



Point-Cache: Test-time Dynamic and Hierarchical Cache for Robust and Generalizable Point Cloud Analysis



Hongyu Sun^{1,2} Qiuhong Ke² Ming Cheng² Yongcai Wang¹ Deying Li¹ Chenhui Gou² Jianfei Cai²

¹Department of Computer Science, Renmin University of China, China

²Department of Data Science & AI, Monash University, Australia

Motivation

- **Problem:** Existing methods falter when faced with unpredictable shifts or new classes during online inference. They're stuck in a "training-only" mindset—impractical for dynamic, real-world scenarios.
- **Objective:** Enter Point-Cache—a training-free, test-time solution! We built a dynamic, hierarchical cache that learns on the fly from online test data alone. It captures both global shapes and local details of point clouds, adapting robustly to anything thrown its way.




Figure 1. Recognition accuracy comparison on clean and corrupted point cloud datasets. The suffix -C indicates datasets with corruptions. Models experience a severe performance drop when data corruptions arise.

Highlights

- First to tackle test-time point cloud recognition with only test data, generalizing to known and unseen classes, in a training-free manner.
- A plug-and-play hierarchical (global + local) cache boosting large 3D models like ULIP & ULIP-2 & OpenShape & Uni3D.
- Massive gains: +6.18% on ModelNet-C, +6.99% on ScanObjNN-C (hardest), +2.10% on Objaverse-LVIS (1,156 classes) across 8 benchmarks.

Challenges

We can obtain the global feature of a point cloud P and its text encoding easily. However, it is *non-trivial* to represent the local-part features of the point cloud P due to the following challenges.

