

# 000 001 002 003 004 005 RECONVIAGEN: TOWARDS ACCURATE MULTI-VIEW 006 3D OBJECT RECONSTRUCTION VIA GENERATION 007

008  
009 **Anonymous authors**  
010  
011  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053  
054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067  
068  
069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083  
084  
085  
086  
087  
088  
089  
090  
091  
092  
093  
094  
095  
096  
097  
098  
099  
100  
101  
102  
103  
104  
105  
106  
107  
108  
109  
110  
111  
112  
113  
114  
115  
116  
117  
118  
119  
120  
121  
122  
123  
124  
125  
126  
127  
128  
129  
130  
131  
132  
133  
134  
135  
136  
137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
303  
304  
305  
306  
307  
308  
309  
310  
311  
312  
313  
314  
315  
316  
317  
318  
319  
320  
321  
322  
323  
324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
639  
640  
641  
642  
643  
644  
645  
646  
647  
648  
649  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
699  
700  
701  
702  
703  
704  
705  
706  
707  
708  
709  
709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809  
809  
810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
859  
860  
861  
862  
863  
864  
865  
866  
867  
868  
869  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
909  
910  
911  
912  
913  
914  
915  
916  
917  
918  
919  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
969  
970  
971  
972  
973  
974  
975  
976  
977  
978  
979  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1039  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
1049  
1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079  
1079  
1080  
1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1129  
1130  
1131  
1132  
1133  
1134  
1135  
1136  
1137  
1138  
1139  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187  
1188  
1189  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1199  
1200  
1201  
1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1239  
1240  
1241  
1242  
1243  
1244  
1245  
1246  
1247  
1248  
1249  
1249  
1250  
1251  
1252  
1253  
1254  
1255  
1256  
1257  
1258  
1259  
1259  
1260  
1261  
1262  
1263  
1264  
1265  
1266  
1267  
1268  
1269  
1269  
1270  
1271  
1272  
1273  
1274  
1275  
1276  
1277  
1278  
1279  
1279  
1280  
1281  
1282  
1283  
1284  
1285  
1286  
1287  
1288  
1289  
1289  
1290  
1291  
1292  
1293  
1294  
1295  
1296  
1297  
1298  
1299  
1299  
1300  
1301  
1302  
1303  
1304  
1305  
1306  
1307  
1308  
1309  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1319  
1320  
1321  
1322  
1323  
1324  
1325  
1326  
1327  
1328  
1329  
1329  
1330  
1331  
1332  
1333  
1334  
1335  
1336  
1337  
1338  
1339  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349  
1349  
1350  
1351  
1352  
1353  
1354  
1355  
1356  
1357  
1358  
1359  
1359  
1360  
1361  
1362  
1363  
1364  
1365  
1366  
1367  
1368  
1369  
1369  
1370  
1371  
1372  
1373  
1374  
1375  
1376  
1377  
1378  
1379  
1379  
1380  
1381  
1382  
1383  
1384  
1385  
1386  
1387  
1388  
1389  
1389  
1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1399  
1400  
1401  
1402  
1403  
1404  
1405  
1406  
1407  
