A Surrogate-Based Strategy for Multi-Objective Tolerance Analysis in Electrical Machine Design

Alexandru-Ciprian Zăvoianu*‡, Edwin Lughofer*, Gerd Bramerdorfer†‡, Wolfgang Amrhein†‡, and Susanne Saminger-Platz*

*Department of Knowledge-based Mathematical Systems / Fuzzy Logic Laboratory Linz-Hagenberg, Johannes Kepler University Linz, Austria
†Institute for Electrical Drives and Power Electronics, Johannes Kepler University Linz, Austria
‡LCM, Linz Center of Mechatronics, Linz, Austria

***Abstract*—By employing state-of-the-art automated design and optimization techniques from the field of evolutionary computation, engineers are able to discover electrical machine designs that are highly competitive with respect to several objectives like efficiency, material costs, torque ripple and others. Apart from being Pareto-optimal, a good electrical machine design must also be quite robust, i.e., it must not be sensitive with regard to its design parameters as this would severely increase manufacturing costs or make the physical machine exhibit characteristics that are very different from those of its computer simulation model. Even when using a modern parallel/distributed computing environment, carrying out a (global) tolerance analysis of an electrical machine design is extremely challenging because of the number of evaluations that must be performed and because each evaluation requires very time-intensive non-linear finite element (FE) simulations. In the present research, we describe how global surrogate models (ensembles of fast-to-train artificial neural networks) that are created in order to speed-up the multi-objective evolutionary search can be easily reused to perform a fast tolerance analysis of the optimized designs. Using two industrial optimization scenarios, we show that the surrogate-based approach can offer very valuable insights regarding the local and global sensitivities of the considered objectives at a fraction of the computational cost required by a FE-based strategy. Encouraged by the good performance on individual designs, we also used the surrogate approach to track the average sensitivity of the Pareto front during the entire optimization procedure. Our results indicate that there is no generalized increase of sensitivity during the runs, i.e., the used evolutionary algorithms do not enter a stage where they discover electrical drive designs that trade robustness for quality.**

***Keywords*—tolerance/sensitivity analysis, global surrogate models, artificial neural networks, evolutionary design optimization.**

I. Introduction and State of the Art

When designing various industrial products, carrying out real-life performance tests is highly impractical due to financial and time-related reasons. Fortunately, the advent of many computer-assisted design (CAD) systems and accurate physics-based simulation software has enabled engineers to perform required (design) experiments much faster and much cheaper in a virtual environment.

Even though simulating the behavior of a product is much faster than constructing physical prototypes and setting up real-life experiments, high-level tasks like design space exploration, optimization and tolerance/sensitivity analysis remain rather impractical for many real-life problems. This is because such high-level tasks require the virtual evaluation (and hence, computer-based simulation) of thousands of different models. Since a single computer-based evaluation is a very computationally-intensive operation that may require anything from several seconds to a few hours, the total computational requirements of performing an optimization or a tolerance analysis remain daunting (even when using parallel/distributed computing).

In order to try and alleviate this burden, several researchers [1]–[4] have proposed the usage of surrogate modeling – i.e., the creation of fast-to-evaluate linear and non-linear regression models that can accurately approximate the results produced by computer/physics-based simulation software.

Concerning the specific task of evaluating the sensitivity (changes in objectives and constraints) of a fixed design (parametric model) to allowed deviations in some of its parameter values, in [5], Stephens et. al present the comparative results of applying three types of surrogate modeling techniques: radial-basis functions [6], artificial neural networks [7] and support vector machines [8]. The surrogates from [5] were aimed at replacing high-fidelity computational fluid dynamics models. The results indicate that, despite proposing a different model architecture, all three surrogate construction methods produce largely identical tolerance/sensitivity results. Other noteworthy examples of using surrogates for performing tolerance analysis can be found in [9] and [10].

There are several ways in which a tolerance/sensitivity analysis can be defined/performed [11]. In this paper we describe a variation based on one of the most common strategies: the **min-max approach**. Given a design encoded as a variable parameter vector $\mathbf{x} = \{x_1, x_2, \ldots, x_n\}$ and m targets (i.e., objectives, constraints) that we wish to analyze $\mathbf{y}(\mathbf{x}) = \{y_1(\mathbf{x}), y_2(\mathbf{x}), \ldots, y_m(\mathbf{x})\}$, the outline of our idea is to:

1) ask the decision maker (i.e, engineer, designer, final client) to supply the allowed deviation/tolerance intervals for each of the n variables. They can be defined using a relative formulation (e.g., $\pm p_c\%$ of the nominal value of variable x_c), or a nominal-value formulation (e.g., a lower-bound and upper-bound interval $[l_c, u_c]$ with $l_c \leq 0$, $u_c \geq 0$ and $P(l_c = 0 \cap u_c = 0) = 0$ that will be centered around x_c);

2) for $i \in 1,\ldots,m$, find (as fast as possible) an accurate estimation for y_i^{min} and y_i^{max} , the nominal min and max values of target y_i that one can expect when the n variables can take any value within their allowed deviation intervals;
3) compute and present the decision maker with several informative metrics aimed at displaying the magnitude of the expected target changes.

We are particularly focused on industrial problems related to the optimal design of electrical drives. A generic overview of such problems and of the need to perform a tolerance analysis of their solutions is presented in the next section. Section III contains the detailed description of our general tolerance analysis approach (i.e., of steps 2 and 3 of the preceding enumeration). Section IV describes how the particular optimal motor designs analyzed in this work have been obtained. Section V presents the results of two different tolerance analyses and offers some interpretation hints. The last part of this work contains the conclusions and some perspectives on future work.

II. Surrogate-Assisted Optimization of Electrical Machine Designs

By optimizing the design of electrical machines, decision makers aim to find assemblies that are highly competitive with respect to several, usually conflicting, objectives like efficiency, material costs, torque ripple, and others. Hence, decision makers in the field are actually faced with multi-objective optimization problems (MOOPs). The result of such problems is a *set* of Pareto-optimal (or non-dominated) solutions (PN). Each member of the PN simultaneously holds two properties:

1) it is not worst across all optimization objectives when compared to any other solution from the PN;
2) it is better than any other solution from the PN with regard to at least one of the objectives.

