Supplementary Material

NeuberNet: a neural operator solving elastic-plastic PDEs at V-notches from low-fidelity elastic simulations

T. Grossi[1], M. Beghini[1], and M. Benedetti[2]

[1]Dipartimento di Ingegneria Civile e Industriale, Università di Pisa, Italy
[2]Dipartimento di Ingegneria Industriale, Università di Trento, Italy

S1 Additional results

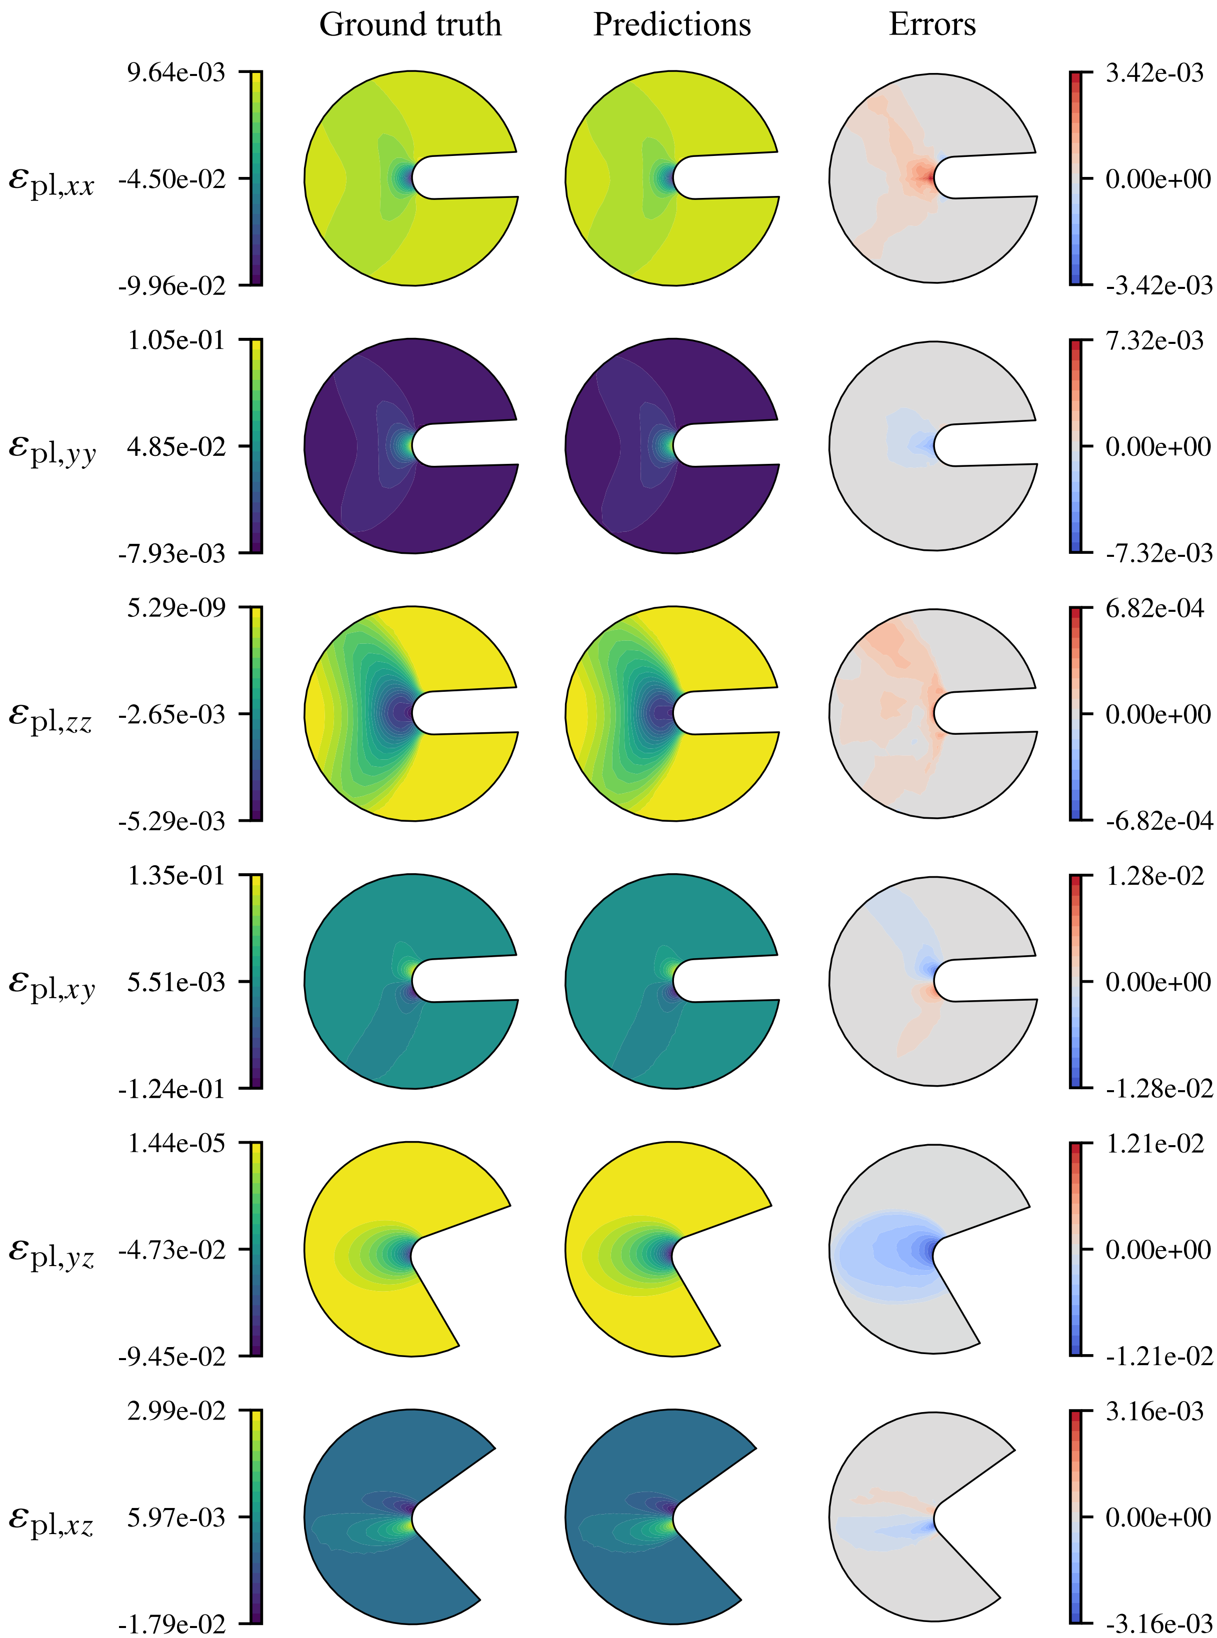