1408  
1409  
1409  
1410  
1411  
1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457  
1458  
1459  
1459  
1460  
1461  
1462  
1463  
1464  
1465  
1466  
1467  
1468  
1469  
1469  
1470  
1471  
1472  
1473  
1474  
1475  
1476  
1477  
1478  
1479  
1479  
1480  
1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1489  
1490  
1491  
1492  
1493  
1494  
1495  
1496  
1497  
1498  
1499  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1509  
1510  
1511  
1512  
1513  
1514  
1515  
1516  
1517  
1518  
1519  
1519  
1520  
1521  
1522  
1523  
1524  
1525  
1526  
1527  
1528  
1529  
1529  
1530  
1531  
1532  
1533  
1534  
1535  
1536  
1537  
1538  
1539  
1539  
1540  
1541  
1542  
1543  
1544  
1545  
1546  
1547  
1548  
1549  
1549  
1550  
1551  
1552  
1553  
1554  
1555  
1556  
1557  
1558  
1559  
1559  
1560  
1561  
1562  
1563  
1564  
1565  
1566  
1567  
1568  
1569  
1569  
1570  
1571  
1572  
1573  
1574  
1575  
1576  
1577  
1578  
1579  
1579  
1580  
1581  
1582  
1583  
1584  
1585  
1586  
1587  
1588  
1589  
1589  
1590  
1591  
1592  
1593  
1594  
1595  
1596  
1597  
1598  
1599  
1599  
1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619  
1619  
1620  
1621  
1622  
1623  
1624  
1625  
1626  
1627  
1628  
1629  
1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1649  
1650  
1651  
1652  
1653  
1654  
1655  
1656  
1657  
1658  
1659  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1669  
1670  
1671  
1672  
1673  
1674  
1675  
1676  
1677  
1678  
1679  
1679  
1680  
1681  
1682  
1683  
1684  
1685  
1686  
1687  
1688  
1689  
1689  
1690  
1691  
1692  
1693  
1694  
1695  
1696  
1697  
1698  
1699  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709  
1709  
1710  
1711  
1712  
1713  
1714  
1715  
1716  
1717  
1718  
1719  
1719  
1720  
1721  
1722  
1723  
1724  
1725  
1726  
1727  
1728  
1729  
1729  
1730  
1731  
1732  
1733  
1734  
1735  
1736  
1737  
1738  
1739  
1739  
1740  
1741  
1742  
1743  
1744  
1745  
1746  
1747  
1748  
1749  
1749  
1750  
1751  
1752  
1753  
1754  
1755  
1756  
1757  
1758  
1759  
1759  
1760  
1761  
1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1779  
1780  
1781  
1782  
1783  
1784  
1785  
1786  
1787  
1788  
1789  
1789  
1790  
1791  
1792  
1793  
1794  
1795  
1796  
1797  
1798  
1799  
1799  
1800  
1801  
1802  
1803  
1804  
1805  
1806  
1807  
1808  
1809  
1809  
1810  
1811  
1812  
1813  
1814  
1815  
1816  
1817  
1818  
1819  
1819  
1820  
1821  
1822  
1823  
1824  
1825  
1826  
1827  
1828  
1829  
1829  
1830  
1831  
1832  
1833  
1834  
1835  
1836  
1837  
1838  
1839  
1839  
1840  
1841  
1842  
1843  
1844  
1845  
1846  
1847  
1848  
1849  
1849  
1850  
1851  
1852  
1853  
1854  
1855  
1856  
1857  
1858  
1859  
1859  
1860  
1861  
1862  
1863  
1864  
1865  
1866  
1867  
1868  
1869  
1869  
1870  
1871  
1872  
1873  
1874  
1875  
1876  
1877  
1878  
1879  
1879  
1880  
1881  
1882  
1883  
1884  
1885  
1886  
1887  
1888  
1889  
1889  
1890  
1891  
1892  
1893  
1894  
1895  
1896  
1897  
1898  
1899  
1899  
1900  
1901  
1902  
1903  
1904  
1905  
1906  
1907  
1908  
1909  
1909  
1910  
1911  
1912  
1913  
1914  
1915  
1

