










# Machine Learned Interatomic Potentials for Ternary Carbides trained on the AFLOW Database

Josiah Roberts <sup>†</sup>, Biswas Rijal <sup>†</sup>, Simon Divilov <sup>‡</sup>, Jon-Paul Maria <sup>§</sup>,  
William G. Fahrenholz <sup>||</sup>, Douglas E. Wolfe <sup>§</sup>, Donald W. Brenner <sup>⊥</sup>,  
Stefano Curtarolo <sup>‡</sup> and Eva Zurek <sup>\*,†</sup>

<sup>†</sup>*Department of Chemistry, State University of New York at Buffalo, Buffalo, NY 14260, USA*

<sup>‡</sup>*Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA*

<sup>¶</sup>*Center for Autonomous Materials Design, Duke University, Durham, NC 27708, USA*

<sup>§</sup>*Department of Materials Science and Engineering, The Pennsylvania State University, University Park, PA 16802, USA*

<sup>||</sup>*Department of Materials Science and Engineering, Missouri University of Science and Technology, Rolla, MO 65409, USA*

<sup>⊥</sup>*Department of Materials Science and Engineering, North Carolina State University, Raleigh, NC 27695, USA*

E-mail: ezurek@buffalo.edu

## Contents

<b>S1 Plan for Robust and Accurate Potentials (PRAPs) Training</b>	<b>S2</b>
<b>S2 Training Data and Root Mean Square Errors</b>	<b>S2</b>
<b>S3 Convex Hulls</b>	<b>S7</b>

# S1 Plan for Robust and Accurate Potentials (PRAPs) Training

Table S1: The number of low energy configurations correctly and incorrectly predicted (A-B) and the number of high energy configurations correctly and incorrectly predicted (C-D) by PRAPs (denoted as A-B/C-D) for the binary carbide systems studied at different MTP levels of theory. The Pre-Robust Potential (Pre-RP) and the Robust Potential (RP) are compared against the original AFLOW data (after removing structures with unphysically short interatomic distances). The Accurate Potential (AP) is used to make two predictions and two comparisons. First it is compared against the low-energy robust-relaxed data (AP-RR), and then compared against the low-energy AFLOW data (AP-Low).

Level 10	Pre-RP	RP	AP-RR	AP-Low
CHfTa	6-0/7-0	6-0/9-0	6-0/8-0	6-0/10-0
CHfZr	3-3/7-0	2-0/9-0	2-0/9-0	2-0/9-0
CMoW	8-0/8-0	6-0/9-0	3-0/10-0	6-0/3-1
CTaTi	7-0/7-0	4-0/4-0	4-0/8-0	4-0/7-0
Level 16	Pre-RP	RP	AP-RR	AP-Low
CHfTa	6-0/8-0	6-0/10-0	6-0/8-0	6-0/10-0
CHfZr	3-2/7-0	6-0/10-0	2-0/9-0	6-0/9-0
CMoW	8-0/4-0	7-0/10-0	7-0/8-0	7-0/1-1
CTaTi	8-0/6-0	7-0/4-0	6-0/9-0	7-0/8-0
Level 22	Pre-RP	RP	AP-RR	AP-Low
CHfTa	3-2/6-0	3-0/9-0	6-0/10-0	3-0/10-0
CHfZr	1-2/6-0	2-0/9-0	2-0/7-0	2-0/9-0
CTaTi	3-3/6-0	4-0/7-0	5-0/9-0	4-0/9-0

# S2 Training Data and Root Mean Square Errors

To better understand the factors contributing to the root mean square errors (RMSE), we plotted the DFT energies of the structures present in the AFLOW training set versus an arbitrarily chosen structure index in Figure S1. Configurations with interatomic distances outside of the range  $1.1 \text{ \AA} < x < 3.1 \text{ \AA}$ , were automatically filtered from the training set. While a majority of the configurations present possessed energies that fell within  $-9.5 \pm 1.0 \text{ eV/atom}$ , 5.6% fell outside this range, some by up to -40 meV/atom, and 30% fell outside of one standard deviation from the mean.

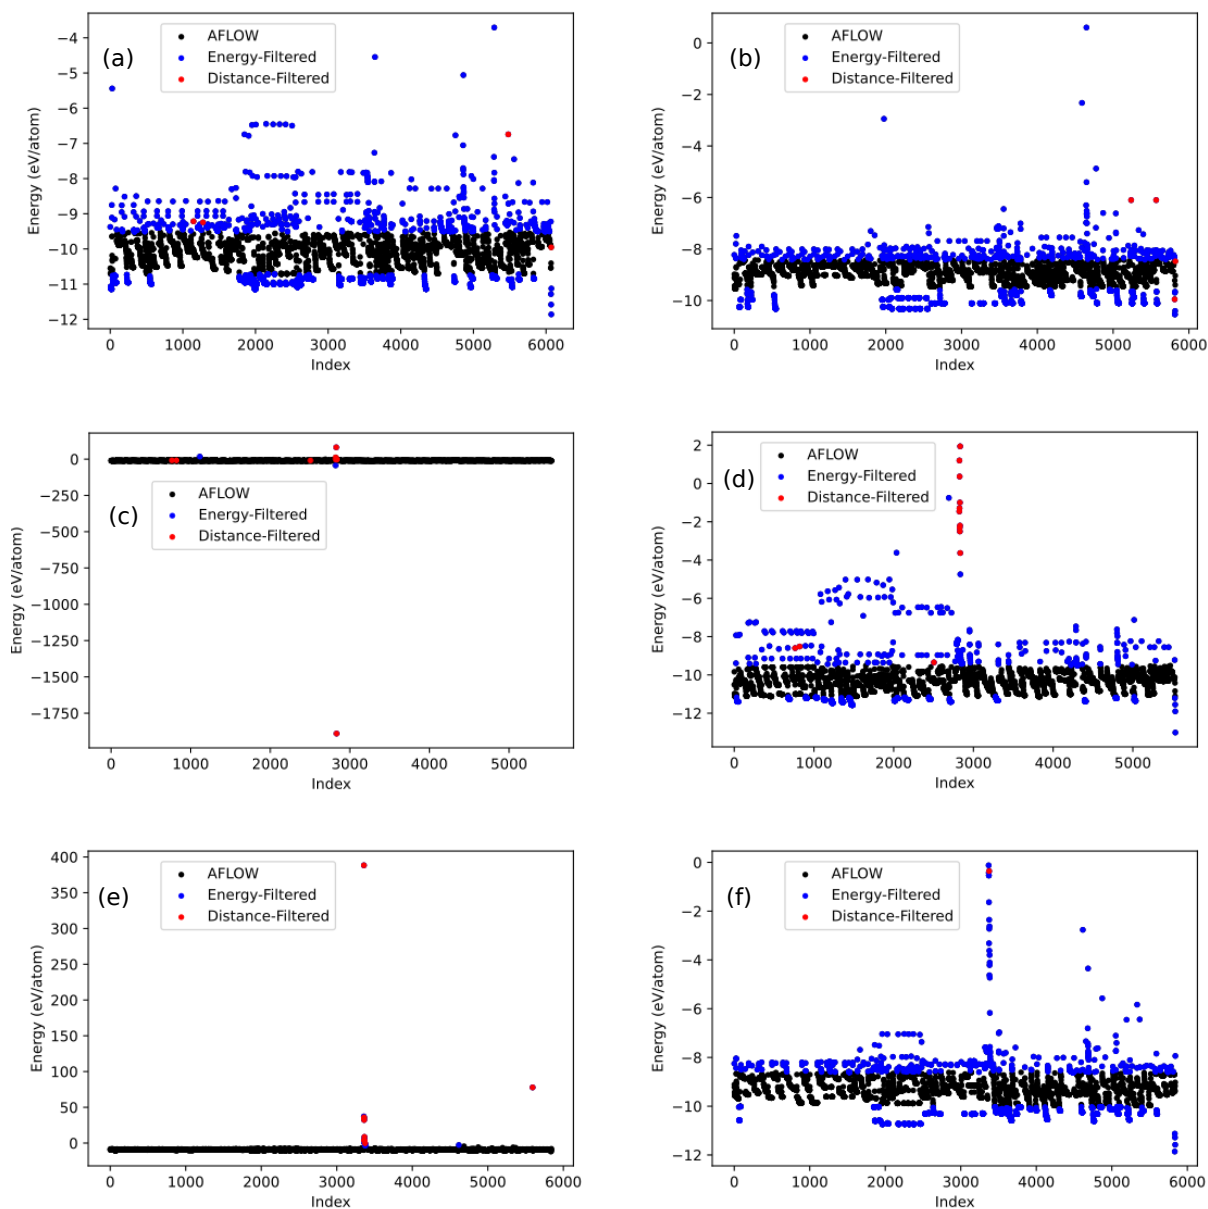


Figure S1: Energy values found in the original AFLOW data for: (a) CHfTa, (b) CHfZr, (c and d) CMoW raw and zoomed-in, and (e and f) CTaTi raw and zoomed-in. Configurations not removed by filtering are marked in black. Configurations removed by the distance filtering are marked in red. Configurations removed during testing of energy filtering are marked in blue.

PRAPs always filters structures by distance, but we considered the possibility of adding energy as a secondary filtering condition. After filtering by distance the remaining filtrate was filtered by only including configurations with energies within one standard deviation from the mean. A new robust potential (Level 10e) was generated from this filtered data and plots were generated to investigate the effect of the filtering on the energy predictions, appearing below in Figure S2.

