ESI LIVE 2023
INNOVATION TALK

Advanced reduced order modeling for more effective structural crash optimization

Gavin England, Senior Product Marketing Manager, ESI Group
Fatima Daim, Team Leader Computational Science, ESI Group

In this article, we're diving into ESI's game-changing technology called ADMORE. This isn't just another tech tool—ADMORE represents a leap forward in how engineers approach structural crash optimization. Imagine having a tool that not only speeds up the design process but also pinpoints the most crucial factors for achieving optimal safety and performance in vehicle structures. That's exactly what ADMORE does. It transforms complex simulations into more manageable, insightful, and actionable data. But before we get into the nitty-gritty and see a real-life use case from our partnership with Renault, let's set the scene with some background on why this technology is such a breakthrough.

What is the current dilemma automakers face when developing new vehicle concepts?

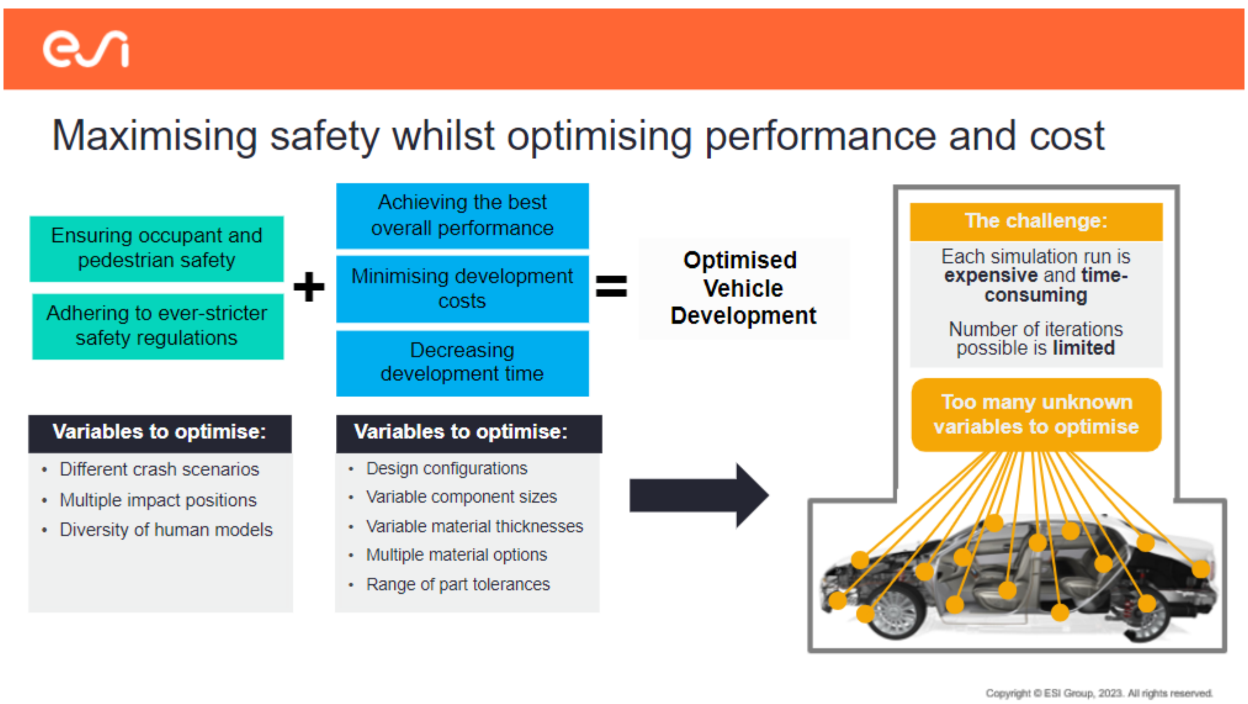

To optimize vehicle safety against performance, many unknowns are considered and assessed. To ensure occupant and pedestrian safety, and adhere to ever-stricter regulations, those unknown variables include many different crash scenarios and impact positions as well as an increasing diversity of human models. At the same time, to achieve the best vehicle performance at minimal cost and in the shortest possible lead times, variables such as different design configurations, component size options, material types, thicknesses, and the impact of the range of component tolerances need to be considered. These variables usually affect each other as they change, so they shouldn't be assessed in isolation. Some combinations of variable values will have a big influence on the design, either positively or negatively, whilst others won't. When you consider the different permutations of these combined variables, there are potentially many thousands of possible outcomes. To fully explore those unknowns, to assess every combination of variables using simulation, to deliver the best product, is simply not possible. The number of simulation runs and design iterations is limited by time and cost.