054 models tend to have holes, artifacts, and missing/blurred geometric details, which severely restrict  
 055 the reconstruction completeness He et al. (2024); Xu et al. (2024c).  
 056

057 Recent advances in diffusion-based 3D generative techniques have shown great promise in addressing  
 058 these limitations. These techniques leverage 3D generative priors learned from large-scale 3D  
 059 data to generate complete 3D outputs from sparse- or even single-view images Li et al. (2024b);  
 060 Zhao et al. (2025); Li et al. (2025a); Zhang et al. (2024b); Ye et al. (2025). Such strong genera-  
 061 tive priors can effectively “hallucinate” the invisible portions of objects with plausible high-quality  
 062 geometry and appearance, thereby showing great potential in 3D reconstruction by filling in the  
 063 missing details and improving the completeness. However, the stochastic nature of the diffusion-  
 064 based inference process introduces significant uncertainty and variability in the generated results,  
 065 making it challenging to achieve high accuracy and reliability, especially the pixel-level alignment  
 066 required in accurate reconstruction. This stochasticity has largely hindered the effective integration  
 067 of diffusion-based 3D generative priors into existing multi-view reconstruction frameworks.

068 Pioneering explorations have been made in the field of 3D diffusion-based generation from multi-  
 069 view images Xiang et al. (2024); Zhao et al. (2025). However, their predictions still suffer from  
 070 inaccurate global structures and inconsistent local details. The inherent key reasons of the failure  
 071 include (i) the insufficiency in constructing cross-view correlations when extracting multi-view im-  
 072 age features as conditions, resulting in inaccurate estimation in both object geometry and texture, at  
 073 the global and local level, (ii) the poor controllability and stability of the denoising process during  
 074 inference, which easily results in inconsistency with input views especially in detailed geometry and  
 075 texture estimation. To address these issues, we present ReconViaGen that innovatively integrates  
 076 multi-view stereo priors into the diffusion-based generative framework for object reconstruction.  
 077 Our solution includes three stages: (i) a pre-trained strong reconstructor Wang et al. (2025a) is de-  
 078 veloped to build a multi-view stereo understanding of the object geometry and texture, aggregated  
 079 into a single global token list and a set of local token lists, for representing the global geometry and  
 080 the detailed per-view appearance, respectively; (ii) a coarse-to-fine 3D generator Xiang et al. (2024)  
 081 first estimates the coarse structure and then produces the fine textured mesh, under the conditioning  
 082 of global and local tokens from the first stage, respectively; (iii) refining the estimated poses from the  
 083 reconstructor using the generation from the second stage, and encouraging the pixel-wise alignment  
 084 with input views using a novel rendering-aware velocity compensation mechanism, where input im-  
 085 ages coupled with estimated camera poses are used to explicitly guide the denoising trajectory of  
 086 local latent representations.

087 Extensive experiments on the Dora-bench Chen et al. (2024) and OminiObject3D Wu et al. (2023b)  
 088 datasets validate that our ReconViaGen can achieve state-of-the-art (SOTA) reconstruction perfor-  
 089 mance in both global shape accuracy and completeness and local details in geometry and textures.  
 090 Our contributions are summarized as follows:

- 091 • We propose a novel framework called ReconViaGen, which is the first to integrate strong recon-  
 092 struction priors into a diffusion-based 3D generator for accurate and complete multi-view object  
 093 reconstruction. A key design is to aggregate image features rich in reconstruction priors as multi-  
 094 view-aware diffusion conditions.
- 095 • The generation adopts a coarse-to-fine paradigm, which leverages global and local reconstruc-  
 096 tion-based conditions to generate accurate coarse and then fine results in both geometry and texture.  
 097 Additionally, a novel rendering-aware velocity compensation mechanism is proposed that con-  
 098 strains the denoising trajectory of local latent representations for detailed pixel-level alignment.
- 099 • Extensive experiments on the Dora-bench and OminiObject3D datasets are conducted that vali-  
 100 date the effectiveness and superiority of the proposed ReconViaGen, which achieves SOTA per-  
 101 formance.

## 102 2 RELATED WORK

103 **Single-view 3D Generation.** Great developments have been made in single-view 3D Object Gen-  
 104 eration. Recent methods can be divided into two groups: 2D prior-based and 3D native generative  
 105 methods. DreamFusion Poole et al. (2022) and its following successors Tang et al. (2024); Qiu et al.  
 106 (2024); Wang et al. (2024b); Lin et al. (2023); Tang et al. (2023a) distill the 3D knowledge from pre-  
 107 trained 2D models. Another line of work develops multi-view diffusion based on pre-trained image

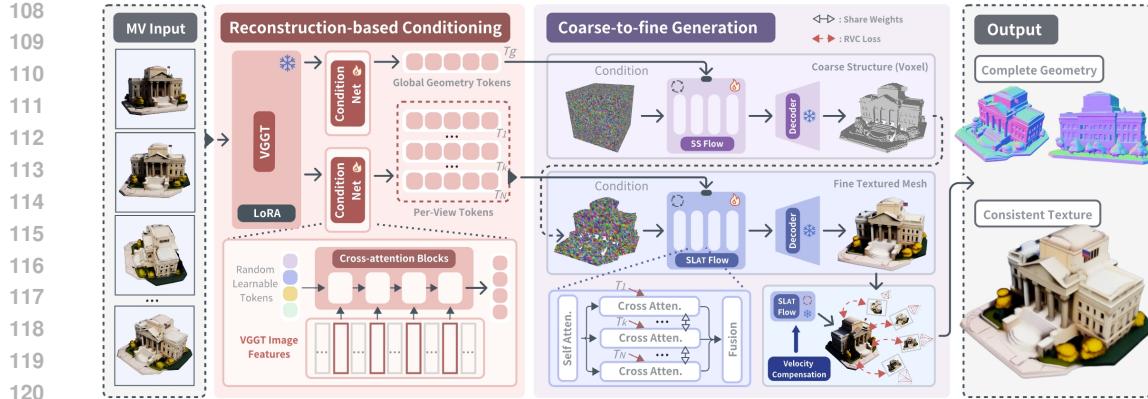


Figure 2: An overview illustration of the proposed ReconViaGen framework, which integrates strong reconstruction priors with 3D diffusion-based generation priors for accurate reconstruction at both the global and local level.

