laws based on their location in the proximity map.

Let the ensemble of newly bred, untested control laws be denoted by $\{\hat{K}^i\}_{i=1}^{N_I}$. The elements of the squared distance matrix between these control laws and already evaluated ones are defined by

$$\hat{D}_{ij}^2 = \langle \hat{K}^i(s), K^j(s)\rangle_{i,j} + \alpha|\hat{J}^i - J^j| \tag{6}$$

where $\hat{\boldsymbol{D}} \in \mathbb{R}^{N_I \times N}$ and $\hat{J}^i$ is the estimated cost of control law $\hat{K}^i$. The first term is given by

$$\langle \hat{K}^i(s), K^j(s)\rangle_{i,j} = \frac{1}{M}\sum_{m=1}^{M} |\hat{K}^i(s^j(t_m)) - K^j(s^j(t_m))|^2 \tag{7}$$

evaluated in the sensor space of the already tested control law. The estimated cost $\hat{J}^i$ in the second term in (6) is determined as the mean of the closest control laws (using the k-nearest neighbor (KNN) search algorithm) in the proximity map given by $\mathbf{\Gamma}_1$. The latter is determined by applying CMDS to the distance matrix $\mathbf{D}_1$ evaluated solely on the first term in (1). Finally, both terms of the distance matrix (6) can be evaluated and the locations of the untested control laws in the proximity map given in the (γ_1, γ_2)−space can be estimated using sLMDS.

2.6 Selection criteria for control laws to be evaluated

In the following, we formulate three criteria based on the proximity of the newly bred individuals to evaluated control laws:

(A) $$\max_i \min_k ||\hat{\boldsymbol{\gamma}}^i - \boldsymbol{\gamma}^k||_2 \tag{8}$$

(B) $$\min_i \#(\boldsymbol{\gamma}^k \text{ s.t. } ||\hat{\boldsymbol{\gamma}}^i - \boldsymbol{\gamma}^k||_2 < R) \tag{9}$$

(C) $$\min_i ||\hat{\boldsymbol{\gamma}}^i - \boldsymbol{\gamma}^{\text{opt}}||_2 \tag{10}$$

where i is the index of the untested and k the index of the already evaluated control law, respectively. Instead of selecting the best control law as defined above, these criteria can also be used as ranking and the best, e.g. 10%, will be evaluated. Both criteria (A) and (B) measure the isolatedness of individuals, i.e. search for control laws that are farthest away from tested control laws or have the fewest number of control laws in a certain radius R around them. These can be interpreted as exploratory criteria. In contrast, criterion (C) exploits regions of well performing control laws by selecting the individual which is closest to the best one $\boldsymbol{\gamma}^{\text{opt}}$ so far.

3. RESULTS

A closed-loop control law optimization using Linear Genetic Programming Control (LGPC) is performed for drag reduction of a car model (Li et al., 2016). The Reynolds number of the model is $Re_H = \frac{U_\infty H}{\nu} = 3 \times 10^5$. The flow is controlled with 4 jet actuators with Coanda surface deflectors at all trailing edges of the model. We consider here particularly single-input multiple-output feedback control, where all four slits are controlled simultaneously. The 16 pressure sensors are distributed on the rear surface. The control objective is a net energy saving from drag reduction which accounts for the actuation expenditure. LGPC is applied for $N_G = 5$ generations, each with $N_I = 50$ individuals (control laws to be tested). A visualization depicting the similarity of these control laws following Sec. 2.1 and 2.2 is shown in Fig. 2. The broad distribution

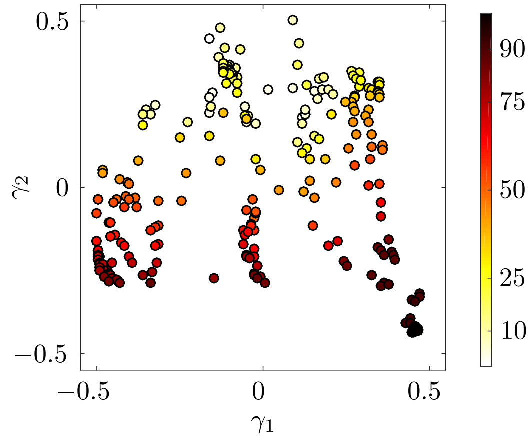

Fig. 2. Control law ensemble color-coded by percentile ranking of their performance. Darker colors refer to better performing individuals, i.e. the nearly black-colored control laws represent the best 10% of all evaluated control laws

of control laws shows that LGPC has successfully explored the control space. The best control law(s) are located in the right bottom corner with $(\gamma_1, \gamma_2) \approx (0.5, -0.5)$. The interpolated fitness topology is displayed in Fig. 3 with the clear optimum in the corner.