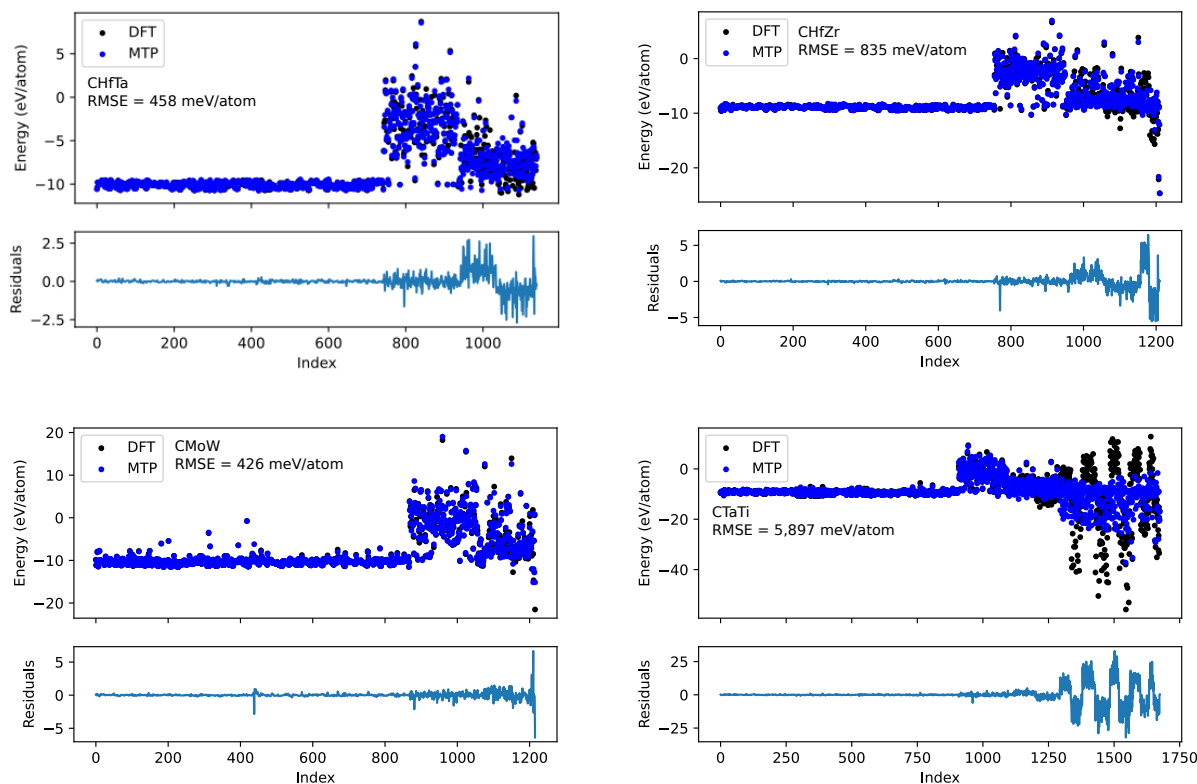


Figure S2: The AFLOW data (Figure S1, all dots) was filtered by both distance and energy (Figure S1, black dots) and submitted for training. Here we show the energies of the Robust Potential training set, as calculated by DFT and the Robust Potential (MTP, Level 10e). The training set contains two regions: the relatively flat region in the first half of the x-axis contains configurations straight from the AFLOW data found in Figure S1, the other region in the second half of the x-axis with the widely spread data contains configurations generated during the active learning procedure, which may include configurations that are not chemically sensible or converged, reflecting the quality of the generating MTP. The associated root-mean-square-errors appear in each plot (compare to 197, 302, 448, and 704 meV/atom, respectively, from the main text, Table II) and the differences between the DFT and Robust Potential appear below as residual traces.

To gain a better understanding of the predictive power of our MTPs and the origin of the previously-stated RMSEs, we examined the differences between the original AFLOW data and the trained potential’s prediction. In Table S2, we counted how many configurations had differences within one standard deviation of the mean and how many configurations had differences less-than 0.5 eV/atom. This was done for each system, for both Robust and Accurate Potentials, using the standard Level 10 data (see main text) and the energy-filtered data from above (Level 10e).

Table S2: Percent of the training set configurations whose predicted energies differed by less-than one standard deviation of the mean and less-than 0.5 eV/atom from the original AFLOW data.

Level 10	RP		AP	
	Stdev	< 0.5	Stdev	< 0.5
CHfTa	77.9%	84.3%	87.4%	99.6%
CHfZr	81.6%	80.5%	90.9%	99.3%
CMoW	81.3%	79.1%	86.4%	98.9%
CTaTi	85.7%	81.1%	87.3%	99.7%
Level 10e	RP		AP	
	Stdev	< 0.5	Stdev	< 0.5
CHfTa	77.2%	75.3%	87.0%	98.7%
CHfZr	79.2%	73.4%	89.9%	99.9%
CMoW	79.7%	78.5%	91.2%	97.7%
CTaTi	78.3%	55.4%	92.2%	72.4%

We performed a set of tests on each of the four systems at level 10 to test the effects of structural diversity. In these, we replaced the set of 6000 RandSPG structures of diverse space groups, with a set of 200 HCP, FCC, and BCC structures from RandSPG. The energy RMS values are reported in meV/atom in Table S3 with the original values from Table II of the main text in parentheses for comparison.

Table S3: Root mean square errors of the energy for training and prediction across the four systems where the lattices of the RandSPG structures were restricted to HCP, FCC, and BCC only. Original values from Table II of the main text are in parentheses.

System	Training Errors (meV/atom)			Prediction Errors (meV/atom)	
	Pre-trained	RP	AP	RP	AP
CHfTa	38 (34)	186 (197)	162 (179)	77 (77)	53 (70)
CHfZr	32 (31)	244 (302)	95 (136)	76 (92)	42 (89)
CMoW	66 (69)	458 (448)	152 (187)	489 (490)	80 (142)
CTaTi	43 (41)	803 (704)	279 (167)	251 (237)	150 (90)

Finally, we performed a small test at level 10 across all four systems to examine the usefulness of the pre-training procedure. The entire PRAPs training procedure was conducted, omitting the pre-training step, and the timings and errors were compared with those obtained from the original tests as described in the main text. The final DFT calculations and convex hull analysis were omitted as that time can be separated from the training time. The results are provided in Table S4. Omitting the pre-training gave significant improvements in speed at the cost of MTP quality, as measured by energy RMS error (values in parentheses are from the original run, Table II in the main text). Wall times (D-HH:MM:SS) were provided instead of CPU times for a more faithful estimate of run time.

Table S4: Comparison of PRAPs performance with and without the pre-training step. Times are given as Days-Hours:Minutes:Seconds. Error values are reported in meV/atom. The parentheses contain the errors for the control group at Level 10 with pre-training as reported in the main text, Table II.

System	Wall Time		Training Errors (meV/atom)		Prediction Errors (meV/atom)	
	Original	No Pre-train	RP	AP	RP	AP
CHfTa	7-20:50:57	6-00:08:55	622 (197)	109 (179)	183 (77)	105 (70)
CHfZr	8-10:27:47	8-02:14:44	613 (302)	94 (136)	160 (92)	136 (89)
CMoW	7-12:24:07	4-21:53:13	760 (448)	201 (187)	504 (490)	434 (142)
CTaTi	12-05:49:09	6-02:22:11	1104 (704)	296 (167)	329 (237)	188 (90)

## S3 Convex Hulls

This section contains a subset of the convex hulls calculated by PRAPs. If desired by the user, PRAPs will calculate eleven convex hulls: two for the DFT data before-and-after relaxation, and nine using the Robust and Accurate Potentials. Six of the hulls are presented here in a table-like format, organized into rows and columns from Figures S3 to S13. The AFLOW data (top row) is combined with the RANDSPG data and filtered to remove unphysical structures as described in the main text. This filtered set is then relaxed using the Robust Potential to form the Robust-Relaxed set (RR, second row). Then the RR set is filtered to obtain only configurations within 50 meV/atom of the lowest energy configuration for each composition and the Accurate Potential is used to relax these structures (AP-RR, bottom row). The Accurate Potential is also used on the AFLOW data to determine if it can reproduce the AFLOW convex hull (AP-v, not shown). The left column contains hulls generated from the MLIP output while the right shows the hulls of the same data after DFT relaxation via AFLOW’s standard procedure. To calculate the remaining hulls, PRAPs will concatenate two files together: the Post-Relaxation AFLOW convex hull (shown below in the top-right of each figure) and the Post-Relaxation Accurate Potential hulls (shown below in the middle-right and lower-right of each figure). This allows a direct comparison of the “literature” and “experiment” and will tell the user whether or not structures predicted by PRAPs are still on the convex hull when considering the original AFLOW data. These hulls are not shown below as almost all of the hulls remain unchanged from those already shown.

