Surrogate Modeling using Physics-guided Learning

Ali Ozdagli
Florida Gulf Coast University
Fort Myers, FL, USA
aozdagli@fgcu.edu

Peter Volgyesi
Vanderbilt University
Nashville, TN, USA
peter.volgyesi@vanderbilt.edu

Xenofon Koutsoukos
Vanderbilt University
Nashville, TN, USA
Xenofon.Koutsoukos@vanderbilt.edu

ABSTRACT

Computer simulation models are used extensively in scientific and engineering problems for complex design tasks and decision processes. Surrogate models generated using data-driven techniques can approximate the behavior of complex simulation models with high fidelity and can accelerate the design process. This paper presents a physics-guided learning architecture that integrates parameters extracted from physics-based simulations into the intermediate layers of a neural network to constrain the learning process during the training of surrogate models and to improve their generalization. The proposed architecture is used to develop a surrogate model for evaluating the structural integrity of the hull of an unmanned underwater vehicle. It is shown that physics-guided learning can improve generalization in less explored regions of the design space compared to black-box models. In addition, the architecture improves the explainability of the model predictions using physics-based parameters and allows the designer to make decisions based on the input and physics-based intermediate parameters.

KEYWORDS

surrogate modeling, physics-guided learning, system design

ACM Reference Format:
Ali Ozdagli, Peter Volgyesi, and Xenofon Koutsoukos. 2023. Surrogate Modeling using Physics-guided Learning. In *Cyber-Physical Systems and Internet of Things Week 2023 (CPS-IoT Week Workshops '23), May 09–12, 2023, San Antonio, TX, USA*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3576914.3587532

1 INTRODUCTION

Computer simulation models are used extensively in scientific and engineering problems to accelerate complex design tasks and decision processes. While through model calibration and parametric sensitivity analysis, it is possible to design high-fidelity simulations, such exercise can be computationally exhaustive. Especially, for design problems that involve multiple heterogeneous domains, running a large number of complex simulations, which is often a hard requirement, is not time- and cost-effective [13].

Surrogate modeling (also known as metamodeling [5] and digital twin modeling [6]) focus on employing computationally efficient surrogates. In essence, the main motivation is the effective utilization of limited computational resources. The surrogate models often rely on data-driven techniques to approximate the behavior of the complex simulation model. Surrogate models based on neural networks, in particular, are a viable method since they are capable of learning complex nonlinear relationships at a relatively low computation cost [12, 14, 16].

Typical surrogate modeling methods based on neural networks employ black-box models resulting in an inference procedure that is non-transparent to the user or the designer. While black-box models can be used to generate efficient and accurate surrogates for predicting simulation responses, they often fail in generalizing over the less explored regions of the design space. Further, even if the input/output relation is represented accurately by the neural network, the designer may have difficulty in interpreting and explaining the results for those regions since black-box layers provide very limited information about the physics of the actual system. The objective of this paper is to use physics-guided learning to improve the generalization and transparency of surrogate models used in the design of cyber-physical systems. Physics-guided learning has been used before in various scientific and engineering domains to improve generalization [4, 11, 17].

The paper presents a physics-guided learning (PGL) architecture that integrates parameters extracted from physics-based simulations into the intermediate layers of the neural network to constrain the learning process during training and improve the generalization. The architecture is used to develop a surrogate model for evaluating the structural integrity of the hull of an unmanned underwater vehicle (UUV) and it is shown that it can improve generalization in less explored regions of the design space compared to black-box models. The architecture also allows the designer to improve the explainability of the model predictions using physics-based parameters. In particular, the proposed approach uses layer-wise propagation (LRP) [10] to determine the most relevant physics-based parameters and inputs that contribute to the prediction for a given design. Such information can aid design space exploration and the generation of new designs more effectively.

The paper first presents a brief discussion of the problem and the goals of surrogate modeling. Then, the PGL architecture used for training the surrogate model and the LRP approach for improving the explainability of the model is presented. Next, the proposed approach is evaluated using various designs of the hull of a UUV. Finally, conclusions and directions for future work are presented.