So how do you know which of these unknowns have the biggest influence on the design, and, therefore, which variable values should you focus those precious simulation runs on to achieve the best crash simulation results whilst optimizing vehicle performance? This is the challenge facing many automotive manufacturers.

To solve this, ESI developed a unique solution, called 'ADMORE'. What is it about and how does it help solve this problem?

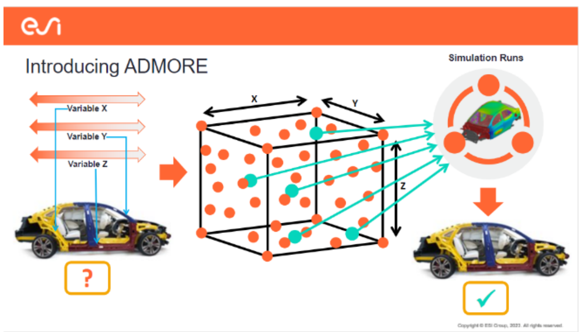

To help solve this and other similar problems, ESI has developed ADMORE, an intelligent model order reduction technology. ADMORE quickly guides engineers to the variables that have the most influence on the product, and so are the most important to assess, so that they know exactly where to focus their simulation runs. It does this using advanced model order reduction, driven by intelligent machine learning technology. It is a powerful decision-support tool for engineers.

- ADMORE enables engineers to build an intelligent parametric crash simulation model and carry out real-time design space exploration to quickly examine the impact of changing multiple parameters, to assess which have the greatest influence on crash scenarios. For example, which combination of parameter values is likely to optimize the strength-to-weight ratio of the part of the vehicle, or, conversely, could potentially cause a safety or regulatory issue in a crash situation?
- Whether those parameters are material thicknesses of structural components, limits of manufacturing tolerances, or different impact positions, ADMORE enables fast parametric design space exploration to quickly assess those unknowns. It provides quick insight for engineers, who will have a much clearer understanding of the direction they should take those limited number of simulation runs to deliver the most meaningful, beneficial crash simulation results.
- ADMORE enables both live manual and automated design space exploration to quickly and easily determine the priority load cases and design parameters to analyze in more detail. The nature of ADMORE, with its intelligent automation and real-time assessment capabilities, means that it can be used by both experienced and non-experienced engineers, helping to democratize simulation tasks.
- ADMORE optimizes simulation cost v accuracy. It enables a better, more optimized product design for the same investment and effort.

With regards to crash simulation, what are the details on this specific use case?

ADMORE can be applied to many different use cases, particularly where there are multiple combinations of variables to assess. For crash simulation, it brings value in three main areas:

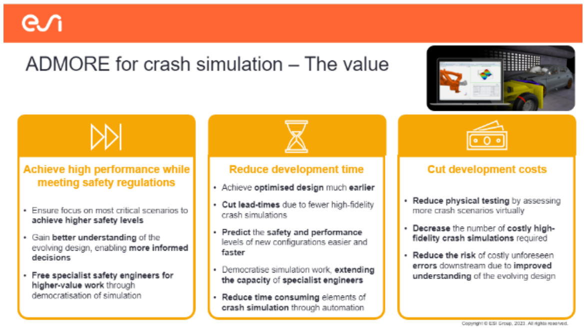

Firstly, it enables the development of a high-performing vehicle whilst ensuring safety regulations are easily met. Being able to focus simulation runs on the most critical scenarios results in improved safety outcomes. Having a better understanding of the evolving vehicle design enables more informed development decisions to be made. And the democratization of simulation activities frees specialist engineers to work on other high-value activities, for a more optimized vehicle design.

Secondly, it helps reduce vehicle development time. Lead times are cut due to the decrease in high-fidelity crash simulations required, new vehicle configurations can be assessed for safety much faster, democratization of the simulation processes helps release additional specialist engineering capacity, and automation further helps cut the need for time-consuming, laborious tasks.

