

Appendix

A. Implementation Details

A.1. Downstream Tasks

ScanRefer [8]: The ScanRefer dataset contains 51,583 sentences written by humans to describe 800 scenes in ScanNet. We used the official split and allocated 36,665 and 9,508 samples for training and validation, respectively. The dataset is categorized into unique and multiple subsets based on whether the target object is a unique class in the scene. In this task, we need to find the target object described by a sentence. The evaluation metric for this task is accuracy under intersection over union (IoU) 0.25 and 0.5.

Nr3D/Sr3D [1]: The Sr3D dataset comprises of 83,572 utterances that are automatically generated using a template that focuses on the target-anchor spatial relationship. The Nr3D contains 45,503 human utterances. Both Sr3D and Nr3D are split by “Easy”/“Hard” and “ViewDep”/“ViewIndep”. Hard samples are the ones with two or more distractors in a scene. The view-dependent samples contain language descriptions that rely on viewing directions. These two datasets are also used for visual grounding like ScanRefer. But grounding accuracy with ground truth object proposal is evaluated in this setting.

ScanQA [3]: ScanQA is a dataset for 3D question answering with 41,363 questions and 58,191 answers. Different from 2D QA, ScanQA focuses more on spatial relations. We follow [3] to use exact matches EM@1 and EM@10 as the evaluation metric. EM@K means the percentage of top K answers from the model matches one of the ground-truth answers. Also, we include text similarity metrics to evaluate answers, including BLEU-4, ROUGE, METEOR, and CIDEr.

SQA3D [36]: SQA3D is a benchmark for scene understanding of embodied agents with 6.8k unique situations, 20.4k descriptions, and 33.4k diverse reasoning questions. Given a situation, an embodied agent must understand embodied activities, navigation instructions, and common sense, and perform multi-hop reasoning. The evaluation metric is answer accuracy under different types of questions.

A.2. Model Architecture

For the scene encoder, we use a three-layer Pointnet++ [39] with radius 0.2, 0.4, and sample all points to aggregate a 768-dimension feature. For all text and object tokens, the dimension is 768 in the following multi-modal fusion layers. In the unified encoder, the number of attention heads is set to 12 and the dimension of feedforward layers is set to 2048. For the visual grounding head, we use a two-layer MLP with a hidden dimension of 384. For the question-answering head and the situated reasoning head, we use a two-layer MLP with input dimensions 512 (from

the attention flat layer) and 768.

A.3. Training settings

The settings of pre-training including mask ratio, and optimization hyperparameters are introduced in the main paper. In this part, we elaborate on the fine-tuning details.

3D Visual Grounding: We only use a cross-entropy loss for fine-tuning 3D-VisTA on ScanRefer, Nr3D, and Sr3D. For all these grounding tasks, we set the batch size to 64, and the learning rate to $1e-4$. We multiply the learning rate of the text encoder by 0.1 to stabilize the training process. We fine-tune the pre-trained 3D-VisTA for 100, 100, and 50 epochs for ScanRefer, Nr3D, and Sr3D, respectively. AdamW with $\beta_1 = 0.9, \beta_2 = 0.98$ is chosen as the optimizer. We use a warmup of 5,000 steps and a cosine annealing learning rate schedule.

3D Question Answering: We use a cross-entropy answer classification loss and a visual grounding loss for ScanQA. The batch size is 64 and the learning rate is $1e-4$. 3D-VisTA is fine-tuned for 30 epochs with 2000 warmup steps for this task. Other optimization parameters are the same as the visual grounding task.

3D Situated Reasoning: Answer classification loss is used for SQA3D. We fine-tune 3D-VisTA for 50 epochs. Other optimization parameters are the same as the 3D question-answering task.

A.4. ScanScribe

In the main paper, we introduce our method of generating new scene-text pairs from scene graphs and large language models. More examples and cases are provided in this section. We support 40 relations and the mapping of relations to descriptions for the template-based generation is shown in Table A1.

With these relations, we can use templates like “This is a object, a neighbor is relation to object” and utilize GPT-3 to increase text diversity. During pre-training, to balance the proportion of template and GPT-3 generated texts in the 3R-Scan dataset, we duplicate texts from GPT-3 to 15 times for pre-training. Examples from both template-based generation and GPT-3 are presented in Fig. A1. We can observe that given entities and relations in a scene, GPT-3 can summarize them into a fluent and natural sentence.