Figure S1: *Worst-case* performance on the test dataset in terms of maximum absolute error `Max(AE)` over the subdomain Ω for the six plastic strain components, expressed in the local reference system $Oxyz$ of Figure 3a.

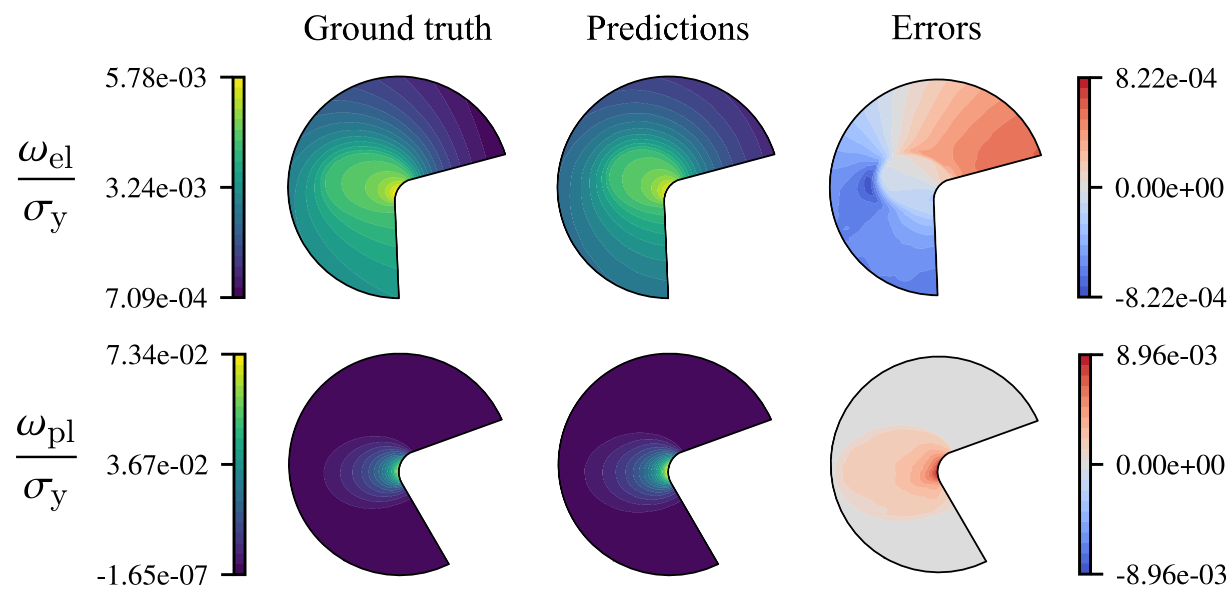

Figure S2: *Worst-case* performance on the test dataset in terms of maximum absolute error `Max(AE)` over the subdomain Ω for the normalized elastic strain energy density and the normalized plastic work per unit volume.