And, finally, it helps to decrease development costs. The need for expensive physical testing and high-fidelity simulations is further reduced, and the improved understanding of the evolving design reduces the risk of potentially costly unforeseen downstream errors.

Looking at the mechanics behind ADMORE: How is a high fidelity model typically used in simulation and how is optimization study commonly carried out?

Let us start by introducing the classical numerical simulation scenario, which we refer to as the high-fidelity model.

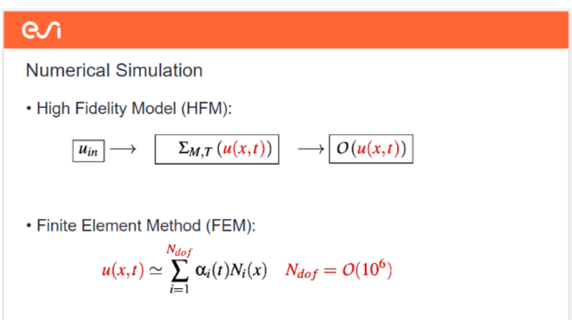

This High Fidelity model (HFM) describes the evolution of a physical model, given by discretized partial differential equations (PDE) denoted here as sigma, on a mesh M during time interval T. This evolution is described by the variable u at each space point x and each instant t, starting from the initial condition denoted Uin. In crash simulation, U is the displacement, velocity, or acceleration. In general, for such simulations, what interests us is the performance of this physical system denoted here by O, which depends on the variable U. In the crash simulation, these performances are a set of intrusions measuring the occupant safety during the crash.

The performances are, in general, obtained by post-processing the variable U. They are usually very quick to calculate once U is computed. What is expensive, however, is the computation of the U variable itself, which in the case of the finite element method, is given as a combination of shape functions denoted Ni. The coefficients of this combination are computed at each time step by solving a linear system. in the case of implicit time discretization. Or under a set of linear systems imposed by the CFL condition in the case of an explicit scheme, which is the case in crash simulation. This combination is strongly linked to the fineness of the mesh. The finer the mesh, the higher the computational cost.

get it right

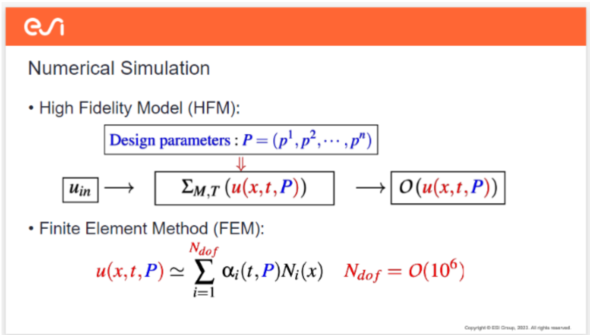

This complexity increases further when we consider design parameters, denoted here as P. These could be, for example, the varying thicknesses of structural parts. Therefore, the variable U depends also on those parameters. In structural car optimization, we need to evaluate the performance for each configuration of parameters.

When carrying out an optimization study, the classical workflow typically looks something like this:

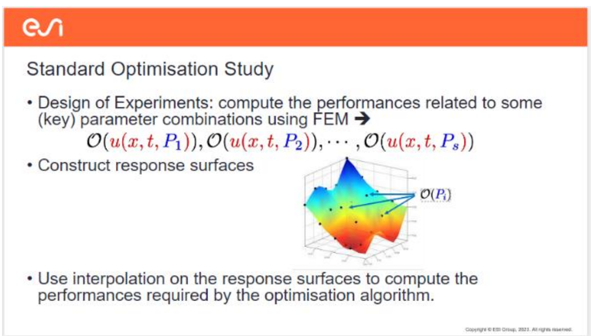

1. We first develop a design of experiments (DOE) where some key parameter combinations are defined, denoted here as P1, P2, etc.
2. We then launch several high-fidelity model computations to obtain the related variable U and then use post-processing to obtain the related performances
3. Next, we construct the response surfaces based on these performances. This surface is built based on an algebraic tool to interpolate between the performances related to DOE.0
4. Finally, we use interpolation on the response surfaces to compute the performances required by the optimization algorithm

How using ADMORE, can a reduced order model be built?