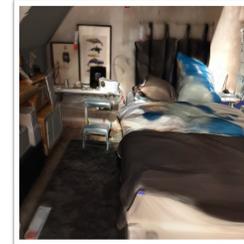
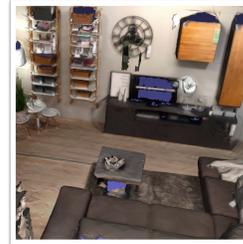
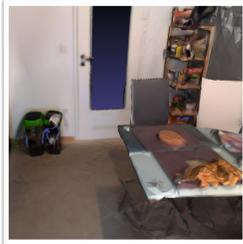
B. Additional Results

We provide ablation studies on the use the template-generated text and GPT-3-generated text. As shown in Table A2, GPT-3-generated text improves Sr3D and Nr3D by 1.0% and 1.5%, while having little impact on ScanRefer and ScanQA. More qualitative results including failure cases are provided in Fig. A2.

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

Scene



GPT-3

The trash can is behind another trash can, and it's located on the left side of a kitchen cabinet, a white chair, and a box. On the right side, there is a rack and a shoe.

The stove is attached to the kitchen cabinet and is positioned on the left side of the sink. It's also located in front of both the kettle and the toaster.

The coffee table is on the right side of the TV stand and close to the gray sofa. It is also on the right side of the stool.

The chair is on the left side of the blue bed, close by the table. It's on the right side of the cabinet and the box, and behind the light. Additionally, there's another chair on its left.

Template

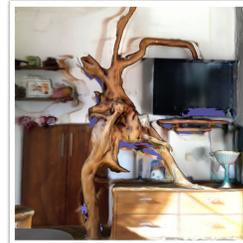
This is a trash can. It is behind the another trash can. It is on the left side of the kitchen cabinet. It is on the left side of the white chair. It is on the left side of the box. It is on the right side of the rack. It is on the right side of the shoe.

This is a stove. It is attached to the kitchen cabinet. It is on the left side of the sink. It is in front of the kettle. It is in front of the toaster.

This is a coffee table. It is close by the gray sofa. It is on the right side of the tv stand. It is on the right side of the stool.

This is a chair. It is on the left side of the blue bed. It is close by the table. It is on the right side of the cabinet. It is on the right side of the box. It is behind the light. It is on the left side of the another chair.

Scene



GPT-3

The doorframe is on the right side of the counter and on the left side of the toilet.

The purple curtain is near the rectangular brown window and another purple curtain

The rectangular black TV is standing on the brown TV stand. It is in front of the gray sofa and on the right side of the brown chair.

The brown rocking chair is on the left side of the brown chair, in front of the rectangular white fireplace, on the left side of the square brown box, and on the right side of the green plant.

Template

This is a doorframe. It is on the left side of the toilet.

This is a purple curtain. It is close by the rectangular brown window. It is close by the another purple curtain.

This is a rectangular black tv. It is standing on the brown tv stand. It has the same state as the rectangular white fireplace. It has the same color as the black pillow. It is darker than the white lamp.

This is a brown rocking chair. It is on the left side of the brown chair. It is on the left side of the rectangular white fireplace. It is on the left side of the square brown box. It is on the right side of the green plant

Figure A1: Examples of both template and GPT-3 generated text in ScanScribe dataset. GPT-3 generated text is more natural than template-generated text.