2 PROBLEM FORMULATION

Autonomous vehicles are a significant research area that can potentially impact domains such as transportation and warfare. A key challenging aspect is the design of such cyber-physical systems (CPS) to satisfy operation and mission requirements. The complex

CPS-IoT Week Workshops '23, May 09–12, 2023, San Antonio, TX, USA

ACM ISBN 979-8-4007-0049-1/23/05…$15.00
https://doi.org/10.1145/3576914.3587532

interaction between subsystems that involve multiple domains leads to very inefficient and expensive processes that require repeated executions of high-fidelity computer simulations [3].

Structural integrity is a crucial design concern for autonomous vehicles that can be addressed by performing high-fidelity finite element simulations. The proposed approach seeks to develop a surrogate model for evaluating the structural integrity of a UUV hull. Specifically, the aim is the design of pressure vessels inside the fairing of the UUV (see Figure 1). In a typical UUV, the shape of the hull depends on the placement and the size of components such as the motor, battery, and sensors. Each of these components can be optimized in an iterative fashion to improve the behavior of the vehicle. When the design of such components changes, the geometric properties of the hull need to be readjusted to accommodate the component placement. The structural integrity of a new hull can then be evaluated using finite element simulations. Such a design process requires repeated executions of simulations that can slow down the overall design and optimization process.

For complex problems, this pattern implicates the use of computationally exhaustive simulations. A neural network-based surrogate model can significantly accelerate this process. Moreover, the integration of physics knowledge into the training of the surrogate model can improve the generalization and explainability of the model.

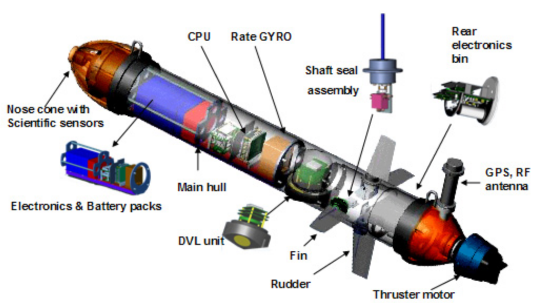

Figure 1: Hull example [15]

3 SURROGATE MODELING USING PHYSICS-GUIDED LEARNING

The main objective of the paper is to design a surrogate model using physics-based learning. The first step of the approach is to define the data generation scheme and design space exploration procedure. Figure 2 illustrates a typical surrogate modeling approach. Here, X is the set of design parameters used to generate data using a simulation model and the output Y is the response generated by the simulator. Analysis of the data can indicate if the selected design parameters result in a design that satisfies the requirements.

In our case study, we consider feasibility as the design requirement. Feasibility is measured in terms of the maximum stress a UUV hull design experiences for a given water depth. If the maximum stress is close to the allowable stress limit dictated by the material properties, the feasibility of the design as measured by the feasibility ratio of the maximum over the allowable stress is low. The maximum stress a design experiences also depends on the geometric shape of the hull. The various parameters, such as geometry, material properties, and feasibility ratio for the UUV hull are defined in Table 1 and 2.

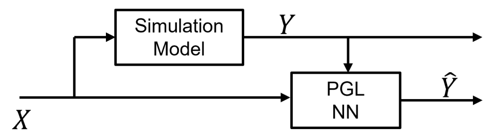

Figure 2: Typical approach for surrogate modeling

Data generation starts with designing a series of experiments. We follow a workflow illustrated in Figure 3. It is assumed that an oracle generates the parameters of a design and we need to establish if the design satisfies the requirements which in this case means that the UUV hull remains structurally intact in a given depth underwater. Based on the design parameters, a finite element model (e.g., Ansys [1]) is used to generate UUV hull designs and simulate their responses. When the response complies with the design requirement, i.e. maximum stress is below the allowable stress limit, the design is considered feasible.