or video generators and conducts view fusion for 3D outputs Li et al. (2023); Xu et al. (2024b); Tang et al. (2025); Xu et al. (2024d); Wang et al. (2025b); Wei et al. (2024); Liu et al. (2023a; 2024); Xu et al. (2024e); Li et al. (2024a); Zuo et al. (2024); Wu et al. (2024a). Differently, 3D native generative methods employ diffusion based on different 3D representations like point clouds Luo & Hu (2021); Zhou et al. (2021); Nichol et al. (2022), voxel grids Hui et al. (2022); Tang et al. (2023b); Müller et al. (2023), Triplanes Chen et al. (2023); Wang et al. (2023b); Shue et al. (2023), 3D Gaussians Zhang et al. (2024a). More recently, 3D latent diffusion has been explored to directly learn the mapping between the image and 3D geometry Zhang et al. (2023); Zhao et al. (2024); Li et al. (2024b); Zhang et al. (2024b); Wu et al. (2024b); Li et al. (2025a); Zhao et al. (2025); Ye et al. (2025), which greatly improves the generation quality. However, multi-view 3D generation is still under-explored, suffering from high variations in generation, easy inconsistency with input images, or strong reliance on input viewpoints Xiang et al. (2024); Zhao et al. (2025), which hinders their direct application in accurate 3D object reconstruction.

**Multi-view 3D Reconstruction.** Traditional methods conduct multi-view stereo (MVS) to reconstruct the visible surface of objects by triangulating correspondences across multiple calibrated images Furukawa et al. (2015); Galliani et al. (2015); Schönberger et al. (2016); Xu & Tao (2019). Learning-based MVS methods Yao et al. (2018; 2019); Chen et al. (2019); Cheng et al. (2020); Gu et al. (2020); Yang et al. (2020); Wang et al. (2021a) employ deep neural networks to enhance both reconstruction quality and computational efficiency. Scene-specific NeRF methods Lin et al. (2021); Wang et al. (2021b) adopt bundle adjustment from conventional SfM pipelines to jointly optimize camera parameters along with radiance field from dense views. Recently, DUSt3R Wang et al. (2024a) and its follow-up works Smart et al. (2024); Leroy et al. (2024); Wang et al. (2025a) together estimate point clouds and camera poses from paired or more views, which releases the reliance on camera parameters, but suffers from incomplete reconstruction results caused by the point cloud representation. Focusing on object reconstructions, large reconstruction models Hong et al. (2023) are explored to produce complete reconstructions via regressing a more compact or structured 3D representation (e.g. 3D Gaussians Kerbl et al. (2023) and Triplane) from multi-view inputs Wei et al. (2024); Xu et al. (2024b); Tang et al. (2025); Xu et al. (2024d), but requiring view inputs from certain camera poses. Follow-up methods further support pose-free reconstructions Wu et al. (2023a); Wang et al. (2023a); Jiang et al. (2023); He et al. (2024), while they tend to predict smooth and blurred details, especially in invisible regions. Differently, our method introduces diffusion-based 3D generation priors to advance pose-free object reconstruction in fidelity and completeness.

**Generative Priors in 3D Object Reconstruction.** Generative priors are introduced into 3D reconstruction frameworks to assist in predicting plausible geometries or textures in invisible portions of objects. Existing methods mainly introduce two kinds of priors: (i) diffusion-based 2D generative prior and (ii) regression-based 3D generative prior. The former is often used in single-view 3D reconstruction by generating plausible multi-view images first and conducting reconstruction Li et al. (2023; 2024a); Xu et al. (2024b); Tang et al. (2025); Xu et al. (2024d); Wang et al. (2025b); Wei et al. (2024); Liu et al. (2023a; 2024); Wu et al. (2024a). For pose-free sparse-view reconstruction, iFusion Wu et al. (2023a) leverages Zero123 Liu et al. (2023b) predictions within an optimization