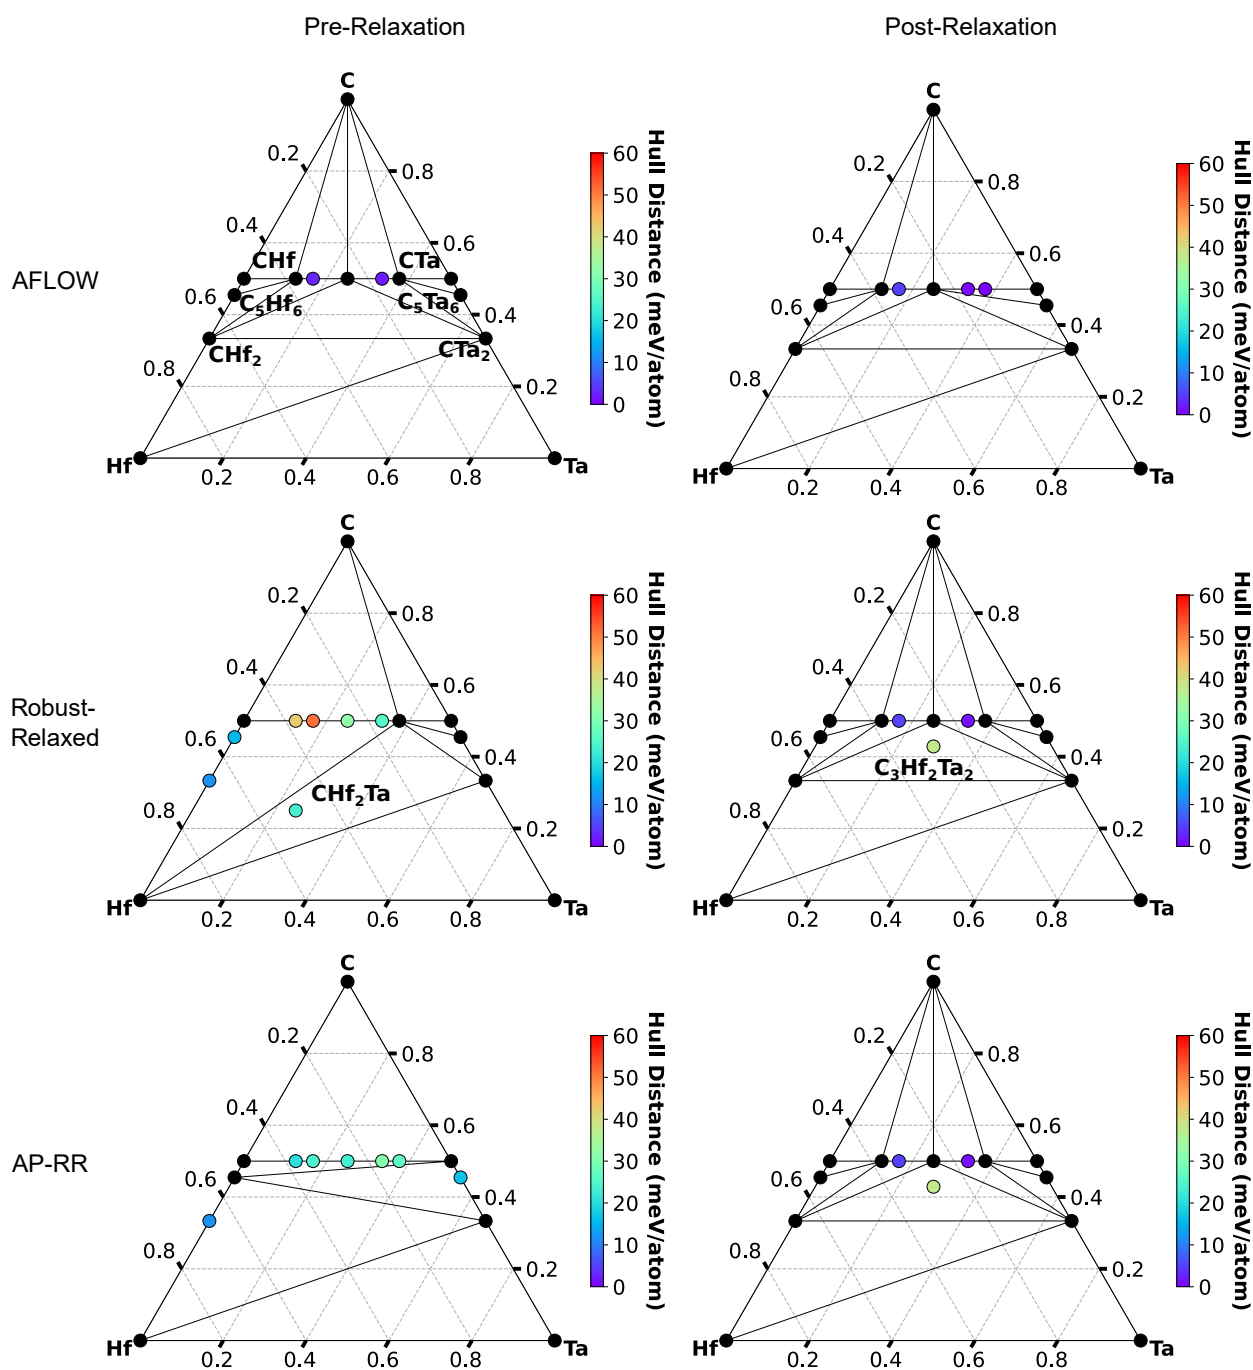


Figure S3: Convex Hulls of CHfTa at Level 10. The unlabeled structures on the line between CHf and CTa are, from left-to-right, C<sub>4</sub>Hf<sub>3</sub>Ta, C<sub>3</sub>Hf<sub>2</sub>Ta, C<sub>2</sub>HfTa, C<sub>3</sub>HfTa<sub>2</sub>, and C<sub>4</sub>HfTa<sub>3</sub>.

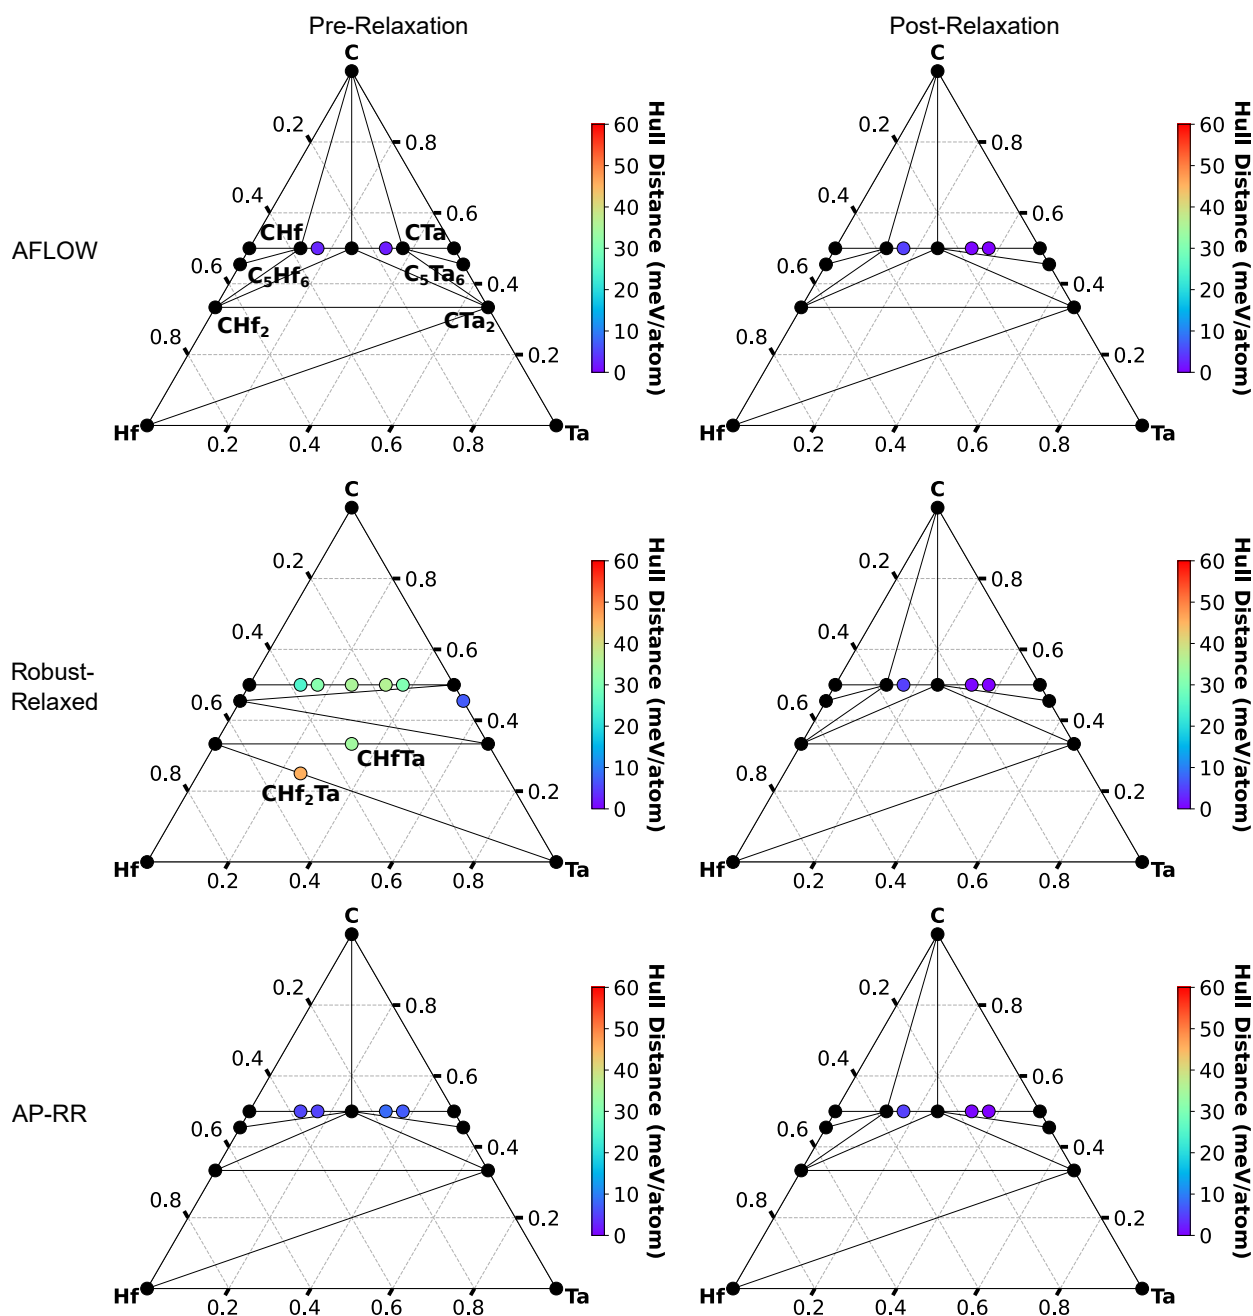


Figure S4: Convex Hulls of CHfTa at Level 16. The unlabeled structures on the line between CHf and CTa are, from left-to-right,  $C_4Hf_3Ta$ ,  $C_3Hf_2Ta$ ,  $C_2HfTa$ ,  $C_3HfTa_2$ , and  $C_4HfTa_3$ .