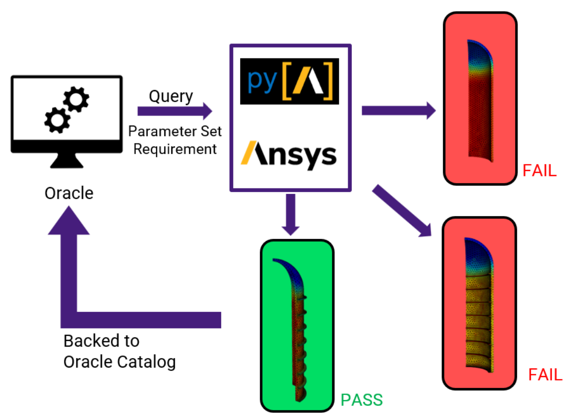

Figure 3: Automatic hull generation

For this study, we consider the UUV hull design as a simple cylindrical capsule model with spherical end caps as shown in Figure 4. The design parameters are listed in Table 1. The first four parameters (grayed) are related to the properties of the material selected. The next three parameters dictate the geometry and determine the hull shape. The external hydraulic pressure σ_{hyd} determines the pressure that acts on the vessel and depends on the water depth for which the simulation is performed. It should be noted that this parameter is directly related to the design requirement that indicates the pressure that needs to be sustained by the UUV. The last four parameters are used for capsule designs that have stiffeners and specify the stiffener shape, number, and location. Examples of inner and outside stiffened designs are illustrated in Figure 3.

The outputs of the FEM simulations are shown in Table 2. An important output is the magnitude of the maximum Von Mises

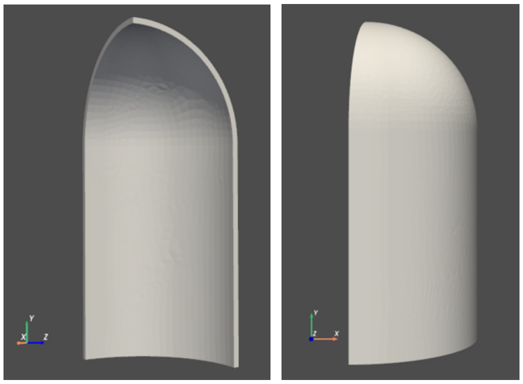

(a) (b)

Figure 4: A representative capsule model (1/8 cut): (a) inside view; (b) outside view

Table 1: Design Parameters

Parameter	Description
E	Material Elasticity
σ	Material Strength
v	Poisson Ratio
ρ	Material Density
r_i	Inner Diameter
t	Thickness
l	Cylinder Length
σ_{hyd}	External Hydraulic Pressure
n	Number of Stiffeners
t_s	Thickness of Stiffener
h_s	Height of Stiffener
l_s	Location of Stiffener

Table 2: Simulation Outputs

Parameter	Description
f_r	Feasibility ratio
$\sigma_{x,y,z}$	Maximum Directional Stresses
$\sigma_{xy,yz,xz}$	Maximum Plane Stresses

stress. For simplicity, we consider the feasibility ratio which is defined as the ratio of the maximum Von Mises stress σ_{max} over the nominal material strength σ and used directly for evaluation of the design requirement. Typically, it is desirable to include an additional safety margin represented by a feasibility ratio smaller than one. We also consider internal stresses that can be generated by the simulator. These parameters are used as intermediate variables in the surrogate model to improve generalization and explainability. Our motivation for selecting these parameters comes from the fact that the computation of the Von Mises stress relies on first determining these internal stresses [9]. The process involving the generation of design parameters, simulation, and data collection can be automated and repeated multiple times to generate the training and testing data for the surrogate model.

In our previous work, we have developed a PGL architecture for structural health monitoring [11]. The objective is to close the gap between the simulation and experimental domains by introducing parameters extracted from the physics-based simulation model into the intermediate layers of the neural network to constrain the learning process. This paper follows a similar idea for improving generalization in less explored regions of the design space.

The architecture used for surrogate modeling of the UUV hull is shown in Figure 5. The inputs correspond to the design parameters generated by the oracle and the output is the feasibility ratio between the maximum Von Mises stress and the allowable stress. G_f (green) is the block of neural network layers that extract latent features from the given input and G_y (blue) is the block of layers that are used for predicting the feasibility ratio. The architecture includes an intermediate layer G_{pgl} (yellow) that is associated with the absolute maximum internal stresses denoted by z.

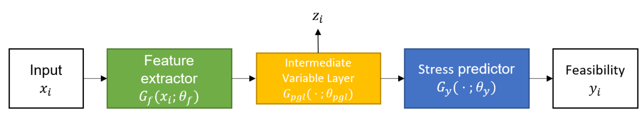

Figure 5: PGL Surrogate Model

The physics-based parameters are generated using the simulator and are incorporated into the training of the surrogate model by considering the following loss function:

$$\mathcal{L} = \frac{1}{n}\sum \mathcal{L}_y(y_i, \hat{y}_i) + \frac{\lambda_{pgl}}{n}\sum \mathcal{L}_{\mathrm{pgl}}(z_i, \hat{z}_i). \quad (1)$$

The first term is used to minimize the error between the predicted output and the output of the simulator and the second term between the predicted and simulated physics-based parameters. The loss function implies that the architecture utilizes a multi-task learning scheme where the first term of the loss corresponds to the regression error between the feasibility ratio computed by the simulation model and the predicted one. The second term corresponds to the mean square error (MSE) for the internal stresses where λ_{pgl} is a regularization parameter. With the latter term, we aim to constrain the learning process so that the intermediate physics-based parameters are close to the simulated ones and generalize the overall predictions.