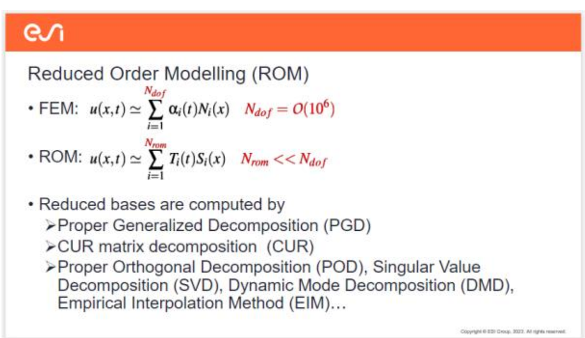

By introducing reduced order models, we aim to drastically reduce the description of the variable U. Instead of looking for U as the combination of a large number of shape functions, a reduced order model (ROM) builds U as a combination of a small number of functions, denoted here by Si, known as reduced basis or modes. The response surfaces reduce the complexity of the performances, which are functions dependent on the variable U. Essentially, the reduced order model tackles the cause of the computational complexity in the calculation of U.

To define such reduced bases, we can use a variety of methods. In this presentation, we will focus on just two, the Proper Generalized Decomposition (PGD) and Column-Row Decomposition (CUR) methods.

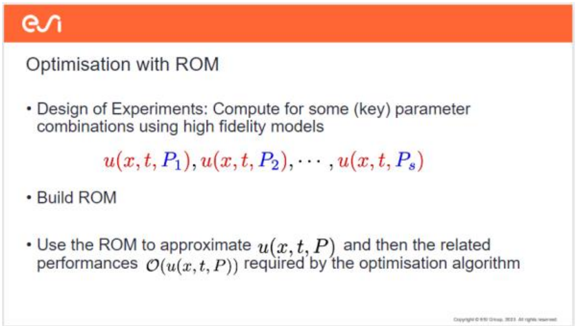

Let's, first, summarise how the reduced order model will be used in an optimization workflow:

1. We start by developing a design of experiments, where some key parameter combinations are defined, denoted here as P1, P2, and so on
2. We then launch the high-fidelity model computations to obtain the related variables U(P1), P2, etc.
3. The reduced order model is then built, which allows the computation of the variable U for any combination of parameters
4. Then we use the reduced order model to approximate the variable U for any combination of parameters and post-process it to obtain the related performances required by the optimization algorithm.

With regards to the sPGD method, how can this approach be applied to an optimization study?

sPGD stands for sparse Proper Generalized Decomposition and was developed by ESI's Scientific Director Prof. Francisco Chinesta. This method considers the parameters as extra coordinates of space and time. It also approximates the variable U as a combination of a separated representation of the space, time, and parameter coordinates. This combination contains a reduced number of functions, or modes, depending separately on space, time, and parameters.

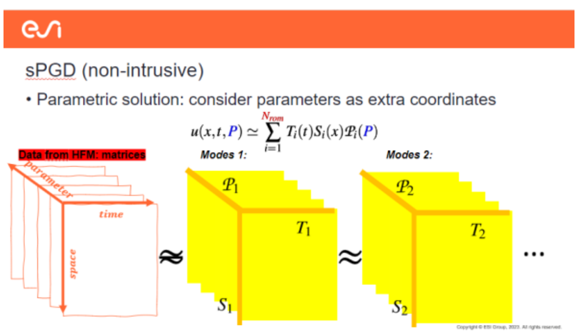

How do we do that? Well, once the design of experiments is defined and the related high-fidelity model computation is launched, we collect, for each computation of parameter combination, the variable U as a matrix. To illustrate this decomposition, we develop a 3D representation of this - one axis for space, one for time, and one for the parameters. The sPGD method collects these modes iteratively. For the first mode, we look for 3 vectors whose product best approximates the cube. For the second mode, the three vectors whose product best approximates the residue are found, and so on. We continue until we reach a good approximation by collecting N reduced order model modes representing the matrices U of the high-fidelity model.

What is the second reduced order model method "CUR" about and how can we apply this approach to an optimization study?