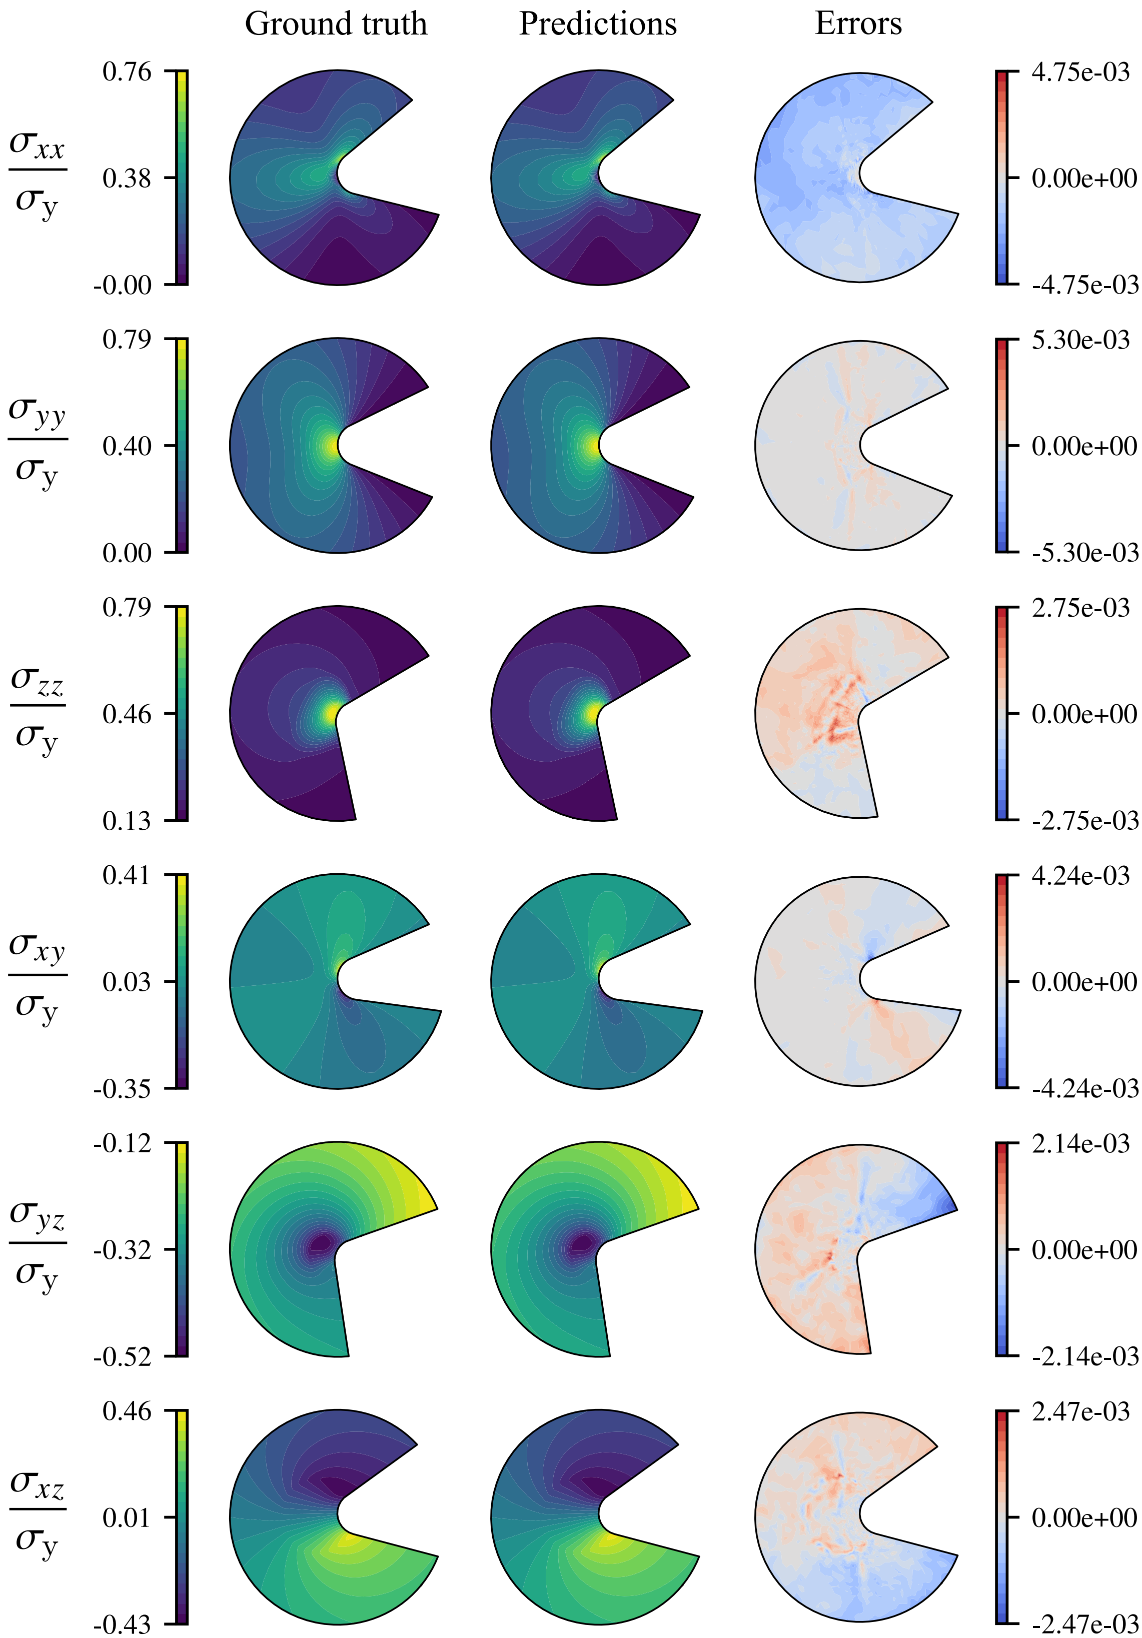

Figure S3: 50^{th} percentile performance on the test dataset in terms of maximum absolute error `Max(AE)` over the subdomain Ω for the six normalized stress components, expressed in the local reference system $Oxyz$ of Figure 3a.