The projection of the PN in objective space is called the Pareto front (PF).

Population-based stochastic optimization techniques in general, and evolutionary algorithms in particular, have proven to be one of the best performing methods for solving MOOPs [12]. One of the characteristics of multi-objective evolutionary algorithms (MOEAs) is that a large number of (fitness) evaluations of the objective function (i.e., target quality estimations) need to be performed during their execution. This is a major problem since the evaluation of the targets of a single electrical motor design (i.e., of an individual from the population) is done via a series of computationally-intensive non-linear finite element (FE) simulations. The FE simulations are usually preceded by a modeling stage in which a CAD software is used to construct a 2D or 3D model/mesh of the assembly from the given design parameter vector. A sketch of this simulation based evaluation strategy is presented in the left hand part of Figure 1.

MagOpt [13] – a state-of-the-art integrated framework for optimizing mechatronic components – features several MOEAs (e.g., NSGA-II [14] and SPEA2 [15]) that can be employed when aiming to optimize electrical drive designs. In order to reduce the prohibitive run times obtained when running these algorithms, we have initially proposed [16] a hybrid optimization strategy that relies on creating global surrogate models on the fly. The best performing surrogates were based on multi-layer perceptrons [7] (MLPs) – a specific type of artificial neural networks. The focal idea was to obtain speed improvements by switching the MOEA to a surrogate-based fitness assessment function (that completely removed CAD and FE dependencies) during the middle part of the optimization runs. The hybridization strategy was successful and in a subsequent work [17] we showed that further reductions in total optimization time could be obtained by switching to ensemble-based surrogate models trained over Pareto-trimmed training sets. For each target, the ensemble surrogates are simply averaging the results of the 10 best individual MLP models obtained after a predefined limited best parameter grid search. The new fast-to-train non-linear ensemble predictors also enabled the integration of multiple *surrogate-based optimization blocks* inside an optimization run as suggested by the *"Repeat option"* line from the right hand part of Figure 1. This was done with the hope of exploring a bigger part of the search domain. An overview of the resulting surrogate-assisted optimization procedure is presented in Figure 1.

The main aim of the present work is to show how the global ensemble surrogate models/predictors can be also used to rapidly perform a tolerance analysis of the optimized electrical drive designs. In general, this is a secondary but very important stage of the overall design process since it is usually demanded that a good assembly must be robust with regard to (at least some of) its parameters. Having an oversensitive motor design would severely increase manufacturing costs or make the resulting physical machine (prototype) exhibit characteristics that are very different than those of its computer simulation model. All in all, if a design that is in the Pareto front is not deemed robust, the design will likely not be preferred by the decision maker.

III. The Proposed Tolerance Analysis Strategy

Regardless of the method used to evaluate a given motor design (FE or surrogates) and finally compute the **min**imum and **max**imum expected target values, we apply the same tolerance analysis strategy. The reason is that, in the present stage of development, after performing the surrogate analysis, we want to double-check it using FE computations. Our aim is to obtain an accurate overview of the estimation imprecision one should expect from opting for the surrogate-based approach. On the long run, if the surrogate-induced imprecision is not causing interpretation issues, the surrogates can be used to (at least) offer valuable insight related to target sensitivities.

In light of the the limited number of evaluations that are generally allowed when performing a (time-constrained) FE-based tolerance analysis, the general analysis we propose is a sequence of two distinct stages:

- a primary estimation of sensitivity with regard to **local changes** that affect a single design variable, i.e., the standard one-factor-at-a-time (OFAT) [18] approach;
- a secondary estimation of sensitivity with regard to **global interactions** between all the considered design variables.

The reasoning is that, although optimistic in nature, after performing a limited number of evaluations, the primary