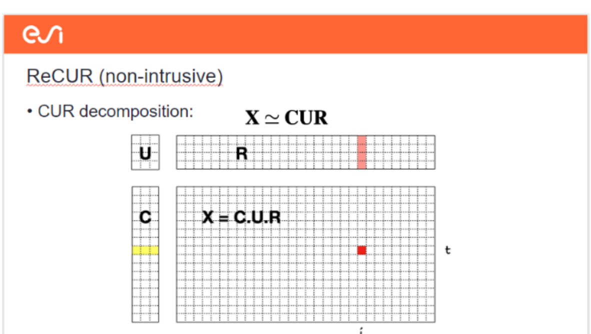

The second method is based on the CUR decomposition usually used to compress matrices. For example, for a matrix denoted X, we approximate it by the products of three matrices C, U, and R. C is a set of selected columns of X, R is a set of selected rows, and U is a small matrix combining these columns and rows. This also means that each cell in X can be explained or approximated by a combination of a product of one column and one row.

For a parametric study, we extend the CUR decomposition to approximate Xp by regression for any parameter combination outside the DoE. Re-CUR comes from the phrase ‘Regression CUR’.

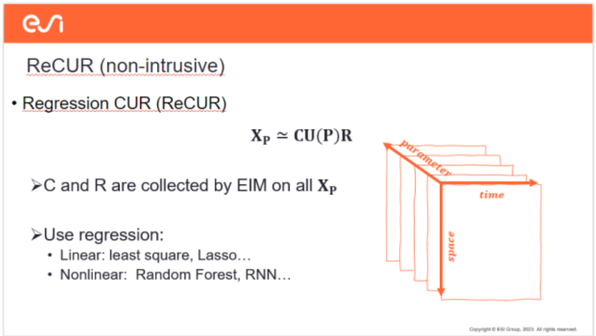

From high-fidelity simulations, we collect the Xp matrices containing displacements of the variable U. We apply the Empirical Interpolation Method to all the Xp matrices to collect columns and rows denoted, respectively, C and R. Each Xp is approximated by a CUR decomposition where the link with parameters is supported by the matrix Up. To interpolate for any other combination of parameters, we have to interpolate the Up matrices. To achieve this, we can use a linear regression, such as a least squares or lasso method, or a nonlinear regression via machine learning. In this study, we used the random forest method.

Is there a real-life case study of applying ADMORE to a specific crash simulation case?

There’s a use case where we deployed this approach for crash simulation, specifically to assess vehicle-side member reinforcement methods. This is based on work we carried out in partnership with Renault.

This is an example of a Mobile Progressive Deformable Barrier or MPBD test. Here we used a high-fidelity model with 6 million nodes and 8 million elements. We are interested in 31 different intrusions to measure and assess occupant safety, and, bearing in mind that one high-fidelity simulation took 12 hours on 144 processors, this number of quantities of interest will have a significant effect on the simulation performance.

The optimization considered in this particular study is a reinforcement method. This method is applied during the project vehicle development when the performance of the nominal design, determined by the crash development team, deviates from the targets and no longer satisfies the safety criteria. This deviation is typically due to constraints imposed by different service teams involved in the project development. This method looks to reinforce or reduce local zones in the structure using triggers such as safety targets.

In the case of the model in our study typically the side member is divided into 23 local zones to be weakened or reinforced by varying each thickness zone by +- 0.2 mm. These local zones are those marked in blue and red.

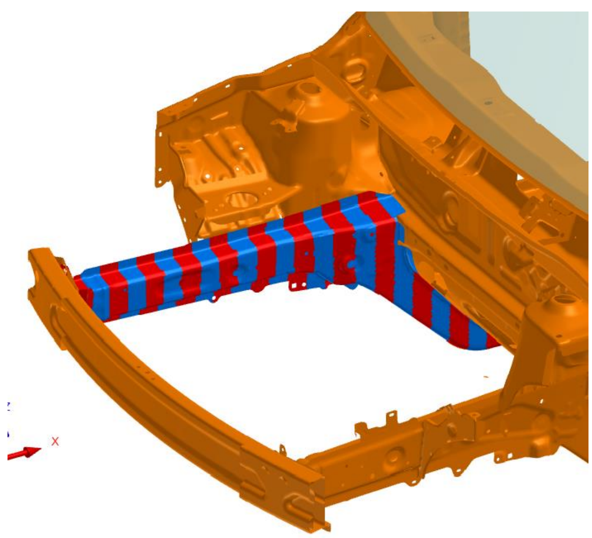

In practice, and to identify the local zone to reinforce or reduce, the engineer starts by varying each zone one at a time using high-fidelity models to find a solution. In this case that equates to 46 high-fidelity simulation runs, which takes a significant amount of time and computing cost. If any solution is found from those 46 runs, it is sent to the optimization team for further work.