In addition to improving generalization, the intermediate physics-based parameters can be used to explain the predicted output by estimating how strongly each parameter contributes to the output. In order to address this objective, we use layer-wise relevance propagation (LRP) [2]. The essential idea in LRP is to back-propagate relevance scores between the layers of a neural network using a set of purposely designed propagation rules. We follow the definition and implementation of the rules presented in [10] as well as the guidelines to choose the propagation rules at each layer.

4 EVALUATION

4.1 Implementation

We consider three capsule designs: plain hull, hull with inner stiffeners, and hull with outer stiffeners. The oracle uses Latin hypercube sampling to select design parameters for generating the simulation data [8]. 10,000 samples are obtained for each hull design by sampling parameters from regions of the design space defined by the lower and upper bounds shown in Table 3. The material selected for these experiments is A36 structural steel. To reduce the simulation time, 1/8 of the hull is modeled and simulated as shown in Figure 4. Symmetric boundary conditions are applied at the cross-section of the reduced capsule model to emulate full model behavior. Ansys is used for the finite element simulation [1] and the design generation and simulation process is automated using PyAnsys [7].

Table 3: Upper and lower bounds for design parameters.

Property	Lower Bound	Upper Bound	unit
r_i	7.5	20	in
t	0.125	1.125	in
l	20	50	in
σ_{hyd}	50	1500	psi
n	5	16	
t_s	0.125	1.125	in
h_s	0.125	1.125	in

After data generation, the data is divided into training and testing datasets with a ratio of 4:1. Using the training dataset, a PGL surrogate model for each of the three hull designs is trained in PyTorch with the architectural layout shown in Table 4. In addition, a black-box model without the physics-based intermediate variables is also trained using a loss function that minimizes only the regression error. Table 4 also shows the rules used in the LRP procedure. Since ReLU layers are absorbed by their preceding layer through the propagation, no special LRP rule is applied.

Table 4: Model Architecture

Layer Count	Layer Type	Notes	LRP Rule
1	Input		LRP-w^2
2	Linear		LRP-γ
3	ReLU		LRP-γ
4	Linear		LRP-γ
5	ReLU		LRP-γ
6	Linear	Intermediate	LRP-ϵ
7	ReLU		LRP-ϵ
8	Linear		LRP-ϵ
9	ReLU		LRP-ϵ
10	Linear		LRP-0
11	ReLU		LRP-0
12	Linear		-

4.2 Generalization

We evaluate the performance and generalization of the PGL surrogate model by (1) using the hold-out testing dataset and (2) generating a new testing dataset by selecting design parameters outside of the region of the design space considered in the training phase and running additional simulations. For the new testing dataset, we generate 100 samples from regions of the design space that are 10% above the upper bounds and 100 samples from regions that are 10% below the lower bounds of the design parameters using Latin hypercube sampling. We evaluate both the black-box and PGL surrogate models using the two testing datasets.

We consider metrics for evaluating the prediction performance of the feasibility ratio that include the mean square error (*MSE*), mean relative absolute error in percentage (*MAE*), and maximum relative absolute error in percentage (AE_{max}). Table 5 presents the performance of the black-box and PGL surrogate models for plain hull design based on these metrics. *Testing Dataset 1* denotes the hold-out testing dataset and *Testing Dataset 2* denotes the testing dataset with inputs outside the region of the design space considered in training. For all the metrics considered, the PGL model provides lower prediction error and better generalization over the black-box model. The results for the two designs with the stiffeners are similar. The metrics for the Testing Dataset 2 using the PGL model may be the result of a smaller sample size.