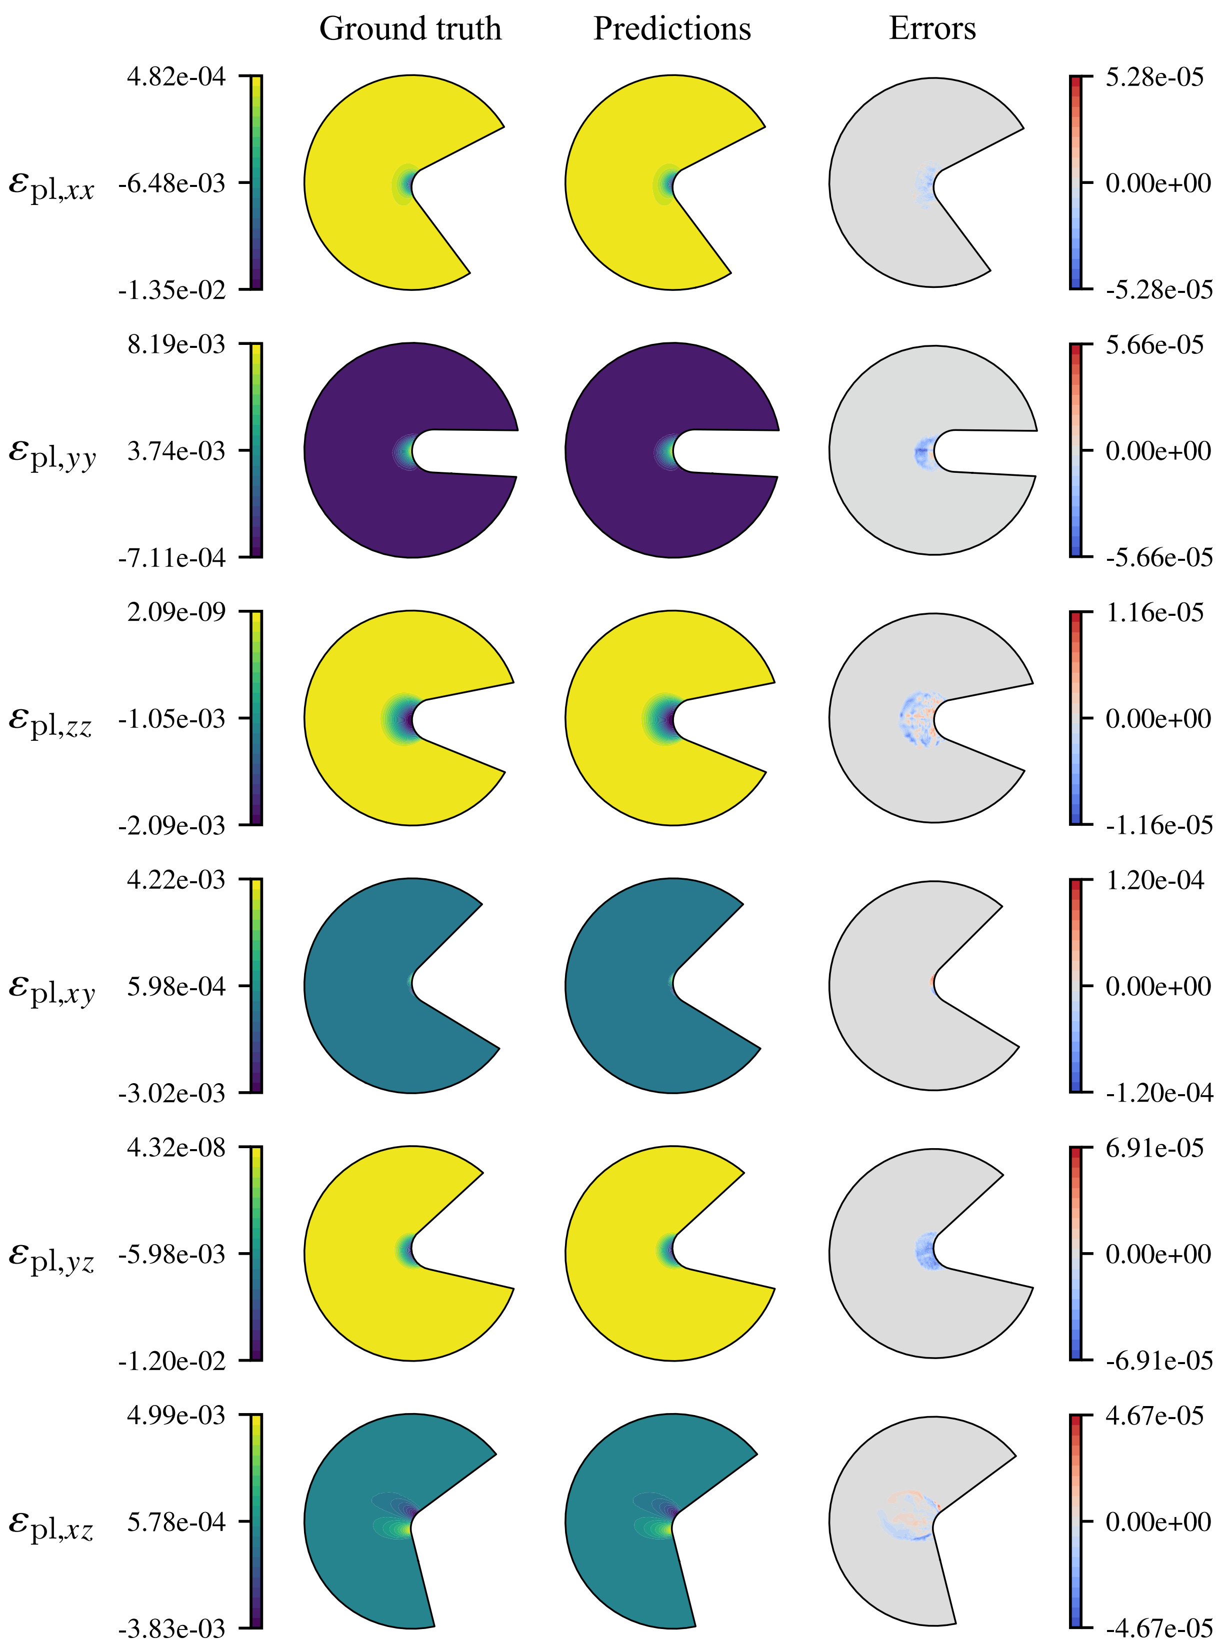

Figure S4: 50^{th} percentile performance on the test dataset in terms of maximum absolute error `Max(AE)` over the subdomain Ω for the six plastic strain components, expressed in the local reference system $Oxyz$ of Figure 3a.