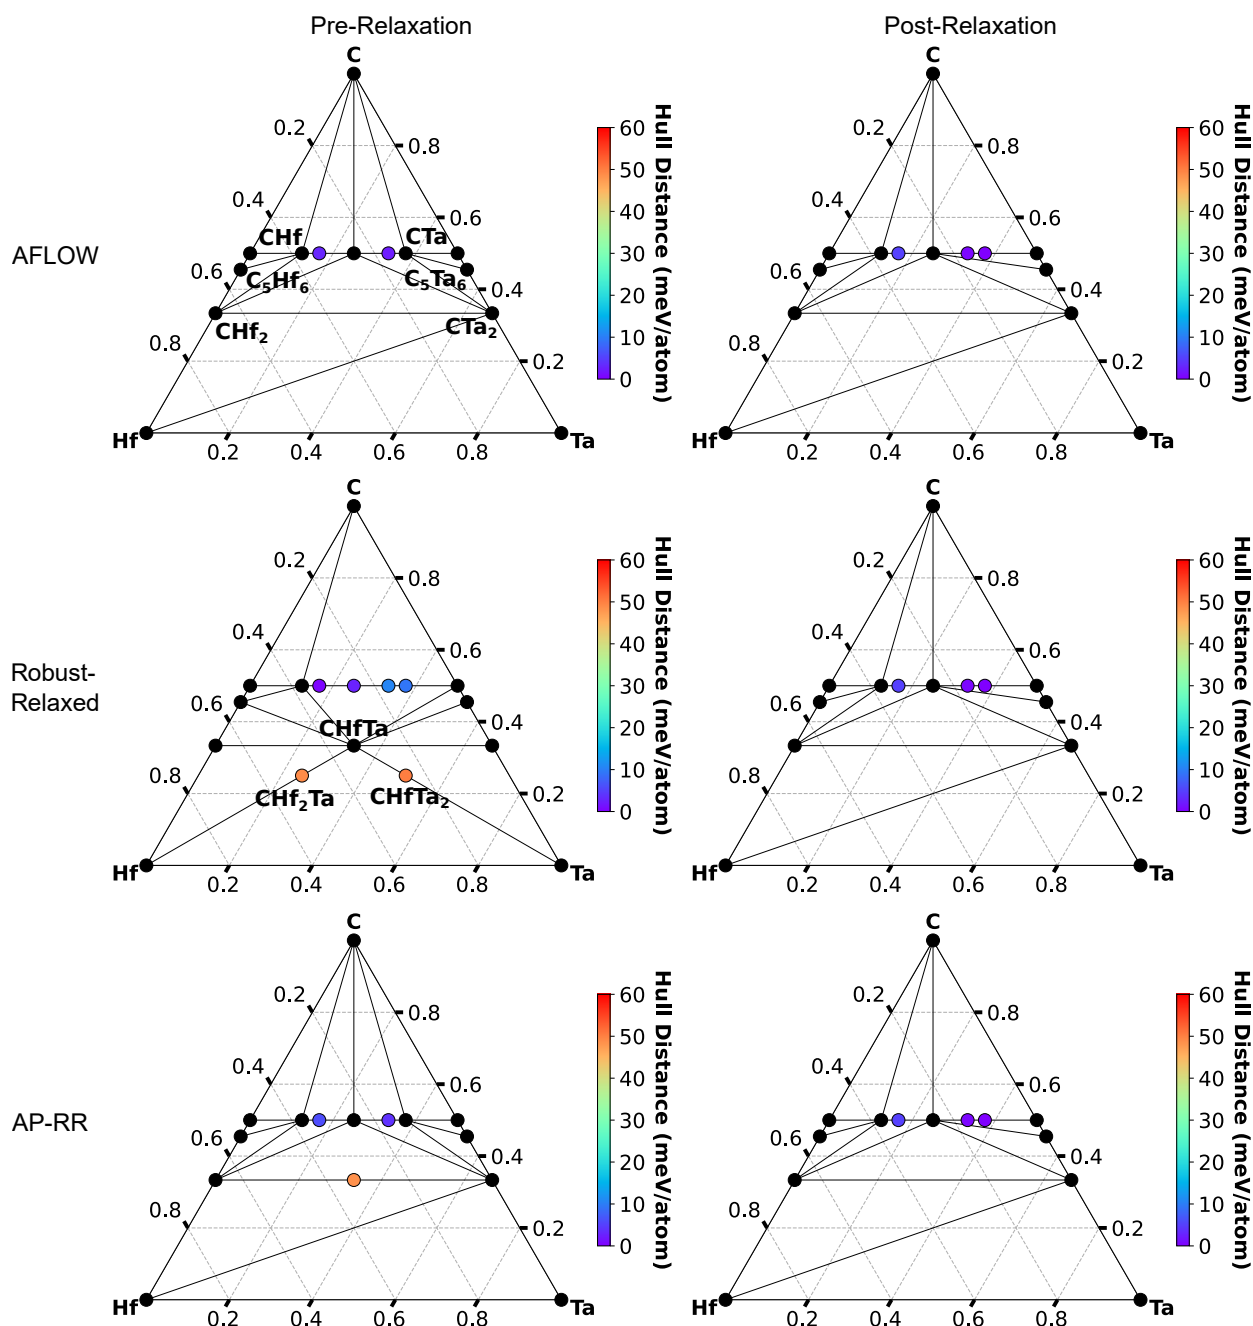


Figure S5: Convex Hulls of CHfTa at Level 22. The unlabeled structures on the line between CHf and CTa are, from left-to-right,  $C_4Hf_3Ta$ ,  $C_3Hf_2Ta$ ,  $C_2HfTa$ ,  $C_3HfTa_2$ , and  $C_4HfTa_3$ .

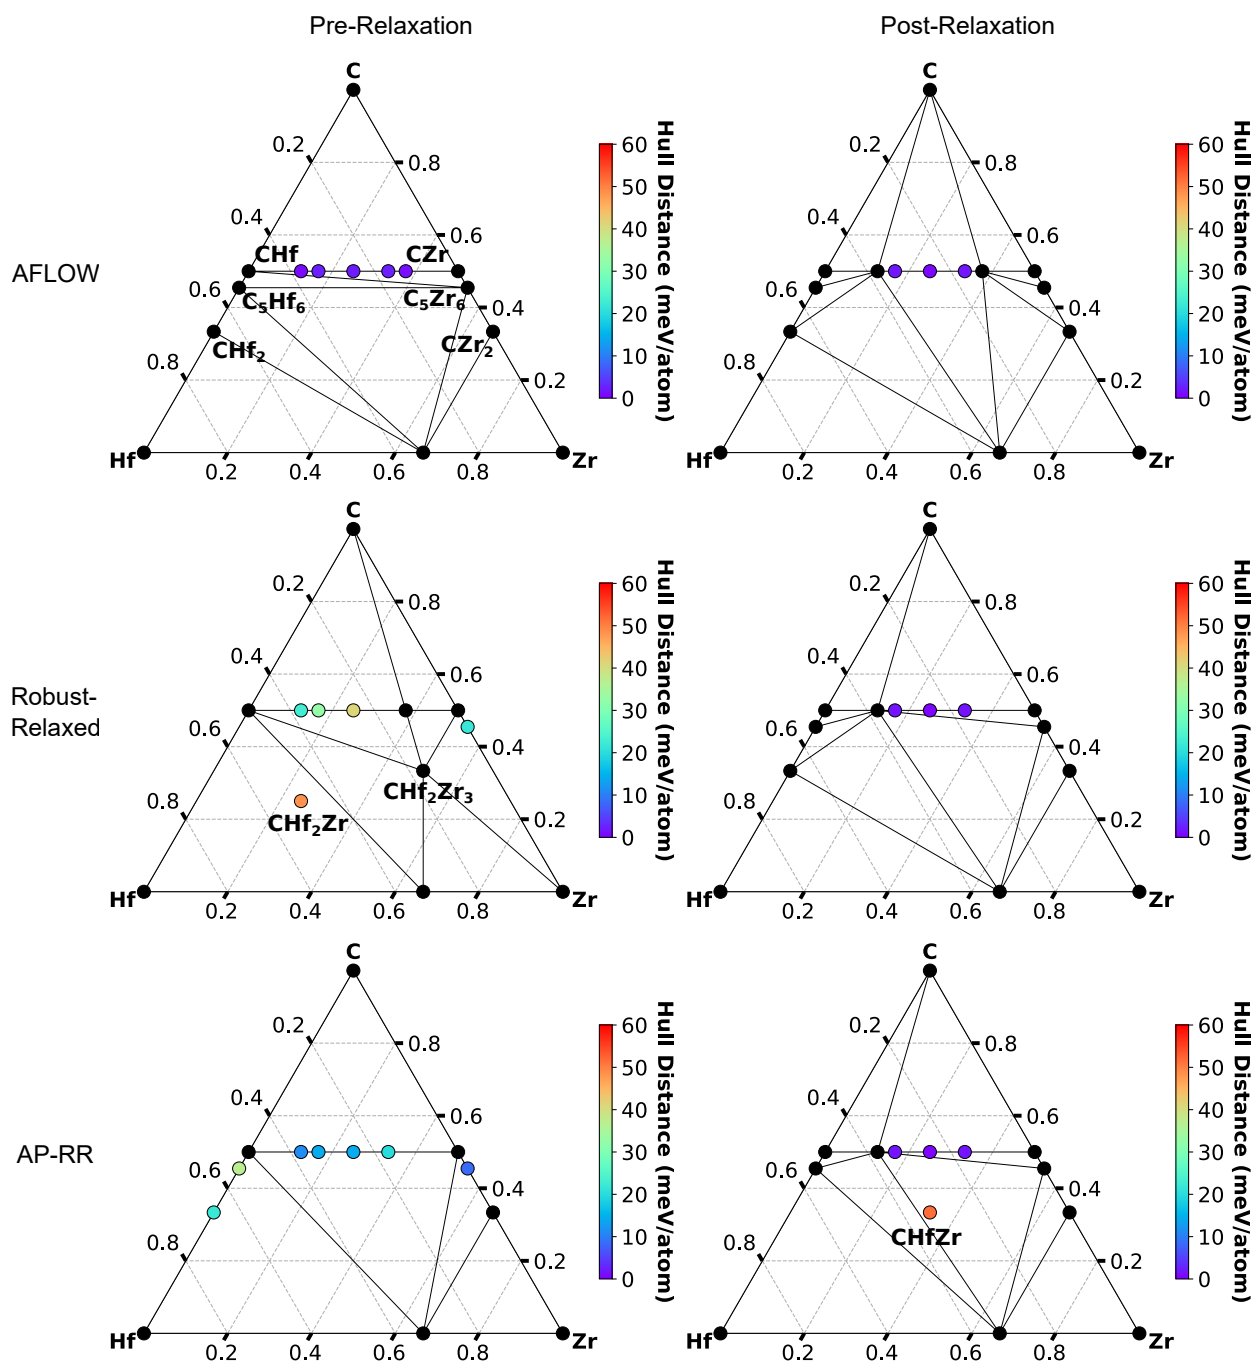


Figure S6: Convex Hulls of CHfZr at Level 10. The unlabeled structures on the line between CHf and CZr are, from left-to-right,  $C_4Hf_3Zr$ ,  $C_3Hf_2Zr$ ,  $C_2HfZr$ ,  $C_3HfZr_2$ , and  $C_4HfZr_3$ .