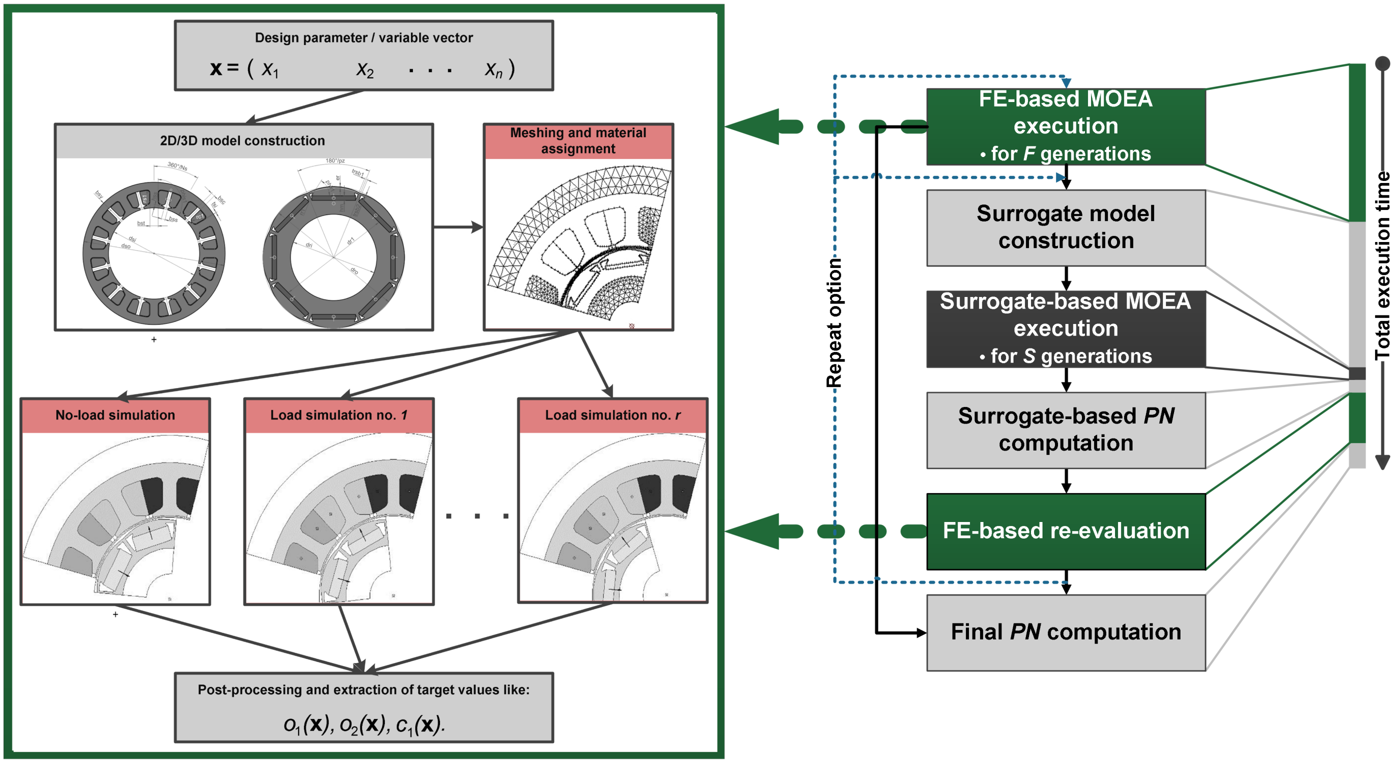

Fig. 1. An overview of the surrogate-based optimization process implemented in the MagOpt [13] framework. There are three hypothetical targets that must be considered: two optimization objectives (o_1 and o_2) and one constraint (c_1).

(local/OFAT) sensitivity estimation will provide the decision maker with a quite accurate picture of the individual influence exhibited on the elicited targets by the various design variables. The more computationally-intensive – and more imprecise/stochastic – secondary (global) sensitivity estimation is aimed at discovering if the possible interactions between the influences of different design variables are likely to amplify the local sensitivities.

A. *Estimation of Local Sensitivity*

Given a design variable vector $\mathbf{x}$ of size n, the local sensitivity of target value $y_i(\mathbf{x})$ to the sole variation of the design variable value x_c, $c \in \{1,\ldots,n\}$ inside the interval $[x_c+l_c, x_c+u_c]$, where $l_c \leq 0$, $u_c \geq 0$ and $P(l_c = 0 \cap u_c = 0) = 0$, is estimated by:

- evaluating 11 different design variations, $\mathbf{x^0}, \mathbf{x^1}, \ldots, \mathbf{x^{10}}$ obtained using the formula:

$$x_j^k = \begin{cases} x_j + l_c + 0.1 \cdot (u_c - l_c) \cdot k, & \text{if } j = c \\ x_j, & \text{if } j \neq c \end{cases}, \tag{1}$$

  with $k \in \{0, 1, \ldots, 10\}$ and $1 \leq j \leq n$;
- selecting the two extreme target values y_i^{min} and y_i^{max} from the aforementioned 11 design variations
- computing the *relative (elasticity) indicator*:

$$\Delta_{loc}(y_i, x_c, l_c, u_c) = \frac{x_c \cdot (y_i^{max} - y_i^{min})}{y_i \cdot (u_c - l_c)} \tag{2}$$

Under the assumption of local linearity inside (narrow) variation intervals, Eq. (2) should provide a good estimation of the relative influence that variable x_c has on target y_i. The interpretation of Δ_{loc} is very simple: for a change of $z\%$ in x_c[1], one should expect a change of $z \cdot \Delta_{loc}(y_i, x_c, l_c, u_c)\%$ in y_i. Thus, a major advantage of Δ_{loc} is that its relative formulation enables a direct comparison between the local sensitivities induced by a variable across several targets or the local sensitivities induced by all variables on a single target.

According to our described OFAT approach, the total number of designs that must be analyzed during the primary estimation of sensitivity for a single target is $11 \cdot n$ where n is the number of considered design variables.

B. *Estimation of Global Sensitivity*

The global sensitivity of target value $y_i(\mathbf{x})$ to the simultaneous variation of all the design variables of $\mathbf{x}$ inside their predefined deviation intervals is estimated by:

- evaluating 1000 "neighboring" designs obtained via a Latin Hypercube sampling [19] (LHS) method and, in some cases, a grid combination of extreme design variable values;
- selecting the two extreme nominal target values y_i^{min} and y_i^{max} among the $1000 + 11 \cdot n$ (global+local) design variations;

[1] Such that $x_c \cdot (1 \pm z) \in [x_c + l_c, x_c + u_c]$ holds.

- computing

$$\Delta_{glob}(y_i) = \frac{x_{spr} \cdot (y_i^{max} - y_i^{min})}{y_i \cdot [x_{spr}^{max} - x_{spr}^{min}]} \tag{3}$$

  where x_{spr} is the base value of the design variable that displays the largest relative deviation when looking at the variable vectors associated with y_i^{min} and y_i^{max}. x_{spr}^{max} and x_{spr}^{min} denote the minimal and maximal values of x_{spr} inside the variable vectors associated with y_i^{min} and y_i^{max}.

In order to understand how x_{spr}, x_{spr}^{max} and x_{spr}^{min} are obtained, let us consider a simple example with:

- a base vector $\mathbf{x} = (0.2, 0.5, 4.8)$ for which we wish to analyze the global sensitivity with regard to target y_i;
- a variation vector $\mathbf{x^{min}} = (0.19, 0.505, 4.776)$ which corresponds to y_i^{min};
- a variation vector $\mathbf{x^{max}} = \mathbf{x} = (0.21, 0.49, 4.824)$ which corresponds to y_i^{max};

The observed relative variable deviations are $|0.21 - 0.19|/0.2 = 0.1$ for x_1, $|0.49 - 0.505|/0.5 = 0.03$ for x_2 and $|4.824 - 4.776|/4.8 = 0.01$ for x_3. Since the largest deviation is observed for variable x_1, $x_{spr} = x_1 = 0.2$, $x_{spr}^{max} = 0.21$ and $x_{spr}^{min} = 0.19$

The 1000 individuals created for the global sensitivity estimation stage are generated according to different approaches depending on the total number of design variables.