4.3 Explainability

To illustrate how the physics-based parameters of the PGL surrogate model are used for improving the explainability of the model predictions, we select a plain hull design with the following parameters: $r_i = 15.0\,in$, $t = 0.375\,in$, $l = 30.0\,in$. Here, we examine the LRP scores for the intermediate physics-based parameters, i.e., the scores that represent how strongly the internal stresses contribute to the value of the feasibility ratio. To generate the LRP scores, we sweep the external hydraulic pressure from its lower bound to the upper bound (50-1500 psi). The results are shown in Figure 6 where the external hydraulic pressure causing the feasibility ratio to exceed 0.5 is marked with a blue vertical line. The intensity of the red and blue colors for each line annotated with the name of the internal stress indicates the positive and negative correlation between the parameter and the output. The results show that for low levels of external hydraulic pressure, the directional stresses affect the feasibility ratio strongly. As the external hydraulic pressure increases, the effect of plane stresses becomes more dominant. We also compute the scores for the relevancy of the design parameters to the feasibility ratio. The scores indicate that the feasibility is driven mainly by the thickness parameter of the hull. The inner diameter has a secondary effect on the prediction, and the cylinder length does not affect the prediction substantially. The LRP results suggest that the designer should focus on calibrating the thickness and the inner diameter to fine-tune the feasibility of the design.

Figure 7 presents the relevancy scores for the hull designs with the stiffeners. While the results are similar to those from the plain model, the plane stresses contribute strongly to the feasibility ratio at higher external hydraulic pressure. The relevancy scores for the design parameters imply that the hull thickness and the inner diameter are the most relevant to the feasibility ratio. Parameters

Table 5: Performance and Generalization capability of black-box and PGL surrogate models

	Black-box			PGL		
Design Space	*MSE*	*MAE*	AE_{max}	*MSE*	*MAE*	AE_{max}
Testing Dataset 1	0.0001395	2.21	9.43	**0.0000967**	**2.18**	**8.96**
Testing Dataset 2	0.0005468	8.66	16.12	**0.0000684**	**3.01**	**6.32**

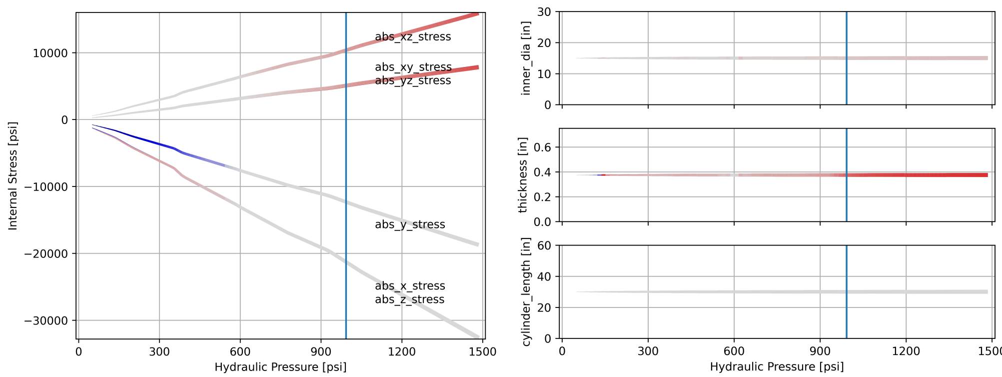

(a) **(b)**

Figure 6: Relevancy scores for plain hull design: (a) intermediate physical parameters; (b) design parameters

such as thickness, height, and number of stiffeners have a secondary relevance. Compared to the plain hull design, the cylinder length has a stronger effect on the feasibility ratio due to the inclusion of the stiffeners.

5 CONCLUSION

The paper presents physics-guided learning for surrogate modeling which improves generalization and explainability. For evaluation, we consider hull design of an unmanned underwater vehicle. The results demonstrate that the proposed approach achieves a better generalization and provides information to the designer about which design parameters are most dominant in the prediction. Future work includes designing a single surrogate model for the hull designs with different types of stiffeners for evaluating if the approach can handle heterogeneous inputs.

ACKNOWLEDGMENTS

This work is supported by DARPA through contract number FA8750-20-C-0537. Any opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the sponsor.

REFERENCES

[1] ANSYS. 2022. *Student Version, 2022 R2.* Ansys, Inc., Canonsburg, Pennsylvania.

[2] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one* 10, 7 (2015), e0130140.

[3] Adam Cobb, Anirban Roy, Daniel Elenius, and Susmit Jha. 2022. Trinity AI Co-Designer for Hierarchical Oracle-guided Design of Cyber-Physical Systems. In *2022 IEEE Workshop on Design Automation for CPS and IoT (DESTION)*. IEEE, 42–44.