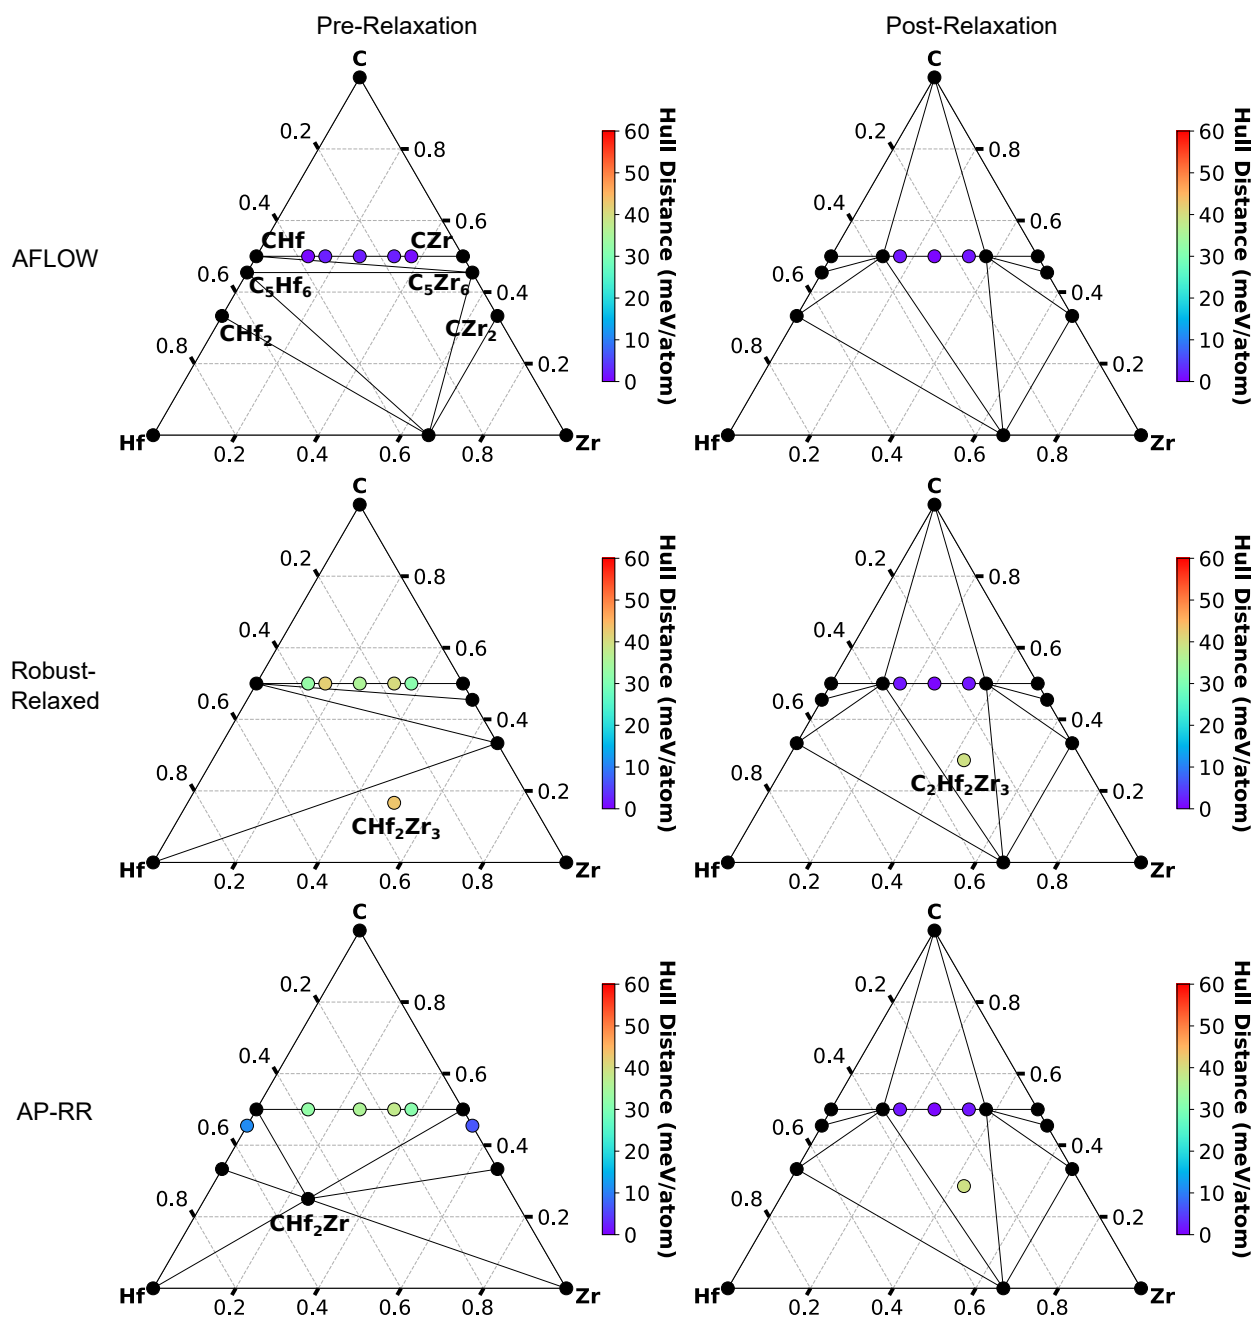


Figure S7: Convex Hulls of CHfZr at Level 16. The unlabeled structures on the line between CHf and CZr are, from left-to-right,  $C_4Hf_3Zr$ ,  $C_3Hf_2Zr$ ,  $C_2HfZr$ ,  $C_3HfZr_2$ , and  $C_4HfZr_3$ .