For the analysis of the operators, the (γ_1, γ_2)-space is discretized into a coarse grid of 6×6 elements, i.e. $N_f = 36$. The transition probabilities that an offspring lies in any of these tiles under the action of the genetic operators are displayed in Fig. 4. Mutation and cross-over should contribute more to the exploration of the control law space which is confirmed by a rather uniform distribution across the grid. In contrast, elitism and replication shall memorize better performing control laws which results in a more local distribution. While the best control law selected through elitism performs in each generation similarly well, control laws selected through replication, i.e. the same control law is evaluated in the next generation again, exhibit a larger variance in their performance. Since only one transition probability is non-zero for elitism, it can be concluded that the best performing control law has been already found in the second generation (in fact it has been found in the first, but this can not be concluded from this

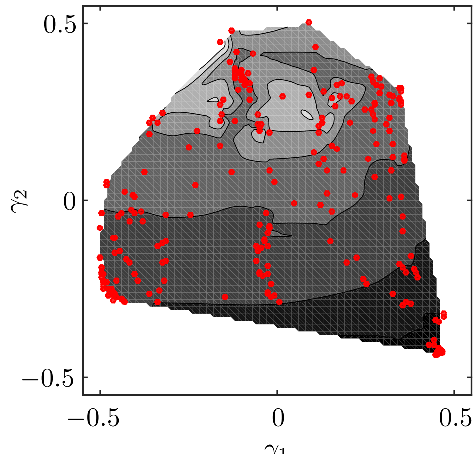

Fig. 3. Control law search space topology obtained from a Delaunay triangulation of the control laws (displayed as small red dots) in (γ_1, γ_2)-space.

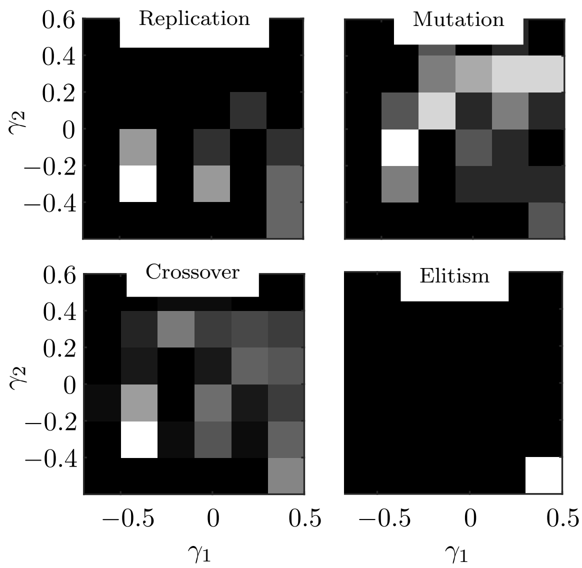

Fig. 4. Probabilities that offspring lands on a specific grid field of the coarse-grained $\boldsymbol{\gamma}-$space due to genetic operations. These probabilities are based on ancestry information of all generations. For a better visualization, probabilities are normalized by the maximum probability, thus high probability is displayed as □ and zero probability as ■.

plot). The analysis of the transition probabilities for each generation (not shown here) suggests that mutation and cross-over do not explore a farther region that is already covered by the first generation.

The transition probabilities that a parent lies in any of these tiles under the action of the genetic operators are displayed in Fig. 5. The single non-zero transition probability for elitism confirms that the best performing control law (or one very close to that) exists indeed already in the first generation. The breeding of a new generation involves a tournament process where better performing individuals are selected with higher probability. As the minimum is located in the lower right corner, this results in a downward shift of the parents from one generation to the next, which is represented by the increasing probabilities