If $1 < n \leq 9$:

- 2^n individuals will be created by performing a grid combination of extreme values. The idea is that, when assuming local linearity, extreme target values correspond to extreme variable values (i.e., $x_c + l_c$ or $x_c + u_c$) and the strongest interactions are also expected between combinations of extreme variable values. As an example, a design variation created in this stage can have the form: $\mathbf{x}^* = (x_1 + l_1, x_2 + l_2, \ldots, x_n + u_n)$. This part of the global sensitivity estimation is *stable* as the same variations are generated whenever fixing the allowed variable deviation intervals.
- $1000 - 2^n$ individuals will be created using the LHS. Obviously, this part of the global sensitivity estimation is *stochastic*.

If $n > 9$, all the 1000 individuals will be created only via LHS and the resulting global sensitivity estimation is expected to be slightly more unstable.

Because of the reliance on statistical sampling and a relatively small pool of neighbors, it should be noted that there is a higher intrinsic uncertainty associated with Δ_{glob} (as opposed to Δ_{loc}) even when using only FE-based target estimates. Nevertheless, the global approach is far more realistic as (especially in production environments) one is far more likely to encounter simultaneous design parameter deviations. Since all the targets of a given design are analyzed based on the same $1000 + 11 \cdot n$ variations, one can use the resulting Δ_{glob} values to compare between the global sensitivities of different targets: Δ_{glob} values of different magnitudes help to rapidly discriminate between robust and highly sensitive targets.

The extreme target values (y_i^{min} and y_i^{max}) discovered during the global sensitivity estimation are also very useful to the decision maker. These values are highly practical as they indicate the expected nominal target values in the worst-case scenario. Using them, one can compute the expected relative loss or gain for each target. By design[2], our optimization problems always require target minimization and thus, the relative loss associated with target y_i is given by:

$$y_i^{loss} = \frac{(y_i^{max} - y_i)}{|y_i|} \tag{4}$$

IV. Experimental Setup

A. The Design Problems and the Elicited Targets

As previously mentioned, the general aim of the present work lies in trying to improve the design process of electrical machines by providing decision makers with the means of performing very fast tolerance analyses. We focus especially on motor designs that are the (Pareto-optimal/non-dominated) results of multi-objective optimization runs of two different industrial problems:

Problem no. 1 concerns a motor with an interior rotor topology with embedded magnets. The stator and rotor cross-sections are presented in Figure 2. The corresponding design parameter vector has a size of 6 and is given by:

$$\mathbf{x} = (\ b_{st},\ \ d_{si},\ \ e_r,\ \ b_{ss},\ \ \alpha_m,\ \ h_m\ )\ , \tag{5}$$

where all parameters are illustrated in Figure 2 except for α_m, which denotes the ratio between the actual magnet size and the maximum possible magnet size as a result of all other geometric parameters of the rotor. The goal is to simultaneously optimize four unconstrained objectives: y_1 – the overall cost of the materials necessary for building the motor, y_2 – the peak-to-peak value of the motor torque for no current excitation, y_3 – the efficiency of the electrical drive and y_4 – the equivalent of y_2 at load operation. All four objectives are considered during the tolerance analysis stage.

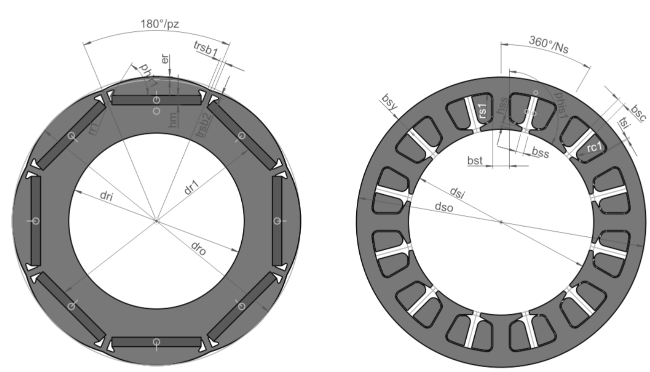

Fig. 2. Cross-sections with the geometric dimensions of the stator and the rotor for Problem no. 1

[2]When y_i conceptually needs to be maximized (because, for instance, it represents the efficiency of the electrical drive), we opt to minimize $-y_i$.

Problem no. 2 is also formulated for a motor that features an interior rotor topology. The associated stator and rotor cross-sections are displayed in Figure 3. The design parameter vector contains 22 real-valued variables that must be configured in order to optimize four unconstrained objectives that regard efficiency (y_2 and y_4) and production costs (y_3 and y_5). The first target of the tolerance analysis (y_1) is a constraint imposed on peak-to-peak value of the motor torque at load operation. The performance of the assembly is evaluated at 3000 revolutions per minute.

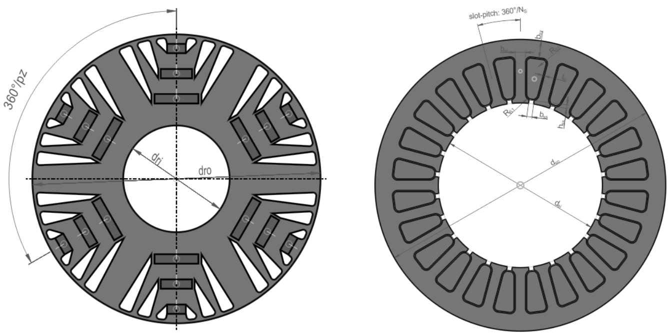

Fig. 3. Stator and the rotor cross-sections for Problem no. 2

B. The Multi-Objective Optimization Runs

The first electrical drive design optimization problem we consider is rather standard and we were able to achieve good optimization results when applying a classic MOEA like SPEA2. We used standard genetic operators: SBX crossover [20] and polynomial mutation [21]. We opted for an equal archive and population size of 100 as this setting has yielded very good results on this type of problem in multiple runs over the past four years. We also used a literature-recommended parameterization with a crossover probability of 0.9, a crossover distribution index of 20, a mutation probability of $1/|\mathbf{x}|$ and a mutation distribution index of 20.

For the second, more challenging problem, we used DECMO [22] – a hybrid multi-objective evolutionary algorithm based on cooperative coevolution that aims to profit from the very good performance exhibited by differential evolution (DE) operators on continuous real-valued search spaces. DECMO was parameterized as described in [23].

Both optimization runs were allowed to perform 10000 FE-based fitness evaluations.