162 pipeline to align poses and generate novel views for reconstruction. However, the inconsistency  
 163 between views still limits the performance of this pipeline. Regression-based 3D generative priors  
 164 are introduced to regress a unified compact 3D representation, avoiding this issue, for example 3D  
 165 neural volume Jiang et al. (2023), Triplane Hong et al. (2023); Wei et al. (2024); Wang et al. (2023a),  
 166 and 3D Gaussians He et al. (2024); Xu et al. (2024c); Smart et al. (2024). Diffusion-based gener-  
 167 ative priors prove superior to regressive ones in generating detailed results in both geometry and  
 168 texture Li et al. (2025a); Zhao et al. (2025); Zhang et al. (2024b); Ye et al. (2025). One2345++ and  
 169 its follow-up work Liu et al. (2024); Xu et al. (2024a) develop 3D volume diffusion conditioned by  
 170 multi-view inputs. However, 3D volume suffer from poor compactness, so a trade-off between the  
 171 diffusion learning difficulty and representation capability limits their performance. Differently, our  
 172 method builds upon strong diffusion-based 3D generative priors Xiang et al. (2024), with powerful  
 173 reconstruction priors Wang et al. (2025a) constraining the denoising process for accurate 3D outputs  
 174 of high-fidelity details.

### 175 3 METHODOLOGY

#### 176 3.1 PRELIMINARY

177 Given a set of  $N$  uncalibrated multi-view images of an object  $I = \{I_i\}_{i=1}^N$ , the task of pose-free  
 178 multi-view reconstruction aims to obtain the complete 3D object  $O$ . Our framework leverages two  
 179 kinds of strong priors to achieve complete and accurate reconstruction results: the reconstruction  
 180 prior from VGGT Wang et al. (2025a) and the generation prior from TRELLIS Xiang et al. (2024).  
 181 In this section, we first introduce these two priors as preliminaries.

182 **Reconstruction prior of VGGT** VGGT Wang et al. (2025a) achieves SOTA results in pose-  
 183 free multi-view 3D reconstruction, providing a powerful reconstruction prior. It adopts a feed-  
 184 forward transformer architecture designed for efficient and unified 3D scene reconstruction from  
 185 single/multiple images. Multi-view images  $I$  are first fed into a DINO-based ViT Oquab et al.  
 186 (2024) simultaneously for tokenization and feature extraction into  $\phi_{\text{dino}}$ . Then, 24 self-attention  
 187 layers further address  $\phi_{\text{dino}}$  into 3D-aware features  $\{\phi_i\}_{i=1}^{24}$  with an alternating attention strategy,  
 188 switching between frame-wise and global self-attention to balance local and global information and  
 189 enhance multi-view consistency. Finally, four prediction heads decode the output of 4 layers (4-th,  
 190 11-th, 17-th, and 23-rd), *i.e.*,  $\phi_{\text{vggt}}(I) = \{\phi_4, \phi_{11}, \phi_{17}, \phi_{24}\}$ , into camera parameters, depth map,  
 191 point map, and tracking feature predictions. To adapt to object reconstruction, we fine-tune VGGT  
 192 on an object-reconstruction dataset (see Sec. 4.1 for details). A LoRA fine-tuning on the VGGT  
 193 aggregator is employed to preserve the pre-trained 3D geometric priors, with a multi-task objective:  
 194

$$\mathcal{L}_{\text{VGGT}}(\theta) = \mathcal{L}_{\text{camera}} + \mathcal{L}_{\text{depth}} + \mathcal{L}_{\text{nmap}}, \quad (1)$$

195 where  $\theta$  is the LoRA parameters,  $\mathcal{L}_{\text{camera}}$ ,  $\mathcal{L}_{\text{depth}}$  and  $\mathcal{L}_{\text{nmap}}$  denote the camera pose loss, the depth  
 196 loss, and the point map loss, respectively. In the following text, we simply use “VGGT” to refer to  
 197 this fine-tuned VGGT.