[4] Arka Daw, Anuj Karpatne, William D Watkins, Jordan S Read, and Vipin Kumar. 2017. Physics-guided neural networks (pgnn): An application in lake temperature modeling. In *Knowledge-Guided Machine Learning*. Chapman and Hall/CRC, 353–372.

[5] Matthew Emerson and Janos Sztipanovits. 2006. Techniques for metamodel composition. In *OOPSLA–6th Workshop on Domain Specific Modeling*. 123–139.

[6] David Jones, Chris Snider, Aydin Nassehi, Jason Yon, and Ben Hicks. 2020. Characterising the Digital Twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology* 29 (2020), 36–52.

[7] Alexander Kaszynski. 2020. *pyansys: Python Interface to MAPDL and Associated Binary and ASCII Files.* https://doi.org/10.5281/zenodo.4009467

[8] Michael D McKay, Richard J Beckman, and William J Conover. 2000. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 42, 1 (2000), 55–61.

[9] R v Mises. 1913. Mechanik der festen Körper im plastisch-deformablen Zustand. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1913 (1913), 582–592.

[10] Grégoire Montavon, Alexander Binder, Sebastian Lapuschkin, Wojciech Samek, and Klaus-Robert Müller. 2019. Layer-wise relevance propagation: an overview. *Explainable AI: interpreting, explaining and visualizing deep learning* (2019), 193–209.

[11] Ali I Ozdagli and Xenofon Koutsoukos. 2021. Model-based Damage Detection through Physics Guided Learning. In *Annual Conference of the PHM Society*, Vol. 13.

[12] V Papadopoulos, G Soimiris, DG Giovanis, and M Papadrakakis. 2018. A neural network-based surrogate model for carbon nanotubes with geometric nonlinearities. *Computer Methods in Applied Mechanics and Engineering* 328 (2018), 411–430.

[13] Nestor V Queipo, Raphael T Haftka, Wei Shyy, Tushar Goel, Rajkumar Vaidyanathan, and P Kevin Tucker. 2005. Surrogate-based analysis and optimization. *Progress in aerospace sciences* 41, 1 (2005), 1–28.

[14] DL Shrestha, N Kayastha, and DP Solomatine. 2009. A novel approach to parameter uncertainty analysis of hydrological models using neural networks. *Hydrology and Earth System Sciences* 13, 7 (2009), 1235–1248.

[15] Mahendra Pratap Singh and Bilas Chowdhury. 2011. *Control of autonomous underwater vehicles.* Ph. D. Dissertation. National Institute of Technology, Rourkela.

[16] J Sreekanth and Bithin Datta. 2010. Multi-objective management of saltwater intrusion in coastal aquifers using genetic programming and modular neural network based surrogate models. *Journal of hydrology* 393, 3-4 (2010), 245–256.

[17] Ruiyang Zhang, Yang Liu, and Hao Sun. 2020. Physics-guided convolutional neural network (PhyCNN) for data-driven seismic response modeling. *Engineering Structures* 215 (2020), 110704.