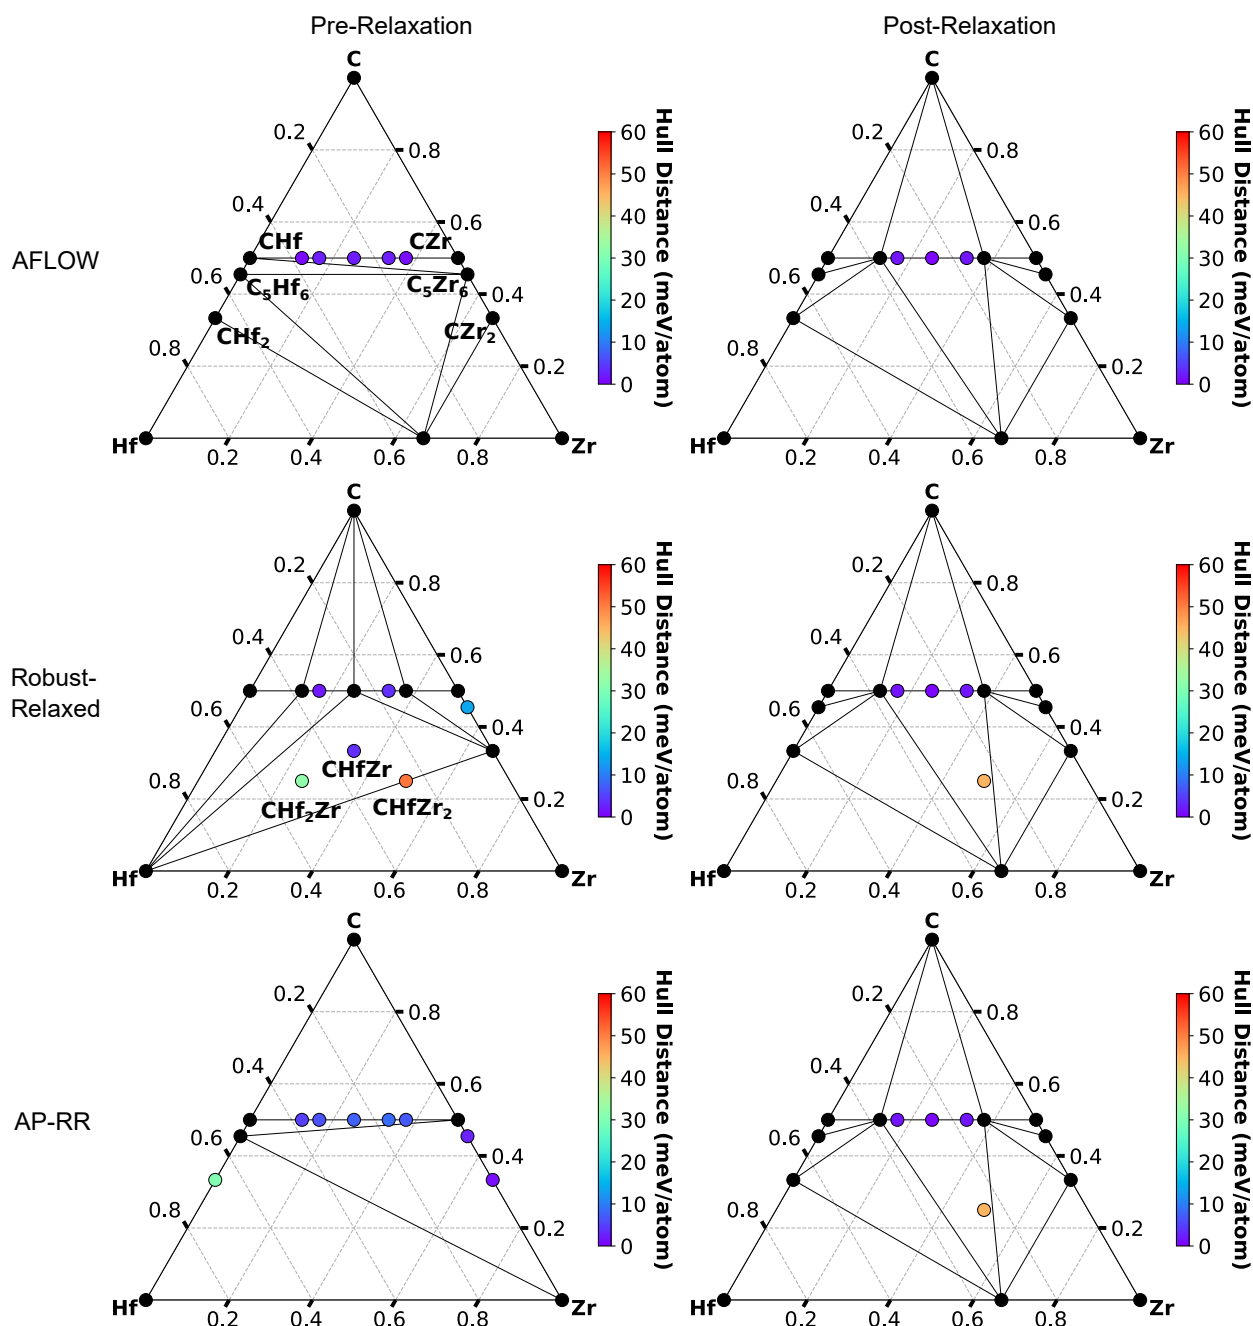


Figure S8: Convex Hulls of CHfZr at Level 22. The unlabeled structures on the line between CHf and CZr are, from left-to-right,  $C_4Hf_3Zr$ ,  $C_3Hf_2Zr$ ,  $C_2HfZr$ ,  $C_3HfZr_2$ , and  $C_4HfZr_3$ .

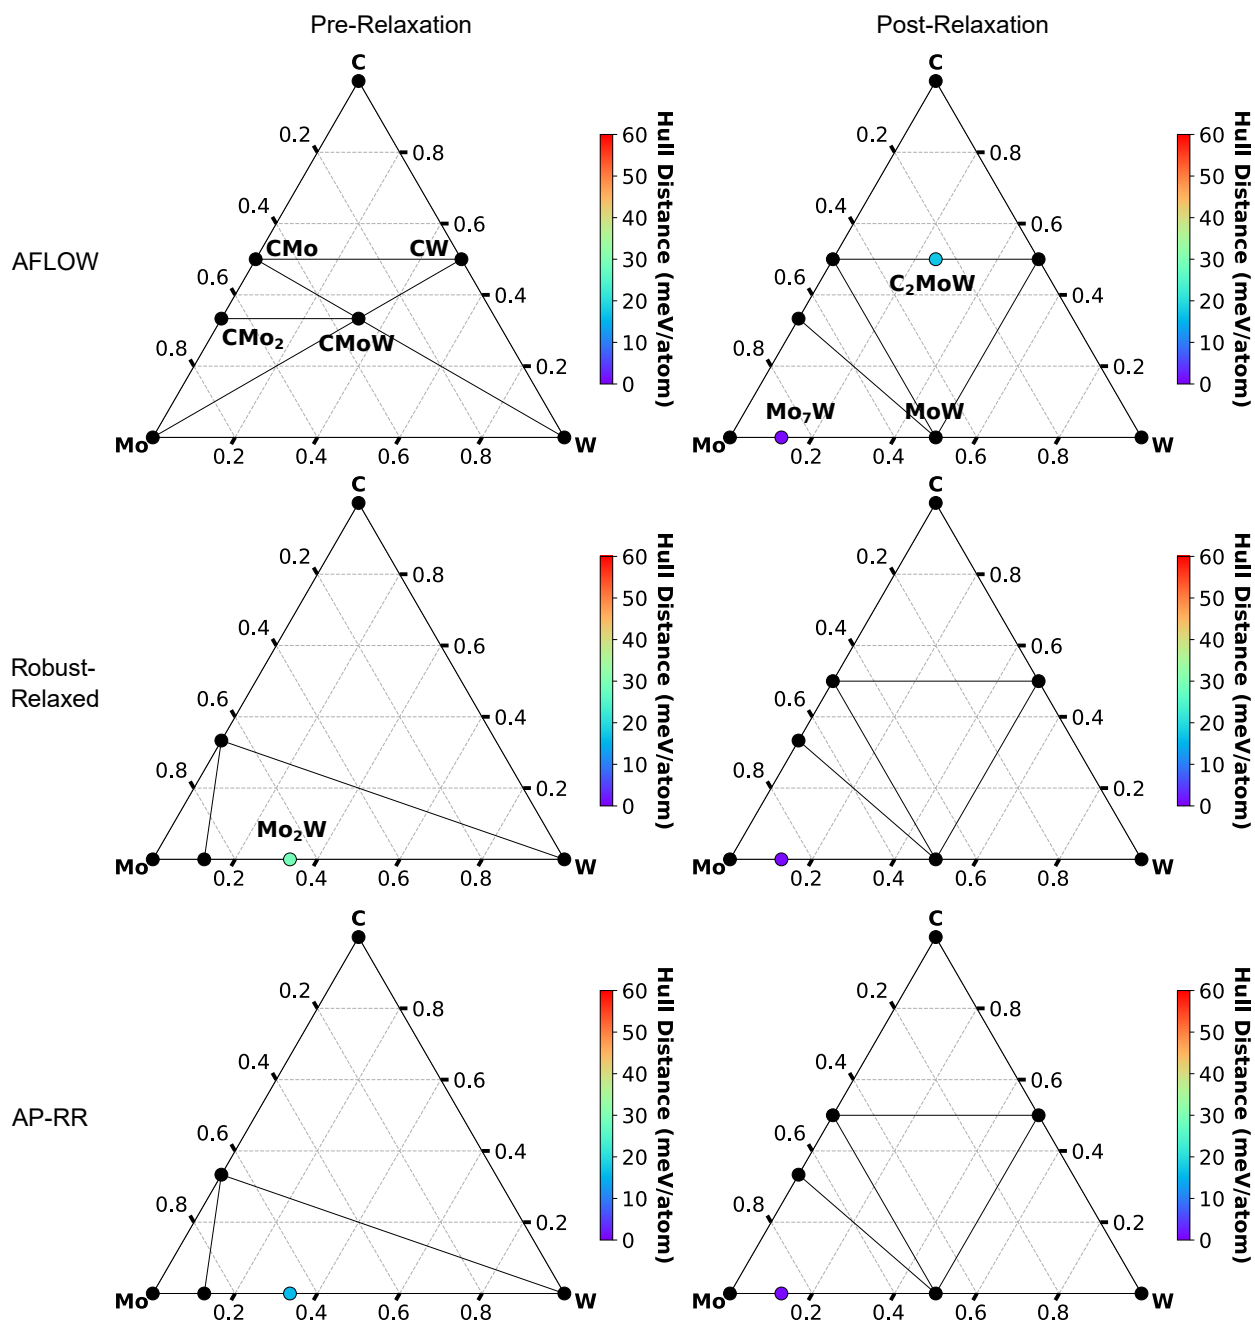


Figure S9: Convex Hulls of CMoW at Level 10.