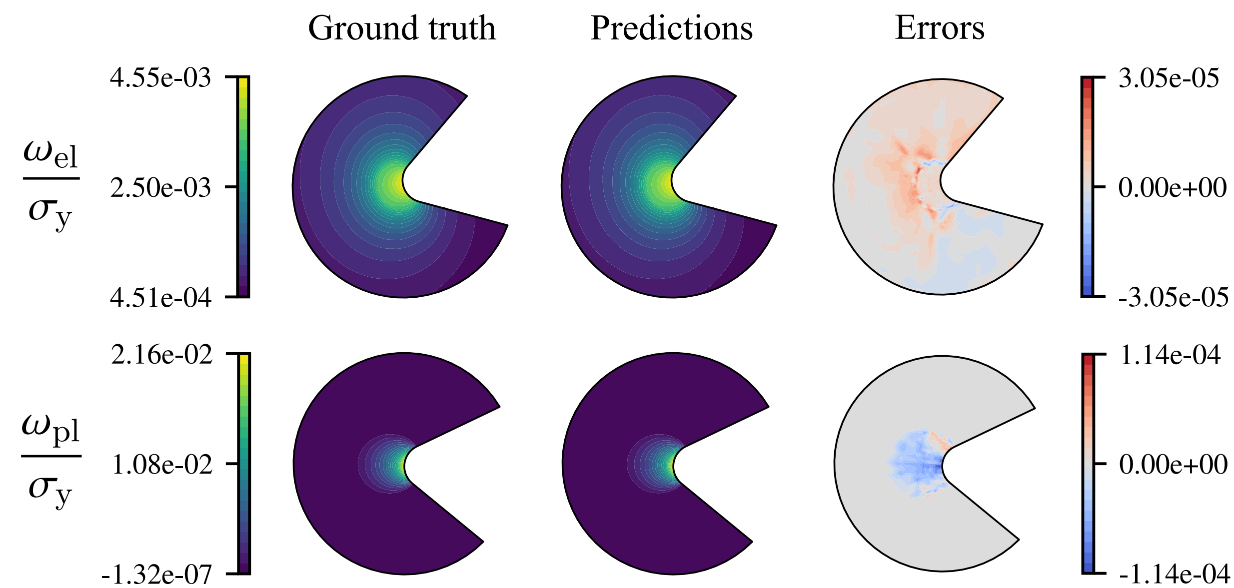

Figure S5: 50th percentile performance on the test dataset in terms of maximum absolute error `Max(AE)` over the subdomain Ω for the normalized elastic strain energy density and the normalized plastic work per unit volume.

S2 FE model details

Every geometry-material combination is uniquely identified by a tuple of scalar parameters, reported in Figure 3a and 3b, whose range of variation is reported in Table S1. This range has been designed to span an extensive library of structural engineering materials and practical use cases. Leveraging Buckingham's π-theorem, linear measurements and stress-like input quantities are normalized, respectively, over the notch radius R_n and over the Young modulus E, reducing the actual dimensions of the parameters space. The diameter of the reentrant corner domain is a design parameter, which has been fixed to 10 times the notch radius. For ease of interpretability of FE models, simulations are carried out with a notch radius of 1 mm and a Young modulus of 200 GPa, with all other variables following accordingly. Far-field loads are specified in terms of the ratios between the tension and torsion components, with their absolute values automatically adjusted in the simulations, as described below. Assuming tension-compression symmetry in the material's plastic behavior and neglecting large-deflection effects, the database can be augmented fourfold by alternatively reversing the signs of the in-plane and out-of-plane components; details are reported in Section S4. As the notch radius R_n becomes negligible compared to its distance R from the axis of cylindrical symmetry, the solution fields corresponding to axisymmetric solid mechanics converge to the plane strain solutions and no longer depend on the distance itself; hence, we crop R/R_n at 100.

The $N = 10^4$ d-tuples of scalar parameters are obtained by generating a $2^{\lceil \log_2 N \rceil}$ scrambled Sobol' sequence on $[0, 1]^d$, truncating it to N, and rescaling it to the desired range. This is done to ensure a low-discrepancy coverage of the parameter space. The tuples are written over separate `txt` files that will be imported at runtime by the APDL solver routine. As a result, the computational burden of FE analyses can be shared across multiple computing servers by simply partitioning the set of input files that are fed to the solver.

The FE models are meshed in ANSYS with 8-node quadratic PLANE183 elements in their axisymmetric formulation; to include torsional behavior, an additional DOF per node is added, representing the rotation about the symmetry axis. The reentrant corner subdomain is meshed using a mapped scheme with a randomly generated number of divisions along the domain boundary. This ensures that the training locations vary between simulations of similar notches and serves as a natural regularizer against model overfitting. The number of divisions is lower-bounded by preliminary convergence tests. The rest of the domain is free-meshed through the smart-sizing tool included in ANSYS. Plastic behavior is introduced with a bilinear isotropic hardening data table, activated for all the model elements. The lower side of the model is fully restrained; the upper side of the model is rigidly constrained with MPC184 elements, and the resulting master node is subjected to the external loads.

The simulation is carried out in steps. First, a linear elastic solution is pursued, and the corresponding stress fields are obtained. Then, the external load is scaled to match the onset of yield in the reentrant corner subdomain, according to a Von Mises stress criterion, and the corresponding displacements fields $\boldsymbol{u}$ at the boundary of the reentrant corner subdomain are stored. Let us call the corresponding external loads vector as $\boldsymbol{f}$. From this point onward, non-linear material behavior is activated. External loads are increased in steps of $0.25\boldsymbol{f}$, and a new solution is searched with an iterative Newton-Raphson solver. Convergence is typically achieved in just a few iterations, given the moderate change of boundary conditions. Once the solution is attained, it is checked whether the plastic zone has reached the boundary of the reentrant corner subdomain. If the plastic zone has reached the boundary, the simulation is terminated; otherwise, the stress, plastic strain, and energy density fields at corner node locations are stored, and another load increment is applied. Tensorial quantities are expressed with respect to the $Oxyz$ reference system in Figure 3a. Since the neural operator is trained to learn the functional mapping between the elastic boundary conditions and the plastic fields, the former are obtained by linearly scaling the ones stored in the first step.