The way in which surrogate modeling was used to speed-up the optimizations is straightforward. In each of the two runs we inserted five *surrogate-based optimization blocks*: after 1000, 2000, 3000, 4000 and 5000 FE-based fitness evaluations. We constructed ensemble surrogate models that employed 10 MLPs trained over *Pareto-trimmed training sets* of at most 1000 individuals. At the end of each surrogate-based optimization cycle we re-evaluated the best 600 (i.e. 6 or 3 generations worth of) individuals discovered during the surrogate-based search. The surrogate-based optimizations were carried out using the base algorithm and parameterization of the run (i.e., SPEA2 or DECMO) and were allowed to generate 5000 individuals. We stopped using surrogate-based optimization blocks after 5600 FE-based evaluations and re-evaluations because the improvements delivered by the enhancement were becoming smaller and smaller. As such, for the last 4400 evaluations, both runs fell back to fully FE-based MOEAs. When also considering the surrogate-based designs, each final PN was obtained after considering 32000 possible designs.

The final MLP-based ensemble surrogate models (used for performing the tolerance analyses described in the next section) were trained – as described in [17] – after the end of the optimizations using all the valid FE-evaluated individuals generated during the runs.

V. Results and Interpretation

A. Design-wise Local Sensitivities

For each design problem, after performing the multi-objective optimization, we proceeded to apply the local and global sensitivity estimation strategies described in Section III on several designs selected by the decision maker from the obtained PNs.

We first performed very fast[3] surrogate-based analyses and then we re-evaluated them (i.e., re-analyzed the created design variations) using lengthy[4] FE simulations. For each design we always considered relative allowed deviations of $\pm 1\%$ for every variable. The main aim of performing dual surrogate/FE analyses was to assess the accuracy of the surrogate-based sensitivity results.

In Figures 4 and 5 we plotted the surrogate-based and FE-based local sensitivity maps (matrices) of two particular optimization solutions – **EvoDesign1** for the first problem and **EvoDesign2** for the second one. These two solutions are the ones deemed by the decision maker as the most interesting with regard to each optimization problem. Each cell in the local sensitivity matrices is obtained by computing Eq. (2) for the considered variable/target pair.

The matrices from Fig. 4, that correspond to the optimized design evolved for Problem No. 1, show that:

- the maximum Δ_{loc} value estimated using the global surrogate models is slightly optimistic (≈ 5) compared to the one computed using FE simulations (≈ 7);
- the overall structure of the two local sensitivity maps is very similar as the surrogate-based analysis is able to correctly identify the more sensitive targets (y_2 and y_4) and the variables that determine the local sensitivities (x_2, x_5 and, to a lesser extent, x_4).

The matrices from Fig. 5, that corresponding to the optimized design evolved for the second (harder-to-model) problem, show that:

- the broad structure of the two local sensitivity maps is very similar as: y_1 is identified as generally sensitive target (that is affected by the local variations of several variables) and y_3 and y_5 are quite sensitive only to x_3 and, to a lesser extent, x_1;

[3] $\approx$ 7 seconds of local computation

[4] $\approx$ 6 hours when distributing over 50 HTCondor™ [24] nodes

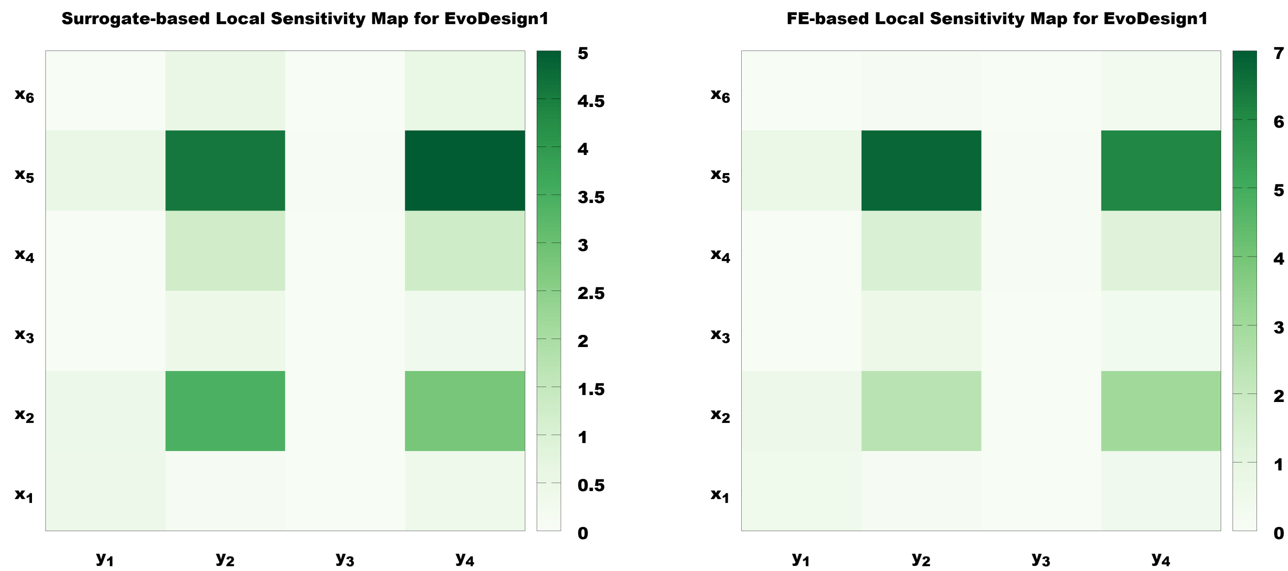

Fig. 4. Surrogate-based and FE-based local sensitivity maps for EvoDesign1 – a Pareto-optimal solution of Problem no. 1