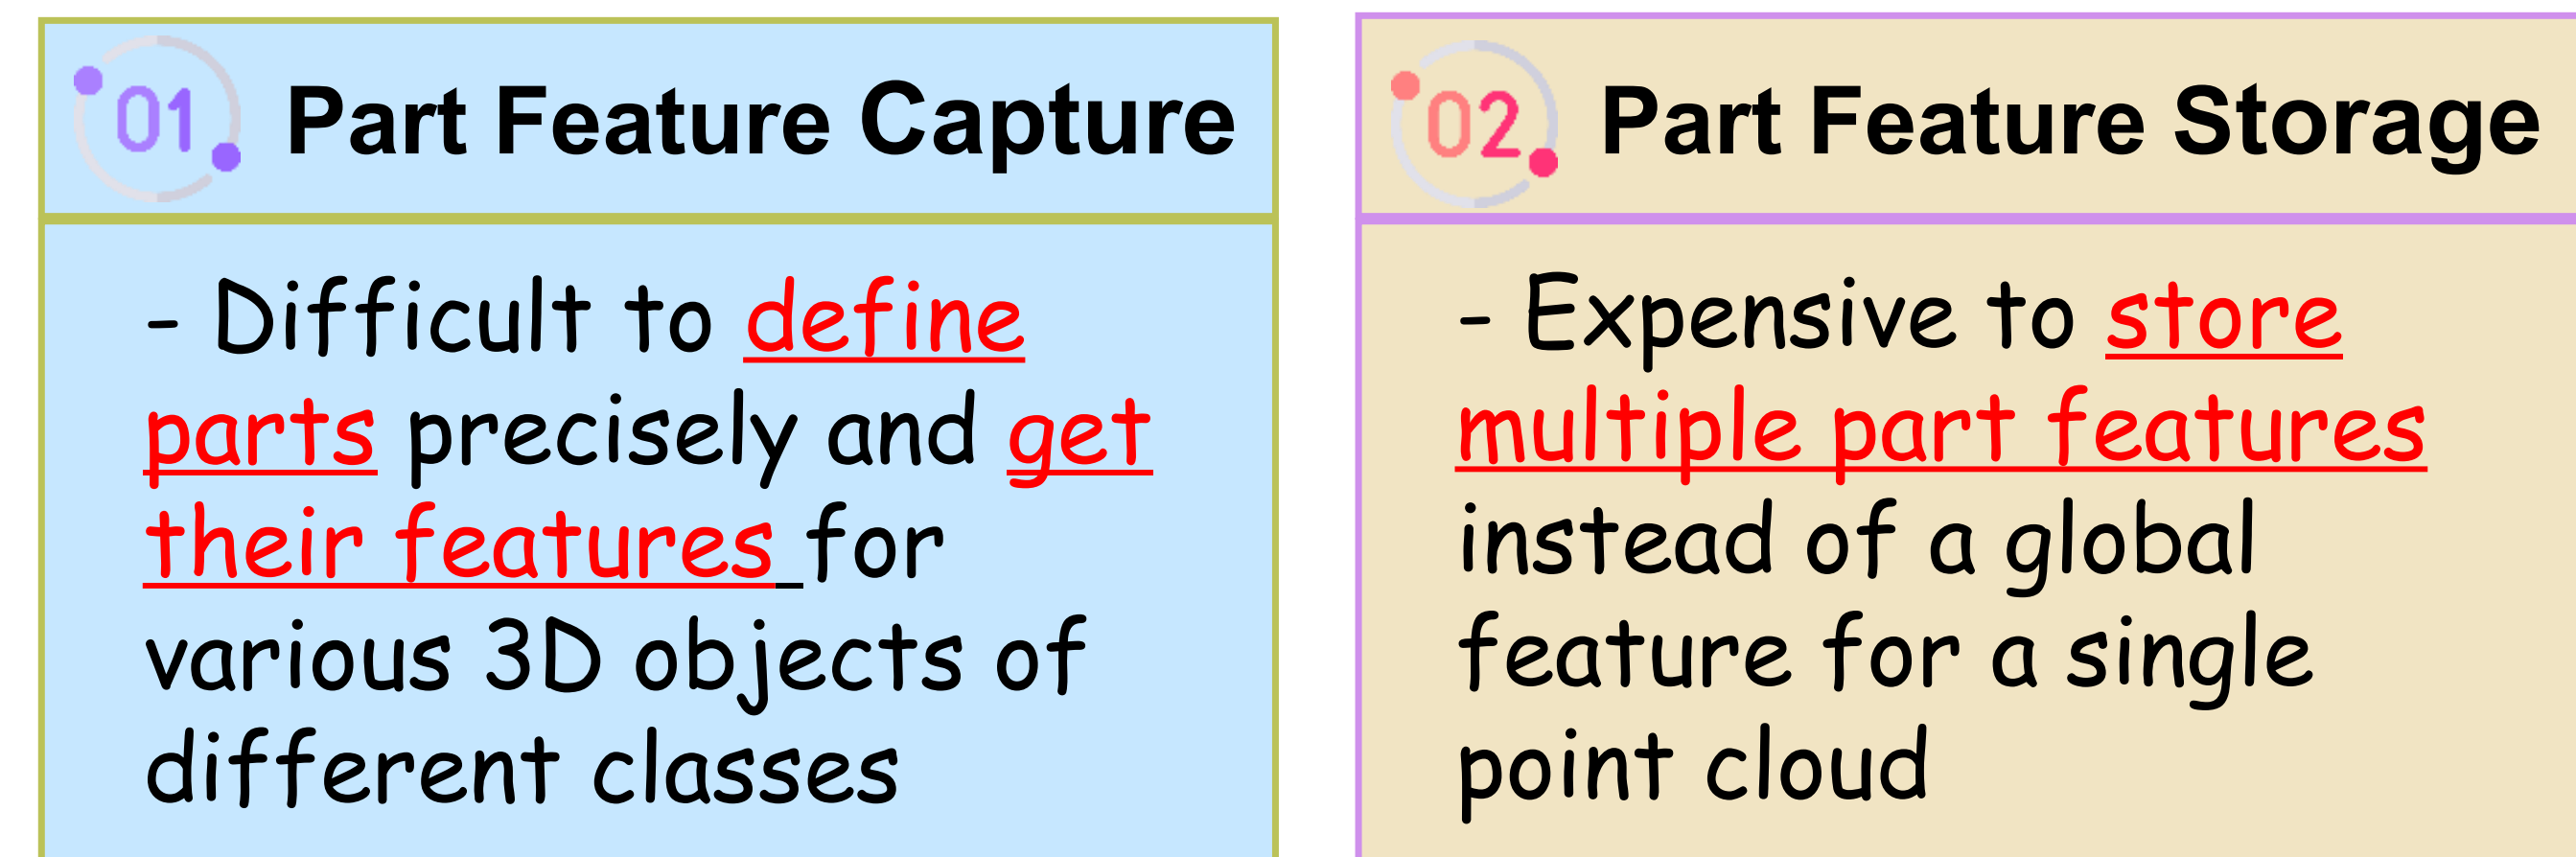


Figure 2. Challenges in encoding part feature for 3D objects of different classes: Part Feature Capture & Storage.

Proposed Methodology

We propose Point-Cache, a cache model to store critical fingerprints (which are essentially key-value or feature-label pairs) of test samples, to fully exploit the data distribution during test time.