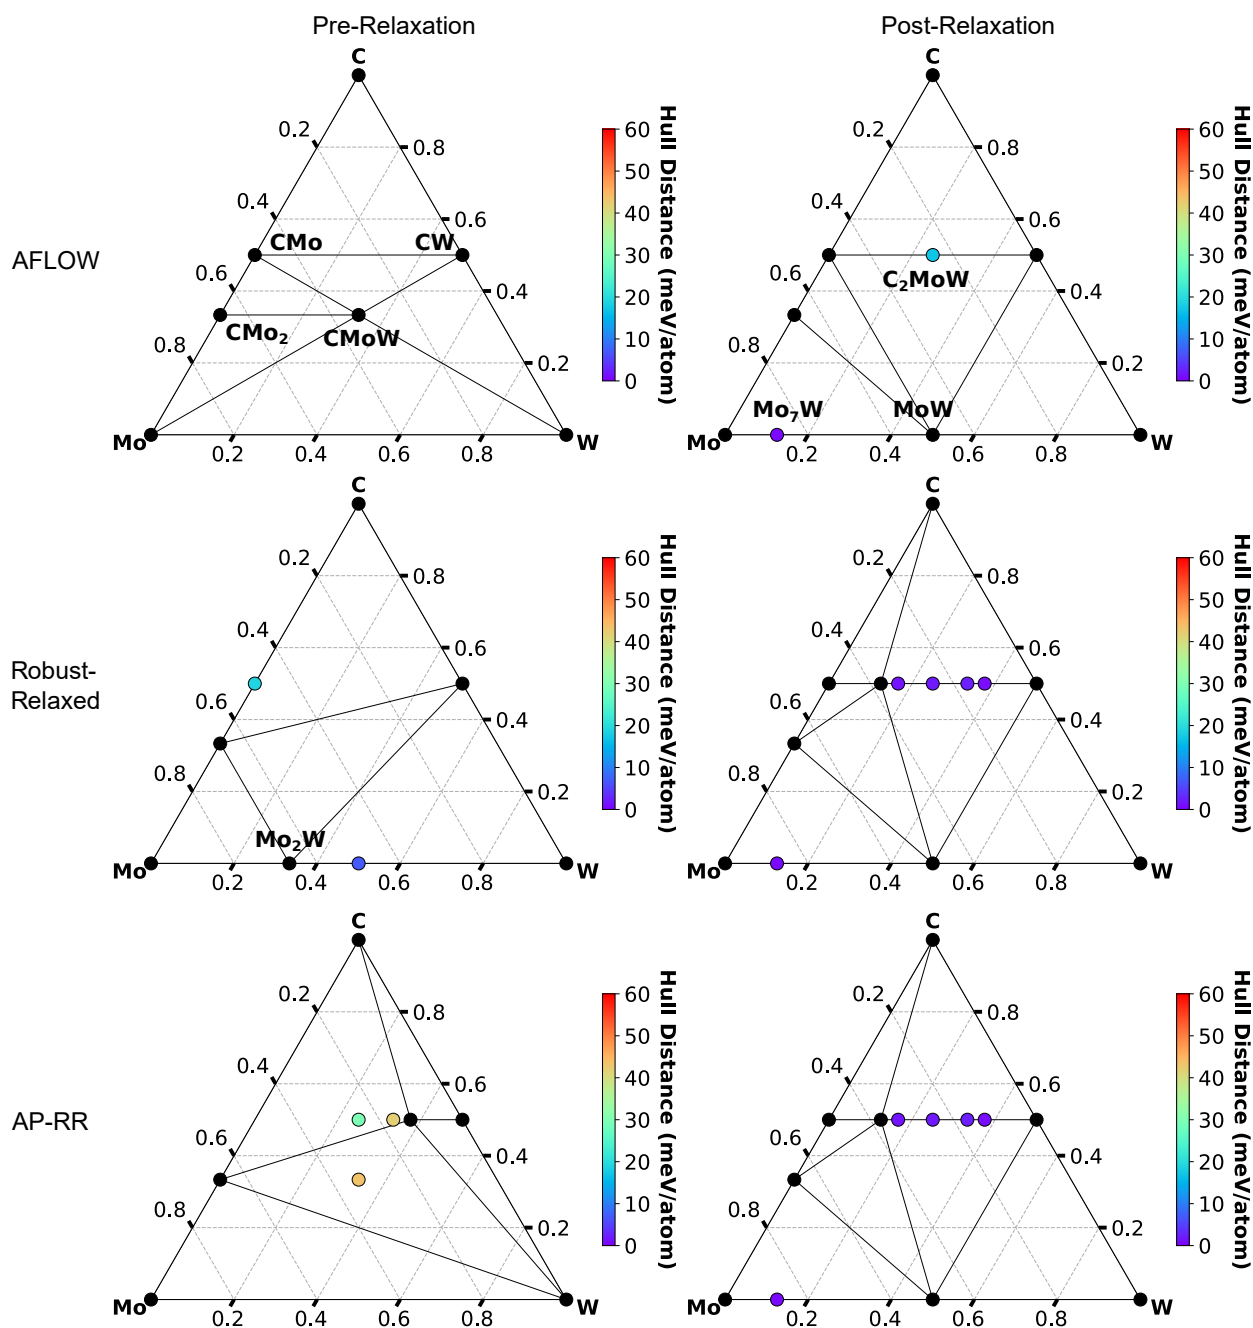


Figure S10: Convex Hulls of CMoW at Level 16. The unlabeled structures on the line between CMo and CW are, from left-to-right,  $C_4Mo_3W$ ,  $C_3Mo_2W$ ,  $C_2MoW$ ,  $C_3MoW_2$ , and  $C_4MoW_3$ .

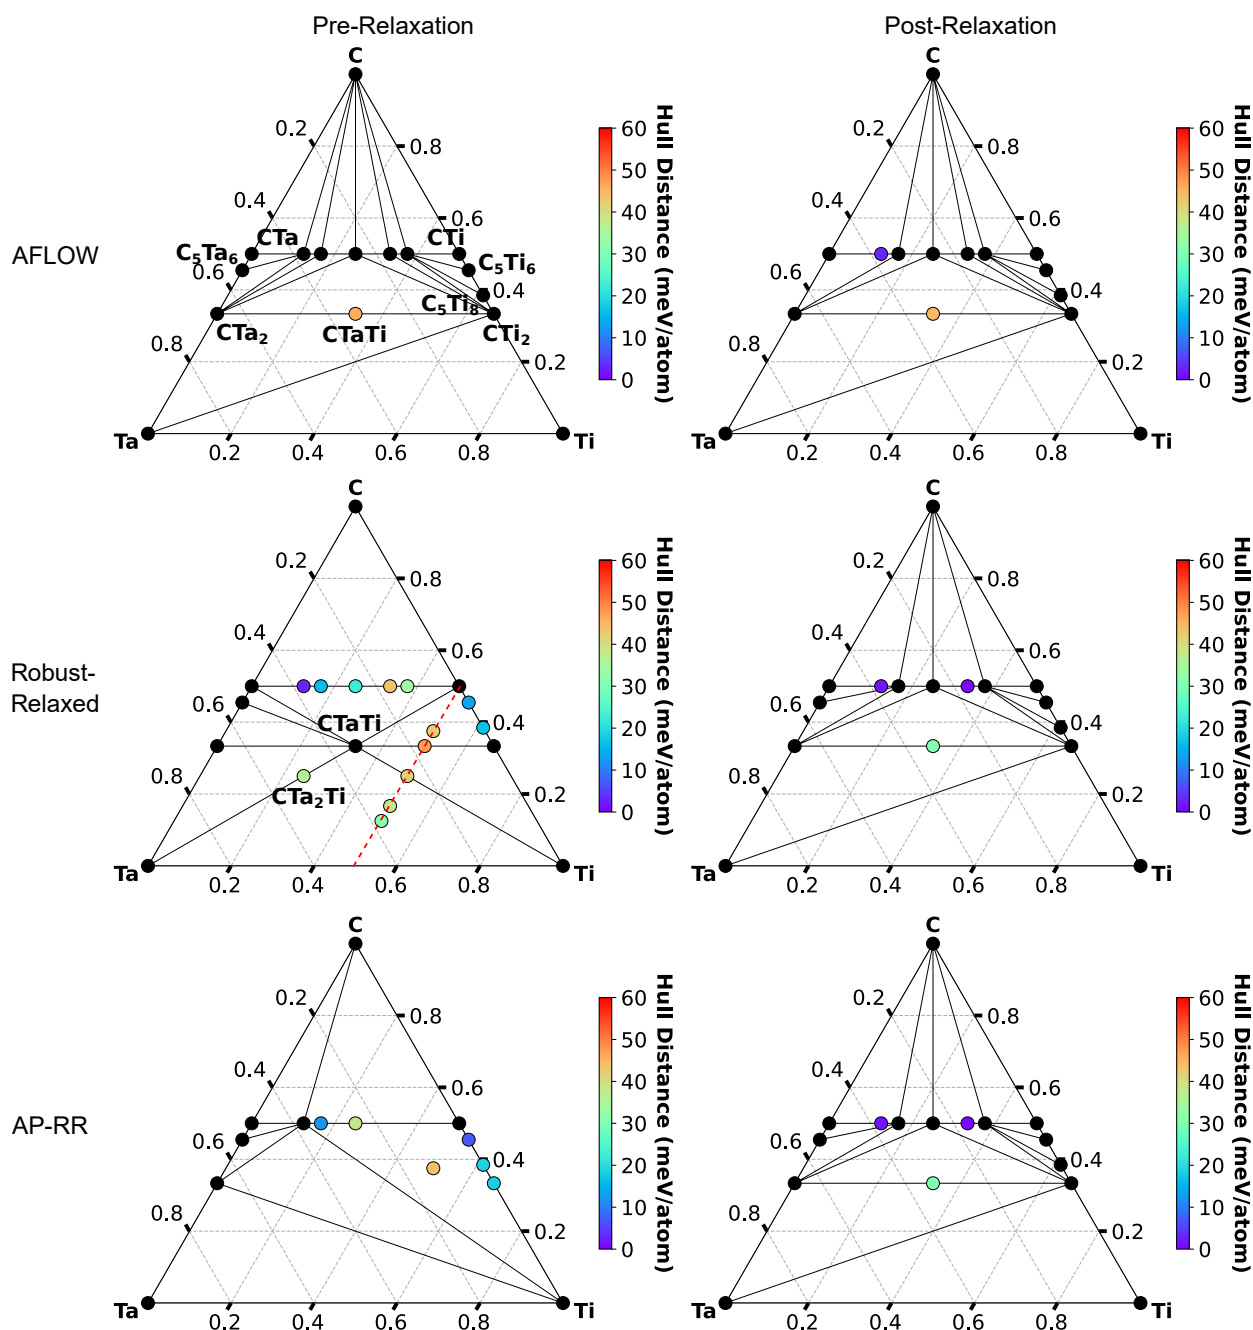


Figure S11: Convex Hulls of CTaTi at Level 10. The unlabeled structures on the line between CTa and CTi are, from left-to-right,  $C_4Ta_3Ti$ ,  $C_3Ta_2Ti$ ,  $C_2TaTi$ ,  $C_3TaTi_2$ , and  $C_4TaTi_3$ . The unlabeled structures on the dotted red line between CTi and TaTi are, from top-to-bottom,  $C_3TaTi_4$ ,  $C_2TaTi_3$ ,  $CTaTi_2$ ,  $CTa_2Ti_3$ , and  $CTa_3Ti_4$ .