What was the goal of this study with Renault?

The objective of this study is to use ADMORE to accelerate this process and find a solution through significantly fewer high-fidelity computations. For this optimization problem, with those 23 parameters, high-fidelity simulation runs need to be first carried out to build the reduced-order model. Using sPGD, 24 runs are required, or, using ReCUR, just 12 runs are needed. We then carry out an iterative, or combinatorial, exploration of the zones to reinforce or to weaken to meet the safety targets:

- In iteration one, we explore configurations that reduce or weaken one zone at a time. So, for those, 23 zones, 46 configurations need to be assessed, using the two reduced-order modeling methods, sPGD and ReCUR. Among these predictions, we select the best candidate giving the best performance and we launch a related high-fidelity simulation to validate this candidate.
- In iteration two we explore configurations that reduce or weaken two zones at a time, which requires 1012 predictions to be carried out with the two ROMs. We, again, choose the candidate giving the best performance and launch the relevant high-fidelity simulation.

How does the solution look like in the software?

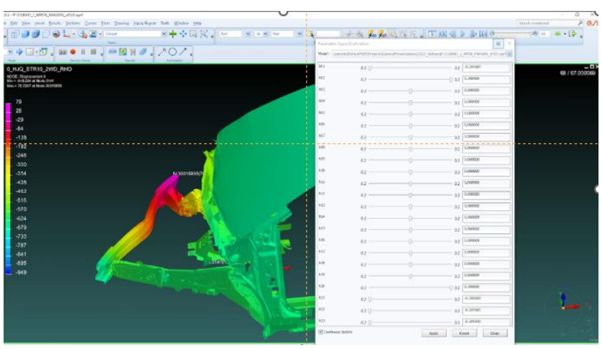

In this video you can see the reduced order model using sPGD, where our 23 parameters can be varied using dynamic sliders. The sliders enable the parameters to be varied, changing the thicknesses of the 23 zones from -0.2 mm, weakening the zone, to 0.2mm, reinforcing the zone. The sliders are all on zero at the start, which represents the nominal case, so we can see the shape of the nominal side members. Moving the sliders for one value to another shows how the side members behave and deform with reinforcing or weakening of the zones.

What are the results that Renault got from this work?

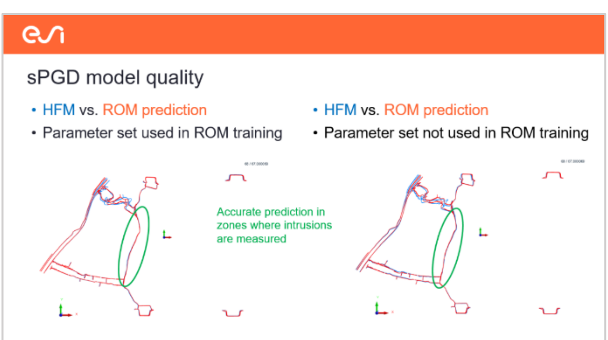

This image shows the quality of the reduced order model built via the sPGD method. We have a section view of the car during the crash simulation, comparing the high-fidelity in blue and, in red, the predicted results using sPGD. The view on the left shows a simulation run where the parameters are set to the same as one of the 24 configurations already used in the design of experiments training. The view on them left uses parameter configurations outside of those used in the design of experiments stage.

If we focus on the frame where the intrusions are at the maximum, we can see that the predictions are quite good, specifically on the footwell area where the intrusions are measured for the occupant safety.

The image below shows a comparison between, in blue, the high fidelity computations of the nominal which does not satisfy the safety targets, and, in red, the high fidelity computations of the best candidates identified by the sPGD model.