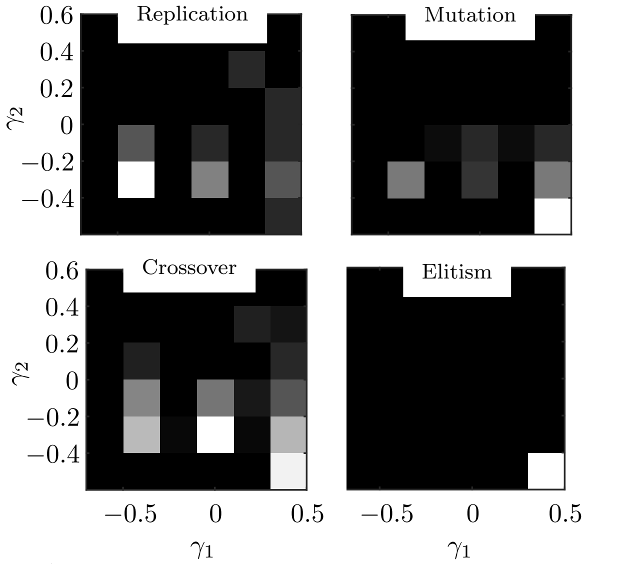

Fig. 5. Analog to Fig. 4 but showing probabilities of origin of parent generation on the coarse-grained $\boldsymbol{\gamma}-$space due to genetic operations.

for decreasing γ_2. This global downward shift indicates the convergence to the top-performing individuals.

In the following, the performance estimator of the untested individuals in the proximity map and the selection criteria are assessed. The approach is applied to the 5th generation of control laws that is assumed to be not evaluated yet. The proximity map computed on the first four generations with the estimated and true locations of the 5th generation is displayed in Fig. 6. Intriguingly, the estimated locations

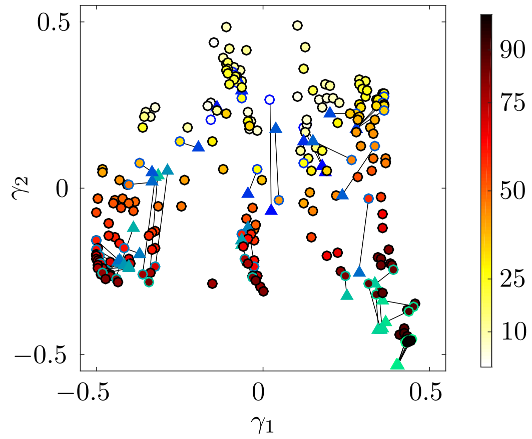

Fig. 6. Comparison of true and estimated locations of individuals in the 5th generation based on information from the four previous generations. True $\boldsymbol{\gamma}^i$'s are displayed as in Fig. 2. Estimated $\hat{\boldsymbol{\gamma}}^i$'s are depicted by color-coded triangles (blue to green). The color of the triangles matches the border color of the circle associated with the true location (also connected by straight line).

are in close proximity to the true ones, considering that the control law distances are only evaluated on the sensor space of the previous generations and the cost can only be estimated from these distances. The mean and median of the relative displacement error $||\hat{\gamma} - \gamma||_2/||\gamma||_2$ are 0.37 and 0.18, respectively.

The selection criteria are employed to rank the 5th generation of individuals in the proximity map in Fig. 6. Then, the best N_s individuals are selected (as for evaluation). The overlap ratio as a function of the N_s best individuals selected based on the true or estimated locations, respectively, is displayed in Fig. 3. All criteria achieve a relatively

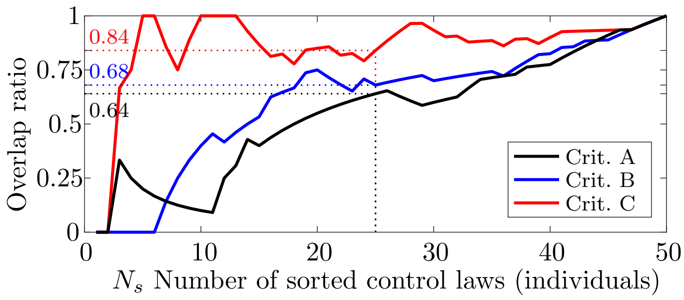

Fig. 7. Assessment of selection criteria w.r.t. the true performance of the control laws of the 5th generation. Overlap of the P best individuals that are selected according to the criteria (A), (B), or (C) applied to the estimated and true locations, respectively.

high overlap of individuals when compared to the true locations. Criterion (C) performs better than the other criteria as it is more robust. The reason is that the location estimation performs sufficiently well, i.e. individuals, which perform well, are estimated to be in the proximity of well-performing control laws and vice-versa. In contrast criteria based on isolatedness are more sensitive to the error of the positioning in the proximity map.