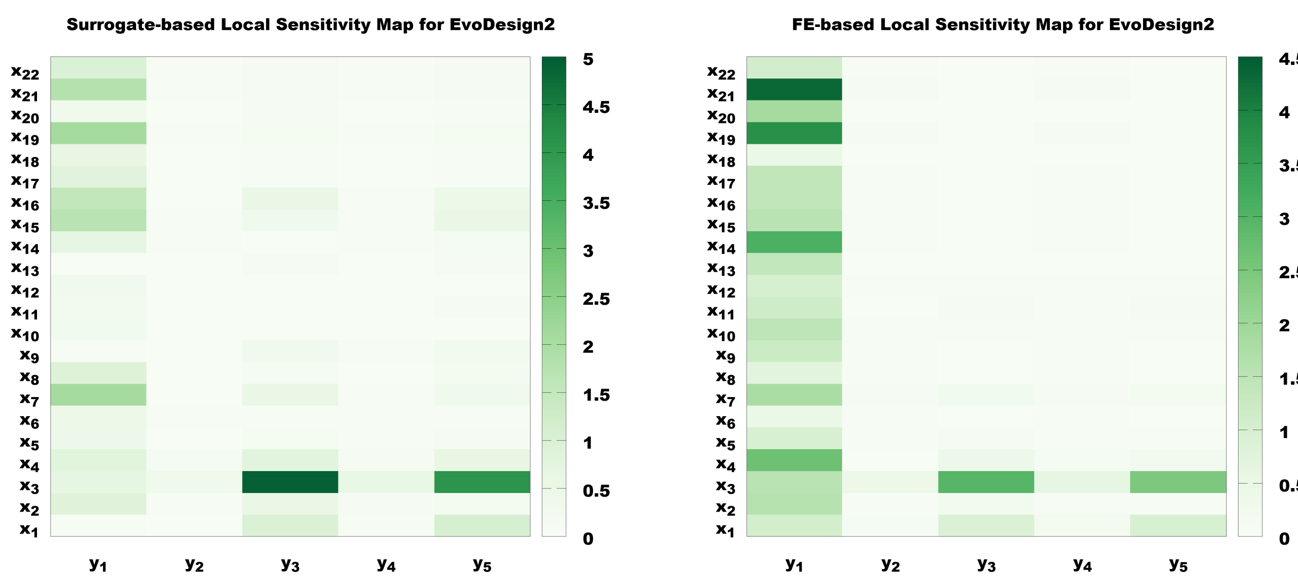

Fig. 5. Surrogate-based and FE-based local sensitivity maps for EvoDesign2 – a Pareto-optimal solution of Problem no. 2

- although maximum Δ_{loc} values are quite similar, on the finer level, the surrogate-based analysis is underestimating the local sensitivities associated with y_1 and slightly overestimating some local sensitivities associated with y_3 and y_5;
- for the targets where a very good surrogate model can be constructed (i.e., y_2 and y_4), the surrogate-based analysis is able to correctly identify a rather slight sensitivity with regard to x_3.

A general conclusion for all the comparative surrogate/FE local sensitivity estimations we have performed is that the surrogate-based variant is able to correctly identify broad sensitivity patterns (and thus offer valuable insights to the decision maker) but (as expected) the precision of the sensitivity estimation is heavily influenced by the quality of the global surrogate model on which the analysis is based.

B. Design-wise Global Sensitivities

Information concerning the global sensitivities of two analyzed electrical drive designs is centralized in Tables I and II. The first row of each table also contains the training R^2 achieved by each ensemble surrogate model during 10-fold cross-validation.

TABLE I. GLOBAL SENSITIVITY INDICATORS FOR EVODESIGN1

Indicator	Tolerance analysis targets			
	y_1	y_2	y_3	y_4
Surrogate R^2	0.9941	0.9664	0.9938	0.9822
Base value	6.9719	0.1062	-0.9169	0.1104
Surrogate-based y_i^{min}	6.8968	0.0989	-0.9171	0.1002
Surrogate-based y_i^{max}	7.0588	0.1140	-0.9165	0.1189
Surrogate-based y_i^{loss}	1.25%	7.34%	0.00%	7.70%
Surrogate-based Δ_{glob}	1.1863	8.4553	0.0306	8.5650
FE-based y_i^{min}	6.8776	0.0998	-0.9171	0.1034
FE-based y_i^{max}	7.0824	0.1207	-0.9142	0.1254
FE-based y_i^{loss}	1.58%	13.65%	0.03%	13.59%
FE-based Δ_{glob}	1.4688	9.8396	0.1617	9.9678

TABLE II. GLOBAL SENSITIVITY INDICATORS FOR EVODESIGN2

Indicator	Tolerance analysis targets				
	y_1	y_2	y_3	y_4	y_5
Surrogate R^2	0.7425	0.7190	0.9651	0.7532	0.9746
Base value	11.1451	-88.4460	37.2550	-84.7884	31.3534
Surrogate-based y_i^{min}	10.3235	-88.8901	34.7290	-85.4092	29.5796
Surrogate-based y_i^{max}	12.2576	-88.0092	39.7738	-84.1261	33.1594
Surrogate-based y_i^{loss}	9.98%	0.49%	6.76%	0.78%	5.76%
Surrogate-based Δ_{glob}	9.0530	0.4925	6.7436	0.7460	5.7603
FE-based y_i^{min}	12.7686	-89.4934	35.8252	-86.2212	30.3409
FE-based y_i^{max}	15.3052	-88.6053	38.6869	-84.9568	32.3845
FE-based y_i^{loss}	37.33%	−0.18%	3.85%	−0.20%	3.33%
FE-based Δ_{glob}	11.3801	0.5021	3.8408	0.7456	3.2586

In the case of EvoDesign1, the magnitudes of both the surrogate and the FE-based indicators show that targets y_2 and y_4 are highly sensitive and target y_1 is slightly sensitive. Target y_3 seems quite insensitive even when considering global interactions between allowed variable deviations. It should be noted that while surrogate and FE-based Δ_{glob} values are rather similar, the FE-based analysis estimates losses for y_2 and y_4 that are nearly twice as large as those predicted using surrogate models.

In the case of EvoDesign2, the magnitudes of the surrogate-based indicators show that targets y_1, y_3 and y_5 are highly sensitive and targets y_2 and y_4 are slightly sensitive. The FE-based indicators confirm these estimations but indicate that y_1 might be more sensitive than initially assessed while y_3 and y_5 seem less sensitive than the surrogate-based analysis indicated. Surprisingly, on y_2 and y_4, two of the three targets of EvoDesign2 where the quality of the surrogate models is quite low, the surrogate and FE-based estimations are far more similar.