Table S1: Range of normalized geometric and constitutive parameters

	Geometric parameters			Material parameters		
	R/R_n	α	β	ν	σ_y/E	E_t/E
max	100	75°	50°	0.45	10^{-2}	10^{-1}
min	10	0°	−50°	0.05	10^{-3}	10^{-3}

S3 NeuberNet loss histories

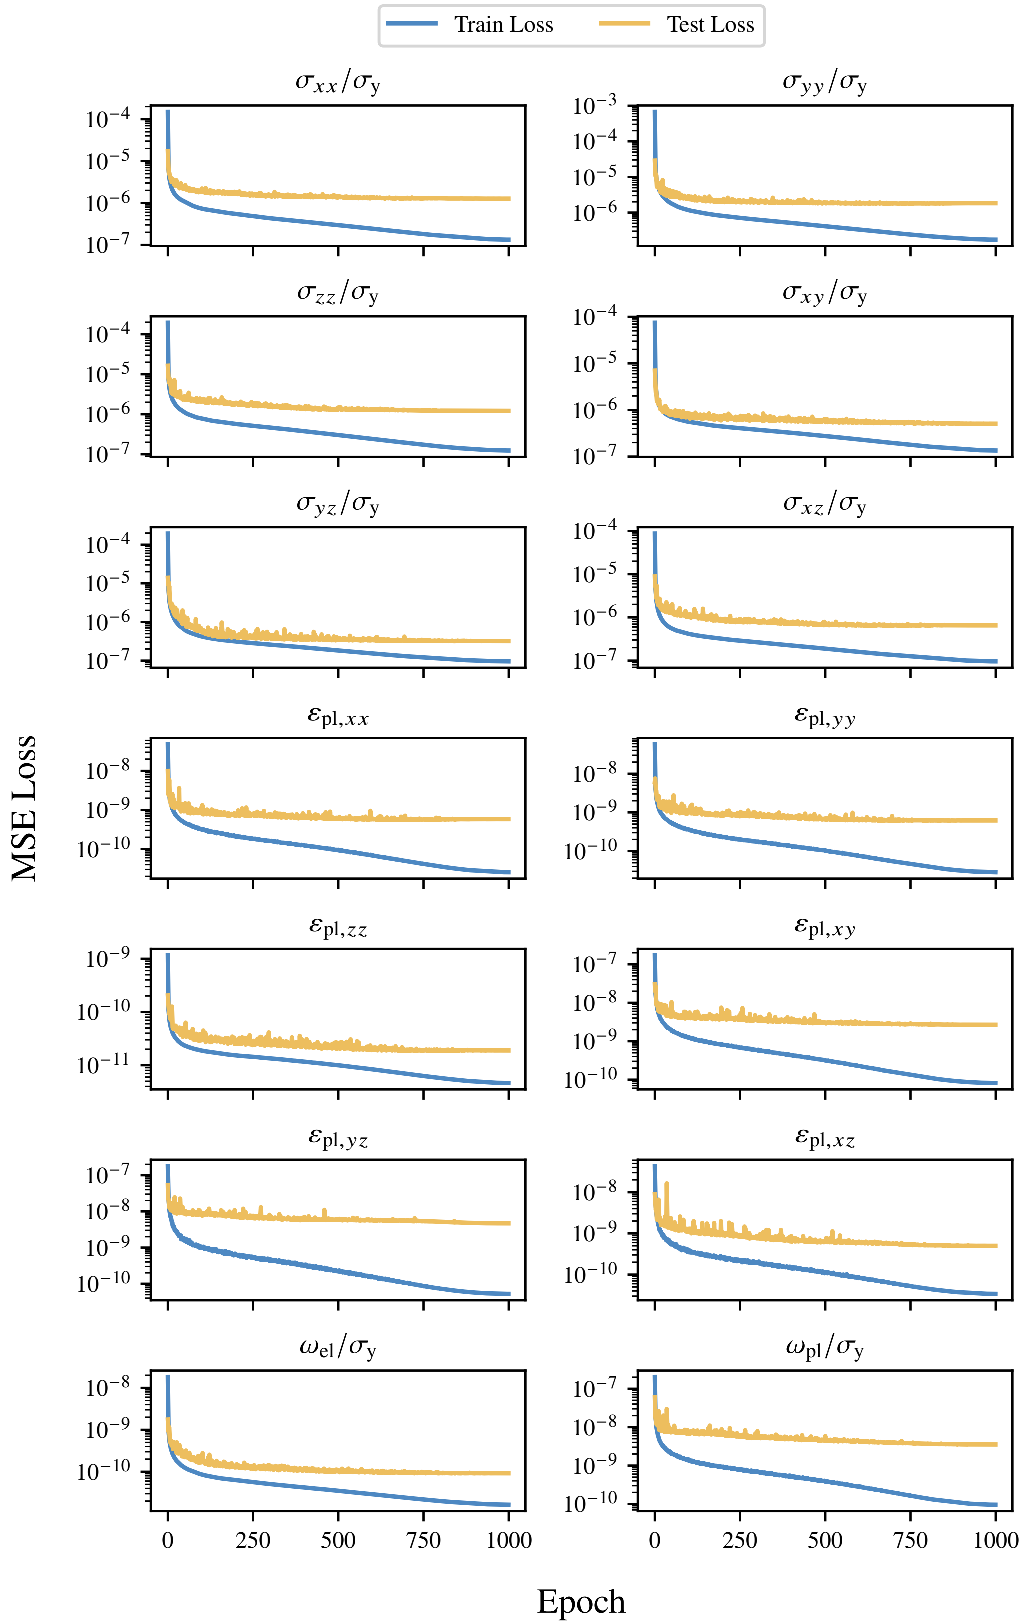

Figure S6: Train and test MSE losses observed for each target variable during NeuberNet training. Losses are reported in their true scale, although training is conducted with unit-variance data.

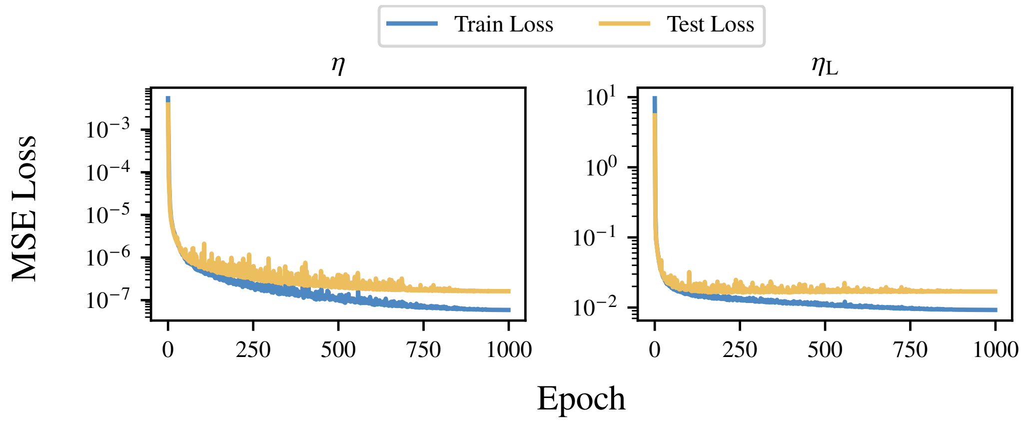

Figure S7: Train and test MSE losses observed for each target variable during YieldNet training. Losses are reported in their true scale, although training is conducted with unit-variance data.

S4 YieldNet training

Starting from the FE dataset, we take the 10^4 input boundary displacements $\boldsymbol{u}$ at yield onset and augment the dataset by randomly scaling them by a number between 0 and 2 times the maximum factor at which the small-scale plasticity condition has been violated. Then, we train a neural network—referred to as YieldNet—that takes $\boldsymbol{u}$ and the geometric/constitutive parameters as inputs and returns two outputs: 1) the ratio $\eta = \max_\Omega \sigma_{\mathrm{vm}}/\sigma_{\mathrm{y}}$ between the maximum Von Mises stress inside the reentrant corner subdomain Ω and the material yield stress; 2) the maximum factor η_{L} by which the boundary displacements corresponding to the yield limit can be scaled without violating the small-scale plasticity hypothesis, according to the FE dataset. For ease of implementation and given the relative