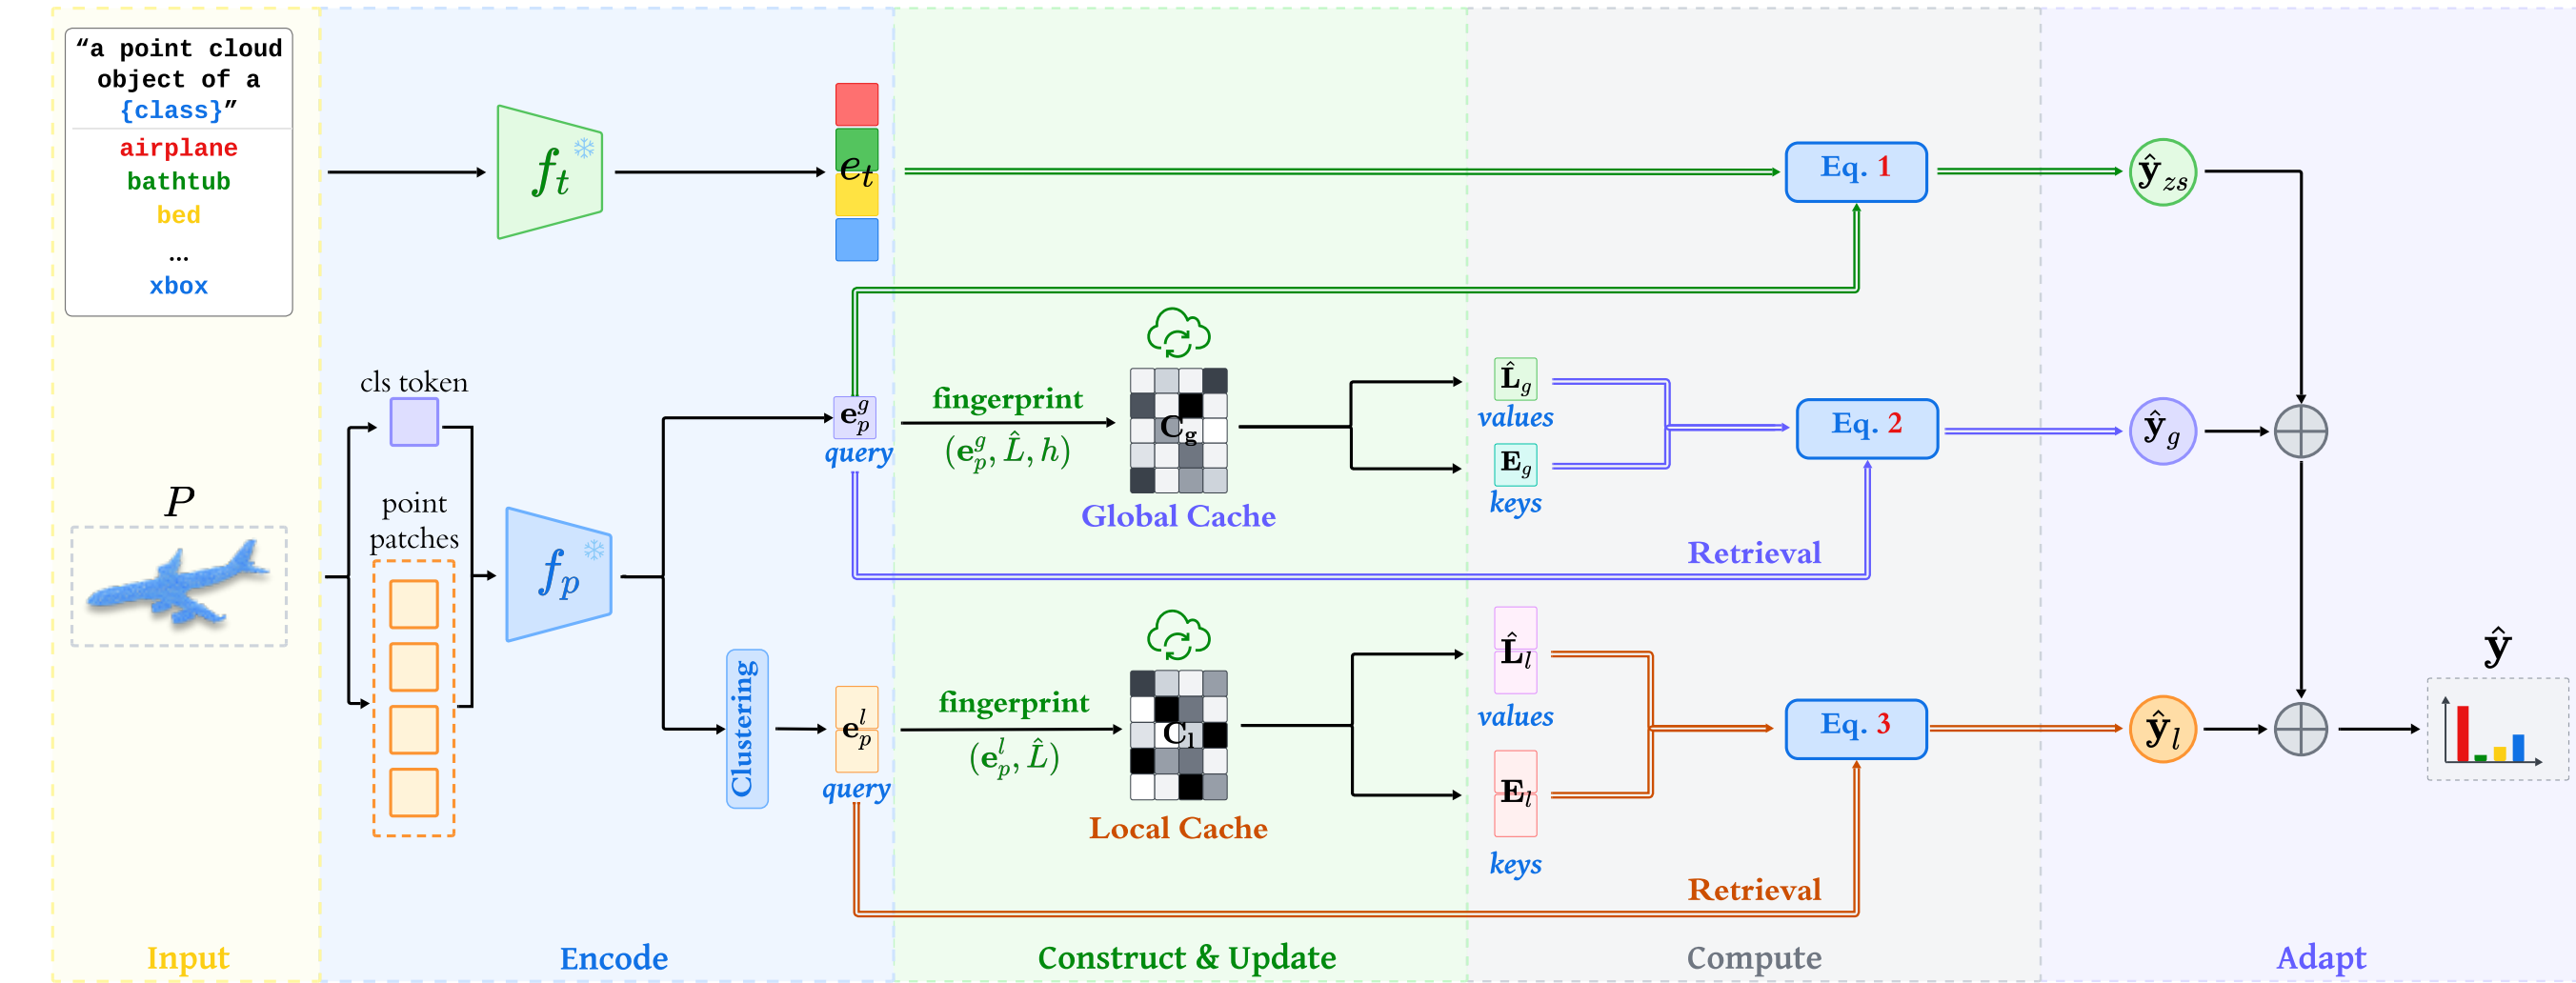


Figure 3. The overall pipeline of Point-Cache. The zero-shot predictions \hat{y}_{zs} of large 3D models are effectively adapted by our global cache logits \hat{y}_g and local cache logits \hat{y}_l to handle the distribution shifts and data corruptions.

Constuction & Update of Hierarchical Cache

Global cache: $C_g = (E_g, \hat{L}_g, h_g)$, storing (feature, label, entropy) triplets of online test samples.
Local cache: $C_l = (E_l, \hat{L}_l)$, storing (feature, label) pairs of online test samples.

Construction & Update of Global Cache and Local Cache:

- K : the **maximum** number of samples per class in the cache. Initially, both C_g and C_l are empty.
- When a new test sample P arrives, we generate its **global** fingerprint (e_p^g, \hat{L}, h) and **local** fingerprint (e_p^l, \hat{L}), and then assign it to class \hat{L} .
- Append the global fingerprint (e_p^g, \hat{L}, h) and the local fingerprint (e_p^l, \hat{L}) of current sample to class \hat{L} **if** the number of stored samples for class \hat{L} is **less than** K .