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

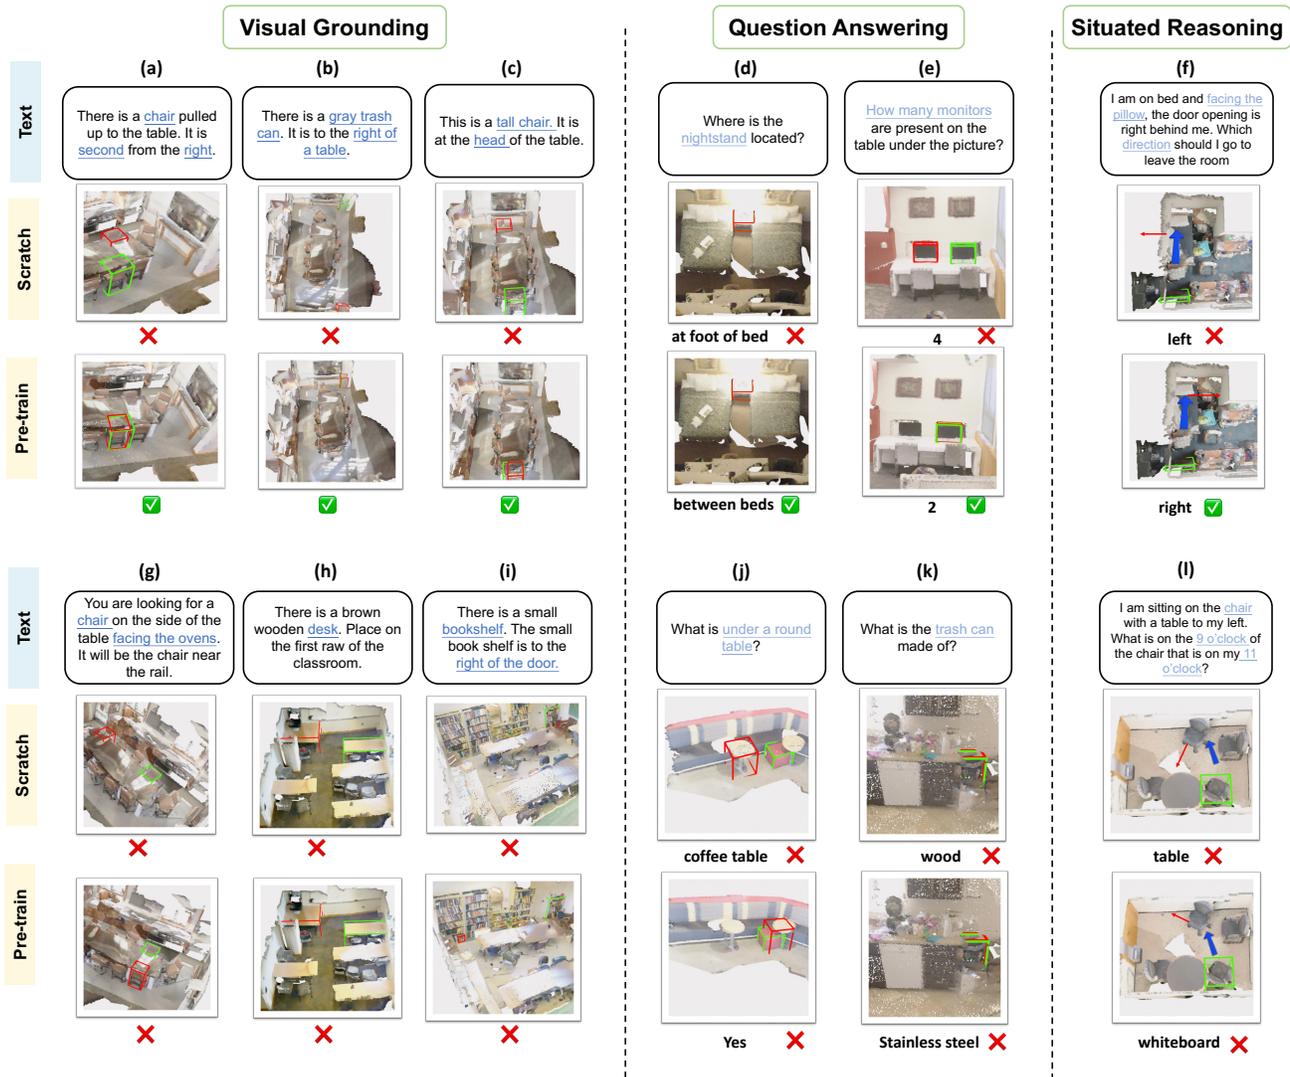


Figure A2: Qualitative results on ScanRefer, ScanQA, and SQA3D. Green and red denote the ground-truth and predicted object boxes, respectively. As shown in (a,b,c,d,e,f), the pre-trained 3D-VisTA shows advantages in spatial reasoning, concept grounding, and situation understanding. In spite of these advantages, (g, h, j) indicate that for some complicated cases with spatial relations, the pre-trained model still cannot understand them. (i, k) show that our model is still limited by the semantic information extracted by point clouds, which fail to locate the right object or understand texture. From (l), we can observe that our model may fail in the case requiring complex multi-hop reasoning.

Table A1: The mapping of relations to descriptions.

Relation	Description
supported by	is supported by the
left	is on the left side of the
right	is on the right side of the
front	is in front of the
behind	is behind the
close by	is close by the
inside	is inside the
bigger than	is bigger than the
smaller than	is smaller than the
higher than	is higher than the
lower than	is lower than the
same symmetry as	has the same symmetry as the
same as	is the same as the
attached to	is attached to the
standing on	is standing on the
lying on	is lying on the
hanging on	is hanging on the
connected to	is connected to the
leaning against	is leaning against the
part of	is part of the
belonging to	is belonging to the
built in	is built in the
standing in	is standing in the
covers	covers the
lying in	is lying in the
hanging in	is hanging in the
same color	has the same color as the
same material	has the same material as the
same texture	has the same texture as the
same shape	has the same shape as the
same state	has the same state as the
same object type	has the same object type as the
messier than	is messier than the
cleaner than	is cleaner than the
fuller than	is fuller than the
more closed	is more closed to the
more open	is more open than the
brighter than	is brighter than the
darker than	is darker than the
more comfortable than	is more comfortable than the

Table A2: Ablation studies on the template and GPT-3 generated text from 3R-Scan. We report the results on ScanRefer, Sr3D, Nr3D and ScanQA.

Template	GPT-3	ScanRefer	Sr3D	Nr3D	ScanQA
✓	×	57.4	75.4	62.7	23.7
✓	✓	57.4	76.4	64.2	23.8