simplicity of this task, we keep the same NOMAD architecture, while we remove the two spatial coordinates (x, y) from the decoder inputs. We train YieldNet for 1000 epochs with an `AdamW` optimizer, using a batch size of 128, an initial learning rate of 10^{-4}, a cosine annealing scheduler and a weight decay of 10^{-4}. Loss histories are available in Figure S7. Recall that the load step size in FE simulations is 0.25 times the yielding displacements, so η_{L} is known up to a random variable that is uniformly distributed over $[-0.25, 0]$. As a result, the MSE loss in predicting η_{L} is theoretically bounded by $0.25^2/3 \approx 0.02$, which is closely matched by YieldNet.

As mentioned in Section S2, the training dataset can be augmented fourfold by reversing the signs of in-plane and out-of-plane displacements and accordingly adjusting the signs of the respective tensorial target variables. However, instead of increasing the dataset size—which would result in longer training times—we opt to include only analyses with predominantly tensile in-plane loading and counterclockwise out-of-plane loading. To ensure consistency with the training dataset, we use a neural network trained to minimize a Binary Cross-Entropy (BCE) loss to classify each input into one of two binary categories: tension/compression and clockwise/counterclockwise. Based on this classification, we adjust the input signs of $\boldsymbol{u}$ accordingly before inference and adapt the solutions post-inference. This step is entirely optional and can be omitted in favor of explicitly augmenting the FE database.