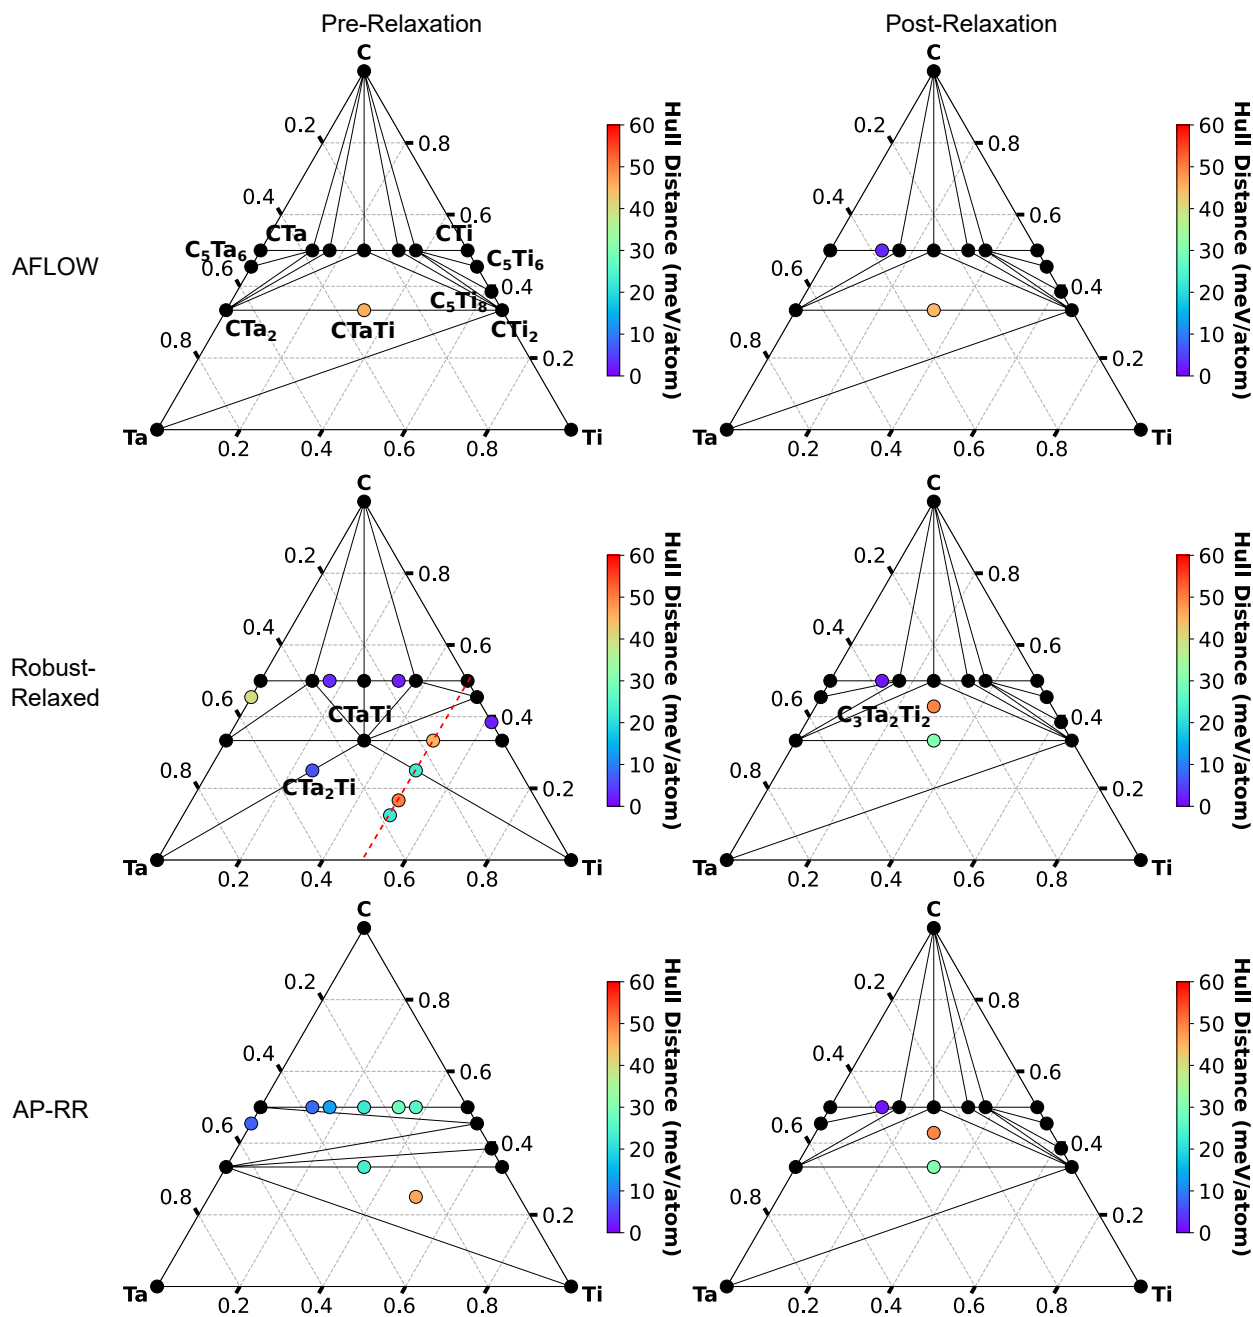


Figure S12: Convex Hulls of CTaTi at Level 16. The unlabeled structures on the line between CTa and CTi are, from left-to-right,  $C_4Ta_3Ti$ ,  $C_3Ta_2Ti$ ,  $C_2TaTi$ ,  $C_3TaTi_2$ , and  $C_4TaTi_3$ . The unlabeled structures on the dotted red line between CTi and TaTi are, from top-to-bottom,  $C_3TaTi_4$ ,  $C_2TaTi_3$ ,  $CTaTi_2$ ,  $CTa_2Ti_3$ , and  $CTa_3Ti_4$ .

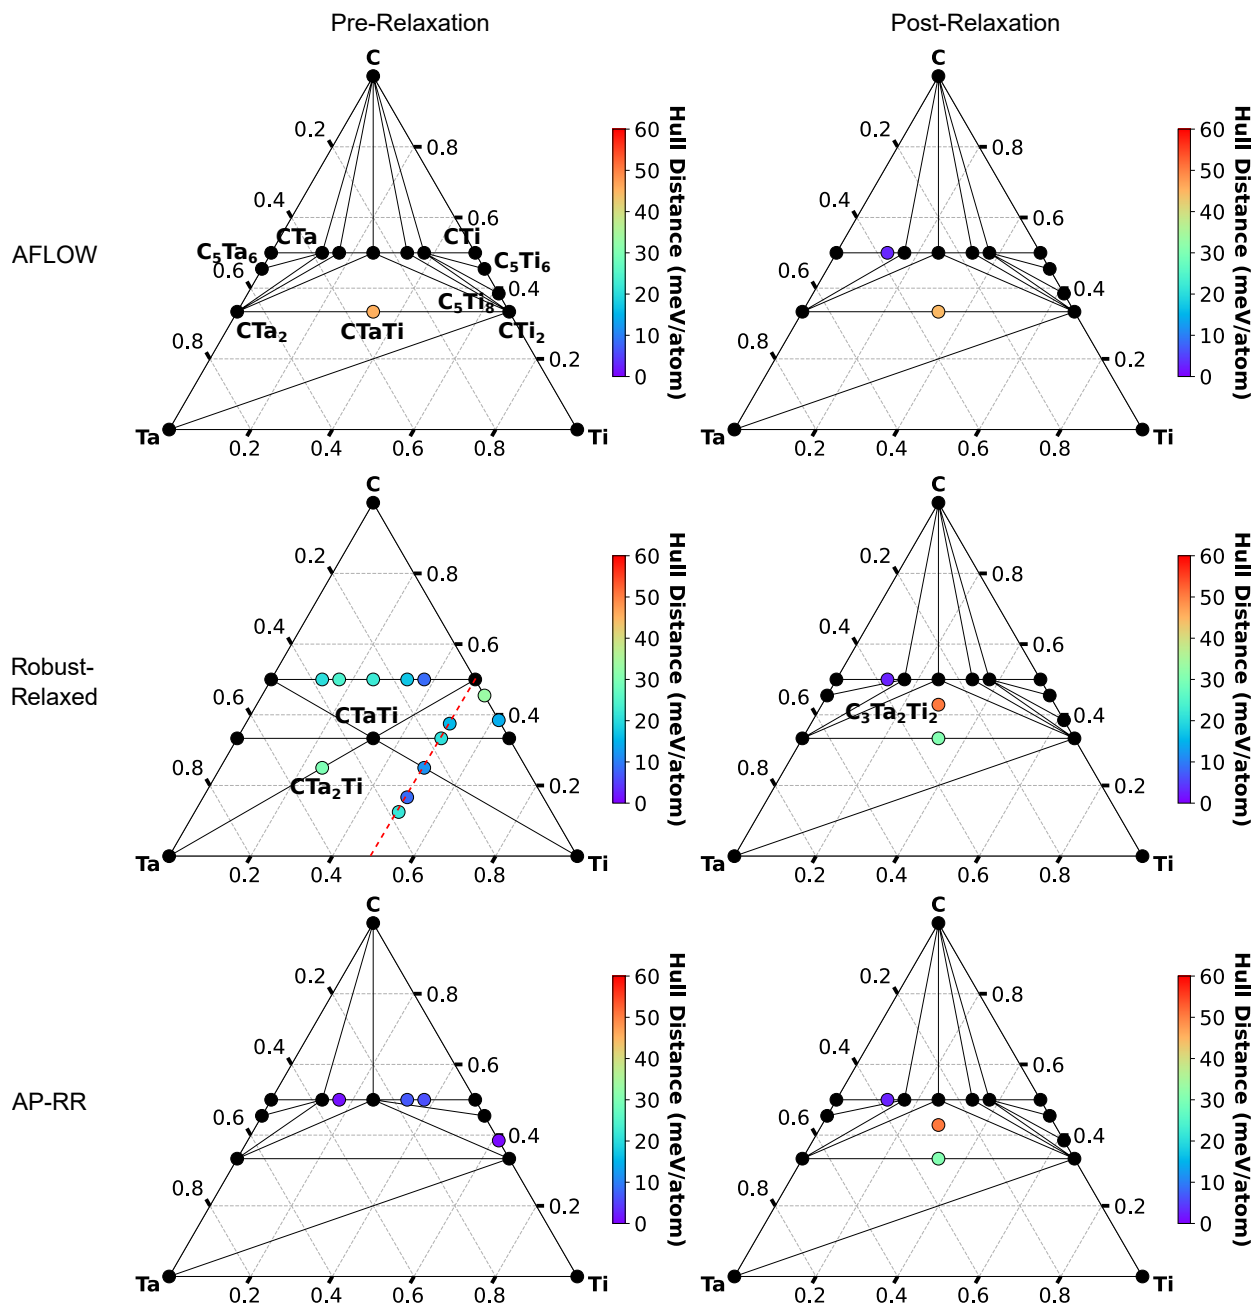


Figure S13: Convex Hulls of CTaTi at Level 22. The unlabeled structures on the line between CTa and CTi are, from left-to-right,  $C_4Ta_3Ti$ ,  $C_3Ta_2Ti$ ,  $C_2TaTi$ ,  $C_3TaTi_2$ , and  $C_4TaTi_3$ . The unlabeled structures on the dotted red line between CTi and TaTi are, from top-to-bottom,  $C_3TaTi_4$ ,  $C_2TaTi_3$ ,  $CTaTi_2$ ,  $CTa_2Ti_3$ , and  $CTa_3Ti_4$ .