The overall conclusion for the global sensitivity tests we performed is that the surrogate-based approach is able to produce meaningful estimations. For example, these fast, surrogate-based, sensitivity results can be used to at least make inferences about the targets of a design at a broad (fuzzy) level: highly sensitive, sensitive, slightly sensitive, etc.

C. Pareto-wise Sensitivity Evolution

Since all our FE-based analyses showed that the surrogate-based global sensitivity indicators are able to at least correctly assess sensitivity magnitudes[5], we attempted to track the **Pareto-wise sensitivity evolution** of different objectives. The reason for this is that we wanted to check if the MOEA-based optimization is inherently prone to discover models that are not only better than their predecessors but also more sensitive. In other words we wanted to check if, at a certain stage during the optimization runs, we would start trading design sensitivity for design optimality.

Our approach is quite simple:

- during the run of the MOEA, at the end of each generation, we stored all the Pareto non-dominated individuals in the population (i.e., all the individuals that define the up-to-date PF);
- at the end of the multi-objective optimization process, we constructed surrogate models for all the elicited targets using as training samples all the FE-evaluated individuals;
- for every target, we analyzed the global sensitivity of every Pareto non-dominated individual in storage and averaged the results over all the designs stored at a given time. We considered $\pm1\%$ deviations for every design variable.

The results we obtained indicate that, for the performed optimization runs there is no systematic increase of the (global) sensitivities associated with PN designs. Nevertheless, some interesting observations can be made regarding the evolution of the average Pareto-wise relative loss indicator:

- for the targets where there is a small average improvement in quality, the average Pareto-wise loss displays a stable behavior. This is exemplified in Figure 6 for target y_3 of Problem no. 1 where the average Pareto-wise quality improves by a factor of only 1.0042 during the entire run (from -89.74 to -90.12).
- for the targets where there is a significant average improvement in quality, the average Pareto-wise loss is showing a steady increase at certain times. This is exemplified in Figure 7 for target y_5 of Problem no. 2 where the average Pareto-wise quality improves by a factor of 4.0405 during the entire run (from 163.19 to 40.39). The average nominal loss (i.e., $y_i^{max} - y_i$) also improves by a factor of 1.3731 during the optimization run from 3.57 in the first Pareto-optimal set to 2.60 in the Pareto-optimal set obtained after 10000 FE-based evaluations. Since for this target the improvements regarding loss do not match those regarding quality, the relative loss is sharply increasing between 1000 and 5000 evaluations. Hence, the behavior of the Pareto-wise relative loss indicator from Figure 7 is explained by the inability to reduce average losses with the same efficiency with which the target is improved.

[5]Between the different targets of the same electrical drive designs, as well as between the same target in different designs.

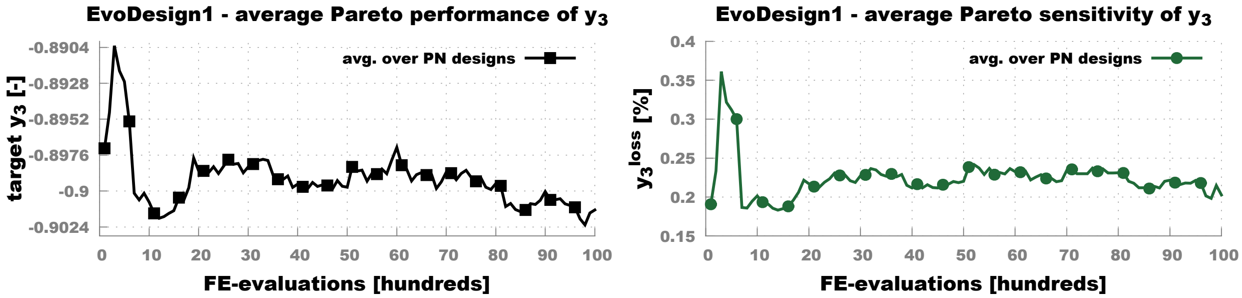

Fig. 6. Evolution of the average Pareto-wise target performance (left) and average Pareto-wise expected relative target loss (right) for y_3 of EvoDesign1.

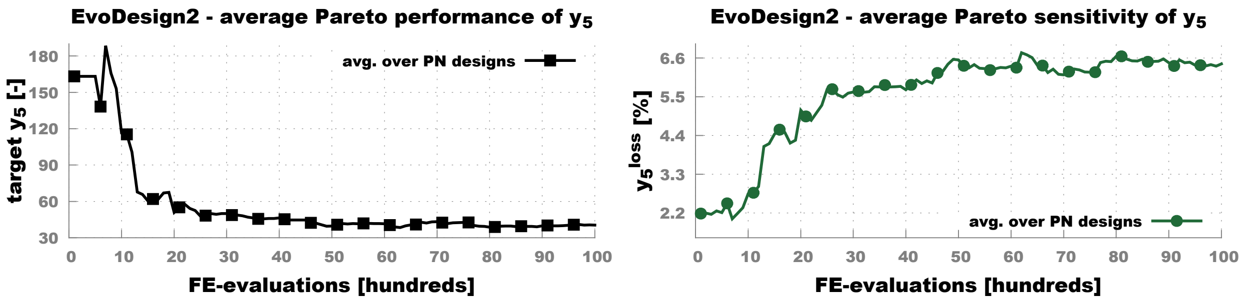

Fig. 7. Evolution of the average Pareto-wise target performance (left) and average Pareto-wise expected relative target loss (right) for y_5 of EvoDesign2.

VI. Conclusions and Future Work

In the present work we have introduced a practical two-stage tolerance/sensitivity analysis strategy based on the min-max approach and we have tested this strategy on industrial problems from the field of electrical drive design. We have also shown that global surrogate (non-linear regression) models intended for speeding-up multi-objective design optimization runs can also be used to rapidly perform very insightful (and to a certain degree accurate) tolerance analyses of the evolved designs.