- Otherwise, **select** the fingerprint with **highest entropy** within class \hat{L} , **replace** this fingerprint with current fingerprint if the entropy of current sample is lower.

Test-time Adaptation by Our Hierarchical Cache

Zero-shot prediction of large 3D models

$$\hat{y}_{zs} = \{\hat{y}_{i|i=1}^C\}, \quad \hat{y}_i = \frac{\exp(\text{sim}(e_p^g, e_i^l)/\tau)}{\sum_{j=1}^C \exp(\text{sim}(e_p^g, e_j^l)/\tau)} \quad (1)$$

Adaptation by global cache

$$A_g = \exp(-\beta_g(1 - e_q^g E_g^T)), \quad \hat{L}_g = \text{OH}(\hat{L}_g), \quad \hat{y}_g = A_g \hat{L}_g \quad (2)$$

Adaptation by local cache

$$A_l = \exp(-\beta_l(1 - e_q^l E_l^T)), \quad \hat{L}_l = \text{OH}(\hat{L}_l), \quad \hat{y}_l = \phi(A_l \hat{L}_l) \quad (3)$$

Overall prediction

$$\hat{y} = \hat{y}_{zs} + \alpha_g \hat{y}_g + \alpha_l \hat{y}_l \quad (4)$$

Experiments

Robustness evaluation on ModelNet-C, which contains 7 types of data corruptions. Each clean point cloud has 1024 points and the corruption severity level is 2. The last column is the average.