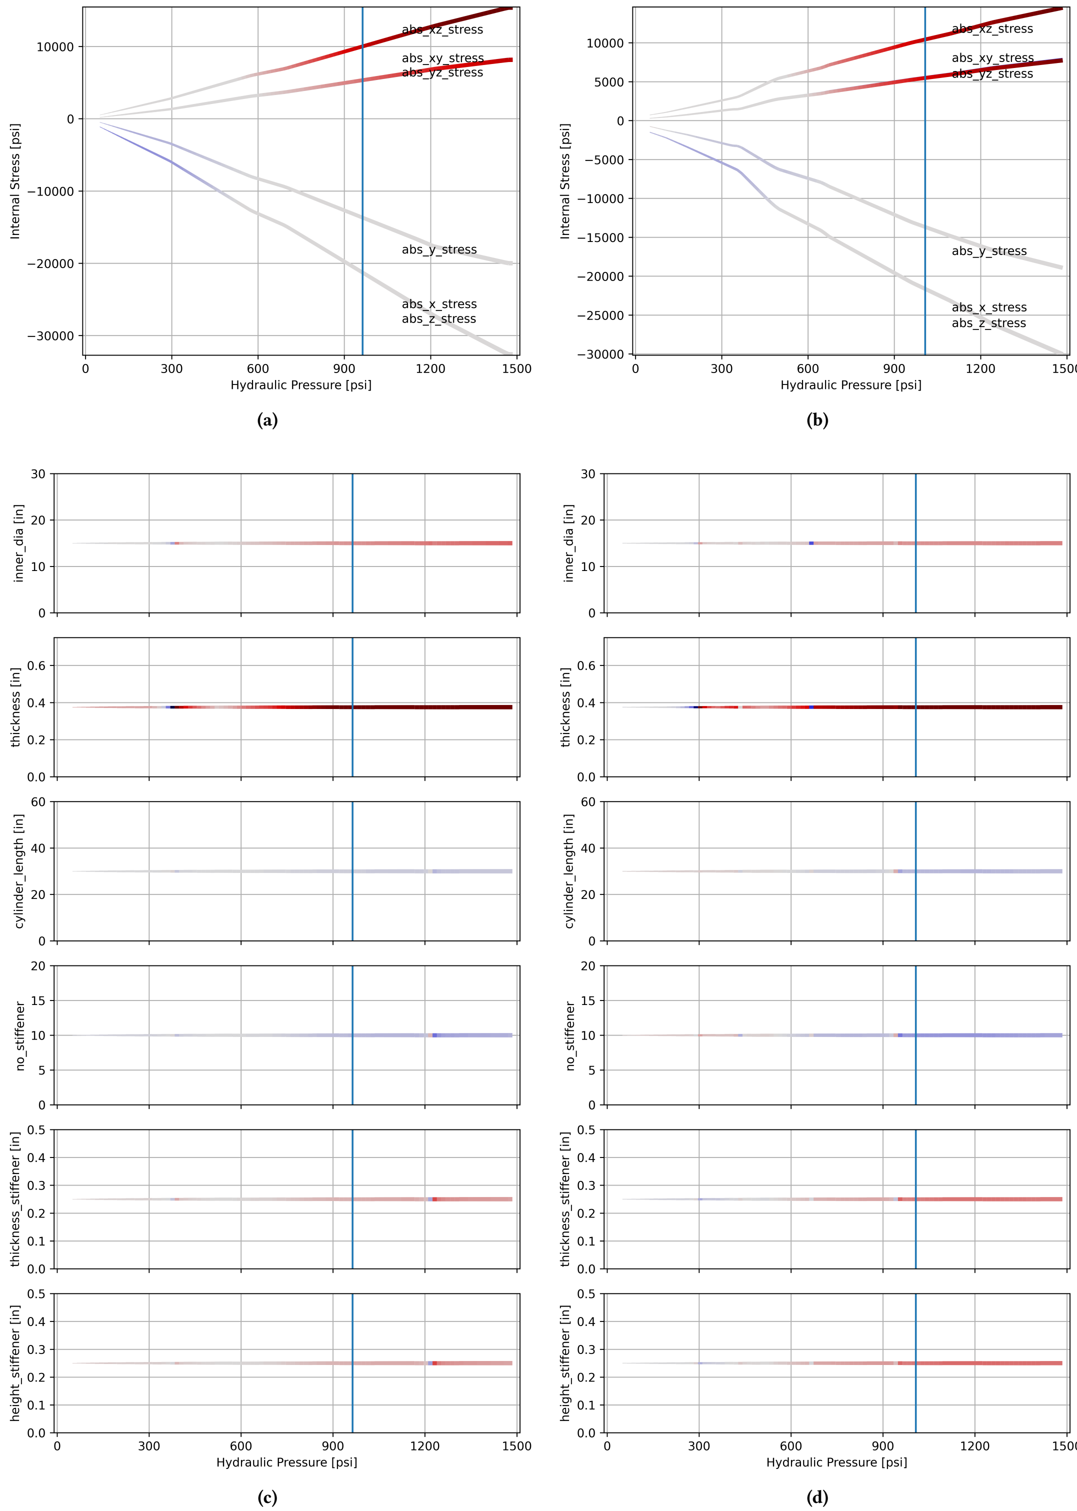

Figure 7: Relevancy scores for hull design with stiffeners: (a) intermediate physical parameters - inside stiffeners; (b) design parameters - inside stiffeners; (c) intermediate physical parameters - outside stiffeners; (d) design parameters - outside stiffeners