Encouraged by the good performance on individual motor designs, we also used our new surrogate-based approach to track the dynamics of the average sensitivity of the Pareto front over the entire optimization run. Our results indicate that the applied evolutionary algorithms are not entering a stage where they discover electrical drive designs that generally trade robustness for quality. Nevertheless, our analysis of Pareto-wise sensitivity evolution revealed that there are certain optimization targets (objectives) for which the average improvements regarding sensitivity are much smaller than the average improvements regarding quality. This suggests that, in the future, it might make a lot of sense to define the tolerance/sensitivity of a design as an optimization objective in its own right.

Another idea that seems very promising for future work is centered around switching the LHS with a more efficient sampling method like the one proposed by Kato et al. in [25].

Acknowledgment

This work was conducted within LCM GmbH as a part of the COMET K2 program of the Austrian government. The COMET K2 projects at LCM are kindly supported by the Austrian and Upper Austrian governments and the participating scientific partners. The authors thank all involved partners for their support.

References

[1] N. Queipo, R. Haftka, W. Shyy, T. Goel, R. Vaidyanathan, and P. Kevin Tucker, "Surrogate-based analysis and optimization," *Progress in Aerospace Sciences*, vol. 41, no. 1, pp. 1–28, 2005.

[2] A. Forrester, A. Sobester, and A. Keane, *Engineering Design via Surrogate Modelling: A Practical Guide*. Wiley, 2008.

[3] Y. Tenne and C. Goh, Eds., *Computational Intelligence in Expensive Optimization Problems*, ser. Adaptation, Learning and Optimization. Springer, 2010, vol. 2.

[4] F. Neri, X. del Toro Garcia, G. L. Cascella, and N. Salvatore, "Surrogate assisted local search in PMSM drive design," *COMPEL: The International Journal for Computation and Mathematics in Electrical and Electronic Engineering*, vol. 27, no. 3, pp. 573–592, 2008.

[5] D. Stephens, D. Gorissen, K. Crombecq, and T. Dhaene, "Surrogate based sensitivity analysis of process equipment," *Applied Mathematical Modelling*, vol. 35, no. 4, pp. 1676–1687, 2011.

[6] M. J. Powell, "Radial basis functions for multivariable interpolation: a review," in *Algorithms for approximation*. Clarendon Press, 1987, pp. 143–167.

[7] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 2nd ed. Prentice Hall Inc., 1999.

[8] R. Collobert and S. Bengio, "SVMTorch: Support vector machines for large-scale regression problems," *Journal of Machine Learning Research*, vol. 1, pp. 143–160, 2001.

[9] C. B. Storlie, L. P. Swiler, J. C. Helton, and C. J. Sallaberry, "Implementation and evaluation of nonparametric regression procedures for sensitivity analysis of computationally demanding models," *Reliability Engineering & System Safety*, vol. 94, no. 11, pp. 1735–1763, 2009.

[10] T. Goel, R. T. Haftka, W. Shyy, and N. V. Queipo, "Ensemble of surrogates," *Structural and Multidisciplinary Optimization*, vol. 33, no. 3, pp. 199–216, 2007.

[11] C. M. Creveling, *Tolerance design: a handbook for developing optimal specifications*. Prentice Hall, 1997.

[12] C. A. Coello Coello, "20 years of evolutionary multiobjective optimization: What has been done and what remains to be done," in *Computational Intelligence: Principles and Practice*. IEEE Computational Intelligence Society, 2006, pp. 73–88.

[13] S. Silber, W. Koppelstätter, G. Weidenholzer, and G. Bramerdorfer, "Magopt-optimization tool for mechatronic components," in *Proceedings of the ISMB14-14th International Symposium on Magnetic Bearings*, 2014.

[14] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.

[15] E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: Improving the strength Pareto evolutionary algorithm for multiobjective optimization," in *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001)*. International Center for Numerical Methods in Engineering (CIMNE), 2002, pp. 95–100.

[16] A.-C. Zăvoianu, G. Bramerdorfer, E. Lughofer, S. Silber, W. Amrhein, and E. P. Klement, "Hybridization of multi-objective evolutionary algorithms and artificial neural networks for optimizing the performance of electrical drives," *Engineering Applications of Artificial Intelligence*, vol. 26, no. 8, pp. 1781–1794, 2013.

[17] A.-C. Zăvoianu, E. Lughofer, G. Bramerdorfer, W. Amrhein, and E. P. Klement, "An effective ensemble-based method for creating on-the-fly surrogate fitness functions for multi-objective evolutionary algorithms," in *Proceedings of the 15th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC 2013)*. IEEE Computer Society Conference Publishing Services (CPS), 2014, pp. 237–244.

[18] C. Wu and M. Hamada, *Experiments: planning, design, and parameter optimization*. Wiley, New York, 2000.

[19] M. D. McKay, R. J. Beckman, and W. J. Conover, "Comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, vol. 21, no. 2, pp. 239–245, 1979.

[20] K. Deb and R. B. Agrawal, "Simulated binary crossover for continuous search space," *Complex Systems*, vol. 9, pp. 115–148, 1995.

[21] K. Deb and M. Goyal, "A combined genetic adaptive search (GeneAS) for engineering design," *Computer Science and Informatics*, vol. 26, pp. 30–45, 1996.

[22] A.-C. Zăvoianu, E. Lughofer, W. Amrhein, and E. P. Klement, "Efficient multi-objective optimization using 2-population cooperative coevolution," in *Computer Aided Systems Theory - EUROCAST 2013*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013, pp. 251–258.

[23] A.-C. Zăvoianu, E. Lughofer, G. Bramerdorfer, W. Amrhein, and E. P. Klement, "DECMO2: a robust hybrid and adaptive multi-objective evolutionary algorithm," *Soft Computing*, vol. [available online] DOI: 10.1007/s00500-014-1308-7, pp. 1–19, 2014.

[24] D. Thain, T. Tannenbaum, and M. Livny, "Distributed computing in practice: the condor experience." *Concurrency - Practice and Experience*, vol. 17, no. 2-4, pp. 323–356, 2005.

[25] T. Kato, K. Inoue, Y. Yano, and S. Oshio, "Parallel tolerance analysis of a power electronic converter by the genetic algorithm with the island model," in *Power Electronics and Motion Control Conference, 2004. IPEMC 2004. The 4th International*, vol. 3. IEEE, 2004, pp. 1659–1664.