Method	Clean Data ModelNet	Corruption Type							Avg.
		Add Global	Add Local	Drop Global	Drop Local	Rotate	Scale	Jitter	
ULIP [72]	56.16	33.55	43.92	54.70	50.89	55.27	50.20	44.08	47.52
+Global Cache(Ours)	<u>62.12</u>	45.79	47.98	56.85	53.89	<u>60.25</u>	<u>54.34</u>	48.91	<u>52.56</u>
+Hierarchical Cache(Ours)	64.22	46.15	<u>47.85</u>	59.16	56.00	61.47	55.35	49.92	53.70
ULIP-2 [73]	71.23	65.15	54.62	68.76	57.98	70.30	67.10	21.76	57.95
+Global Cache(Ours)	<u>73.95</u>	<u>67.02</u>	<u>59.32</u>	<u>71.35</u>	<u>61.59</u>	<u>72.37</u>	<u>68.40</u>	<u>28.20</u>	<u>61.18</u>
+Hierarchical Cache(Ours)	74.53	68.11	61.26	73.22	63.65	73.34	70.42	29.50	62.79
O-Shape [29]	84.56	71.64	67.79	81.56	73.58	82.01	78.48	59.36	73.49
+Global Cache(Ours)	<u>84.52</u>	<u>74.72</u>	<u>72.77</u>	82.41	<u>75.12</u>	83.18	78.93	67.91	<u>76.43</u>
+Hierarchical Cache(Ours)	84.04	74.84	73.70	<u>82.21</u>	76.26	<u>82.66</u>	<u>78.12</u>	68.35	76.59
Uni3D [88]	81.81	72.45	56.36	68.15	67.18	79.94	75.36	56.24	67.95
+Global Cache(Ours)	<u>83.14</u>	<u>76.13</u>	<u>66.49</u>	<u>71.43</u>	<u>69.81</u>	<u>81.52</u>	<u>75.85</u>	<u>61.43</u>	<u>71.81</u>
+Hierarchical Cache(Ours)	83.87	77.51	71.15	72.16	70.75	81.77	77.31	62.52	73.31

Figure 4. Comparison of recognition accuracy on ModelNet-C that contains 7 types of corruptions.

Generalization evaluation on a suite of datasets, where Omni3D and O-LVIS include 216 and 1,156 classes, respectively. The numbers in the bracket are the number of points per 3D object.

Method	ModelNet40 (10000 pts)	S-PB_RS_T50 (2048 pts)	O-LVIS (10000 pts)	Omni3D			Avg.
				(1024 pts)	4096 pts	16384 pts	
ULIP [72]	58.75	46.44	6.24	8.39	7.75	7.28	22.48
+Global Cache(Ours)	<u>61.22</u>	<u>50.21</u>	7.02	<u>10.00</u>	<u>9.36</u>	<u>8.43</u>	<u>24.37</u>
+Hierarchical Cache(Ours)	62.93	51.80	7.02	10.47	9.75	8.90	25.15
ULIP-2 [73]	72.97	47.13	30.26	26.36	29.20	26.58	38.75
+Global Cache(Ours)	<u>74.51</u>	<u>51.70</u>	32.65	<u>28.51</u>	<u>31.10</u>	<u>28.53</u>	<u>41.17</u>
+Hierarchical Cache(Ours)	75.53	54.98	<u>32.36</u>	29.37	31.24	29.44	42.15
Uni3D [88]	88.41	65.19	55.42	31.52	41.98	41.86	54.09
+Global Cache(Ours)	<u>88.86</u>	68.51	<u>53.36</u>	<u>34.97</u>	<u>45.13</u>	<u>45.19</u>	<u>56.00</u>
+Hierarchical Cache(Ours)	89.18	<u>68.24</u>	<u>55.19</u>	35.82	45.60	45.89	56.65

Figure 5. Comparison of recognition accuracy across a suite of datasets. S-PB_RS_T50 is the hardest split of ScanObjectNN. O-LVIS: Objaverse-LVIS. Omni3D: OmniObject3D, where objects have a different number of points.

Visualization of step-by-step adaptation for large 3D models by our Global Cache (GC) and Hierarchical cache (HC). The hierarchical cache model consists of a global and a local cache.