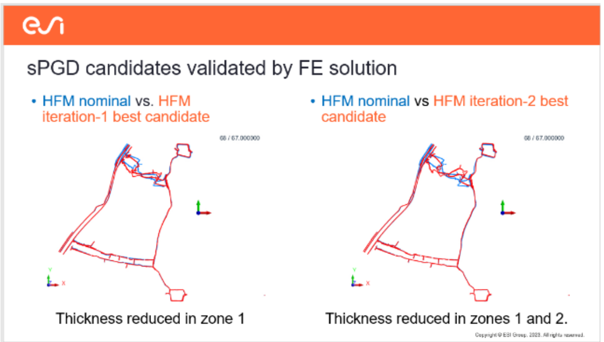

- On the left, showing iteration 1, where we reduce or weaken one zone at the time, zone 1 was selected to be weakened.
- On the right, for iteration 2, where we reduce or weaken two zones at the time, the adjacent zones 1 and 2 were selected to be weakened.

When focusing on the frame where the intrusions are at the maximum, we can clearly see the difference in deformation.

Similar to the previous example, where we used the sPGD method, this shows the quality of the reduced order model built, this time, via the ReCUR method. As before, we have a section view of the car during the crash simulation, comparing the high-fidelity in blue and, in red, the predicted results using ReCUR.

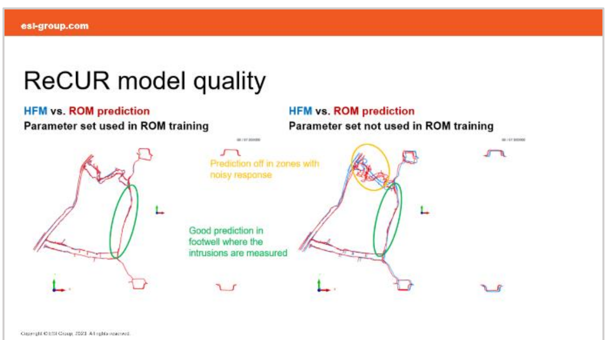

The view on the left shows a simulation run where the parameters are set to the same as one of the 12 configurations already used in the design of experiments training. The view on them left uses parameter configurations outside of those used in the design of experiments stage.

If we focus on the frame where the intrusions are at the maximum, where they are measured, we can see that the predictions are, again, good, specifically on the footwell area where the intrusions are measured for the occupant safety.

Again, as with the previous sPGD example, here we see a comparison between, in blue, the high fidelity computations of the nominal which does not satisfy the safety targets, and, in red, the high fidelity computations of the best candidates identified by the ReCUR model.

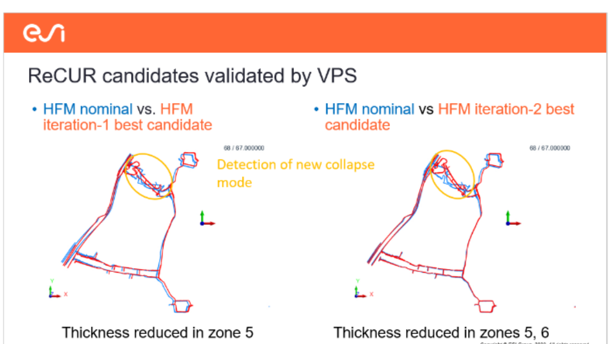

- On the left, showing iteration 1, where we reduce or weaken one zone at the time, zone 5 was selected to be weakened.
- On the right, for iteration 2, where we reduce or weaken two zones at the time, the adjacent zones 5 and 6 were selected to be weakened.

When focusing on the frame where the intrusions are at the maximum, we can clearly see the difference in deformation where new collapse modes were detected by weakening zones 5 and 6 which appear to be critical.

What are the conclusions that were reached during this study?

During their presentation at ESI Live 2023, Fatima summarized the main project success as follows: *"First and foremost, we are happy that it was confirmed that Renault plans to use these methods in car development projects."*

There are three main outcomes that were reached so far:

- Optimisation studies for crash simulations using reduced order models significantly reduces the computational time and costs compared to a classical response surface optimisation approach
- The two reduced order model methods considered for this particular real-life industrial case were both effective in optimising side members, though actual engineering teams will use different solution, for example triggers.
- Both methods have successfully identified the front of the side member as candidate zones to be weakened.