In summary, the analysis shows that LGPC explores successfully the control space, the convergence to top-performing individuals, and the exploratory and exploitative nature of the genetic operators. Moreover, the location estimation of untested control laws in the proximity map performs surprisingly well. This is a critical enabler for a faster GPC, that can be systematically geared towards exploration or exploitation of the best control laws, by selectively evaluating or discarding newly bred individuals. Ongoing work investigates suitable selection criteria and incorporating this approach to explore the control space beyond the region represented by the first generation and to avoid evaluation of similarly performing control laws.

ACKNOWLEDGEMENTS

EK gratefully acknowledges funding by the Gordon and Betty Moore, Alfred P. Sloan, and the Washington Research Foundations and the eScience Institute. The thesis of RL is supported by the OpenLab Fluidics between PSA Peugeot-Citroën and Institute Pprime (Fluidics@poitiers). This work is also supported by a public grant overseen by the French National Research Agency (ANR) as part of the "Investissement dAvenir" program, through the "iCODE Institute project" funded by the IDEX Paris-Saclay, ANR-11-IDEX-0003-02. We appreciate valuable stimulating discussions with Steven Brunton, Francois Lusseyran, Lionel Mathelin and Luc Pastur.

REFERENCES

Brunton, S.L. and Noack, B.R. (2015). Closed-loop turbulence control: Progress and challenges. *Appl. Mech. Rev.*, 67(5), 050801:01–48.

Cox, T.F. and Cox, M.A.A. (2000). *Multidimensional Scaling*. 2nd edition.

de Silva, V. & Tenenbaum, J.B. (2004). Sparse multidimensional scaling using landmark points. Technical report, Standford University.

Duriez, T., Brunton, S., and Noack, B.R. (2016). *Machine Learning Control — Taming Nonlinear Dynamics and Turbulence*.

Emmerich, M., Giotis, A., Özdemir, M., Bäck, T., Giannakoglou, Kyriakos", e.J.J.M., Adamidis, P., Beyer, H.G., Schwefel, H.P., and Fernández-Villacañas, J.L. (2002). *Metamodel—Assisted Evolution Strategies*. Springer.

Gautier, N., Aider, J.L., Duriez, T., Noack, B.R., Segond, M., and Abel, M.W. (2015). Closed-loop separation control using machine learning. *J. Fluid Mech.*, 770, 242–441.

Giacobini, M., Tomassini, M., and Vanneschi, L. (2002). Limiting the number fitness cases in genetic programming using statistics. In J.J.M. Guervós, P. Adamidis, H.G. Beyer, J.L.F.V. nas, and H.P. Schwefel (eds.), *Parallel Problem Solving from Nature - PPSN VII*, 2439, 371–380. Springer.

Jin, Y., Olhofer, M., and Sendhoff, B. (2002). A framework for evolutionary optimization with approximate fitness functions. *IEEE Trans. Evolut. Comput.*, 6, 481–494.

Kim, H.s. (2001). An efficient genetic algorithm with less fitness evaluation by clustering. In *Proc. of the 2001 IEEE Congress on Evolutionary Computation*, 887–894.

Koza, J.R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection.*

Li, R., Noack, B.R., Cordier, L., Borée, J., Harambat, F., Kaiser, E., and Duriez, T. (2016). Drag reduction of a car model by linear genetic programming control. *arXiv:1609.02505v1.*

Mardia, K.V., Kent, J.T., and Bibby, J.M. (1979). *Multivariate Analysis*. Academic Press.

Parezanović, V., Cordier, L., Spohn, A., Duriez, T., Noack, B.R., Bonnet, J.P., Segond, M., Abel, M., and Brunton, S.L. (2016). Frequency selection by feedback control in a turbulent shear flow. *J. Fluid Mech.*, 797, 247–283.

Sastry, K., Goldberg, D.E., and Pelikan, M. (2001). Don't evaluate, inherit. In *Proc. of the Genetic and Evolutionary Computation Conference (GECCO01)*, 551–558. Morgan Kaufmann.

Schoenberg, I.J. (1935). Remarks to Maurice Fréchet's article "Sur la définition axiomatique d'une classe d'espaces distanciés vectoriellement applicable sur l'espace de Hilbert". *Annals of Mathematics*, 38, 724–732.

Young, G. and Householder, A.S. (1938). Discussion of a set of points in terms of their mutual distances. *Psychometrika*, 3, 19–22.

Ziegler, J. and Banzhaf, W. (2003). Decreasing the number of evaluations in evolutionary algorithms by using a meta-model of the fitness function. In *Proc. of the 6th European Conference on Genetic Programming.*