198 **Generation prior of TRELLIS** TRELLIS Xiang et al. (2024) is a SOTA 3D generative model  
 199 that provides a strong generation prior. It proposes a novel representation called Structured LATent  
 200 (SLAT) that combines a sparse 3D grid with dense visual features extracted from a powerful vision  
 201 foundation model, which captures both geometric (structure) and textural (appearance) information  
 202 and enables decoding into multiple 3D representations. We choose TRELLIS as the 3D generator in  
 203 our framework because it has shown great potential in 3D object generation He et al. (2025); Li et al.  
 204 (2025b) and inspired many works in downstream applications Yang et al. (2025); Cao et al. (2025);  
 205 Wu et al. (2024c). It employs a coarse-to-fine two-stage generation pipeline: generating the sparse  
 206 structure (SS), represented as sparse voxels  $\{p_i\}_i^L$ , via SS Flow and then predicting structured latents  
 207 (SLAT) for active SS voxels, represented as  $X = \{(p_i, x_i)\}_i^V$ , via SLAT Flow, where  $p_i$ ,  $x_i$ , and  $V$   
 208 denotes the voxel position, the latent vector, and the number of voxels, respectively. The generation  
 209 in both stages adopts rectified flow transformers Liu et al. (2022) with DINO-encoded image features  
 210 as conditions. The result of SLAT Flow is then decoded into 3D outputs represented by radiance  
 211 fields (RF), 3D Gaussians (3DGS), or meshes, *i.e.*,  $O = \text{Dec}(x)$ . Modeling the backward process  
 212 as a time-dependent vector field  $v(x, t) = \nabla_t(x)$ , the transformers  $v_\theta$  in both stages are trained by  
 213 minimizing the conditional flow matching (CFM) objective Lipman et al. (2023):  
 214

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t, x_0, \epsilon} \|v_\theta(x, t) - (\epsilon - x_0)\|_2^2. \quad (2)$$

216 **Overview** Our ReconViaGen framework conducts reconstruction and generation simultaneously  
 217 and utilizes the two priors in a complementary fashion. It builds upon TRELLIS to generate complete  
 218 3D outputs with strong generation priors to plausibly hallucinate invisible portions to com-  
 219 pensate for the limitation of reconstruction. The proposed ReconViaGen adopts a coarse-to-fine  
 220 reconstruction pipeline. As shown in Fig. 2, in the first stage, we use a pre-trained VGGT to provide  
 221 reconstruction-based multi-view conditions at both the global and local levels. In the next stage,  
 222 we respectively feed the global geometry and local per-view conditions into the SS and SLAT Flow  
 223 transformers, for multi-view-aware generation. Finally, we further refine the estimated camera poses  
 224 from VGGT using the generation and introduce pixel-level alignment constraints only in the infer-  
 225 ence stage for reconstructions highly consistent with input views in detailed geometry and textures.  
 226

### 227 3.2 RECONSTRUCTION-BASED CONDITIONING

228 We first introduce reconstruction priors in VGGT to provide strong multi-view-aware conditions for  
 229 the coarse and detailed shape and texture generation of TRELLIS.

230 **Global Geometry Condition** VGGT learns a strong reconstruction prior to encode explicit 3D  
 231 lifting information into multi-view image features. Therefore, we first aggregate VGGT features  
 232  $\phi_{\text{vggt}}$  into a global geometry representation, serving as SS Flow conditions to generate more accurate  
 233 coarse structures. Note that we did not use explicit reconstruction results like point clouds because  
 234 VGGT features convey richer information including camera poses, depth, point maps, and tracking.  
 235 A fixed-length token list  $T_g$  is aggregated from  $\phi_{\text{vggt}}$  via a proposed Condition Net design shown  
 236 in Fig. 2. Starting from a randomly initialized learnable token list  $T_{\text{init}}$ , four transformer cross-  
 237 attention blocks progressively fuse layer-wise features of  $\phi_{\text{vggt}}$  with the initial token list and produce  
 238  $T_g$ . Formulated as:

$$239 T^{i+1} = \text{CrossAttn}(Q(T^i), K(\phi_{\text{vggt}}), V(\phi_{\text{vggt}})), \quad i \in \{0, 1, 2, 3\}, \quad (3)$$

240 where  $T^0$  is initialized with  $T_{\text{init}}$ ,  $T^3$  is the final output  $T_g$ ,  $Q(\cdot)$ ,  $K(\cdot)$ , and  $V(\cdot)$  are linear layers  
 241 respectively for query, key, and value projection, and  $\phi_{\text{vggt}}$  is the VGGT features that concatenate  
 242 all views on the token dimension. At the training stage of SS Flow, we freeze the VGGT layers and  
 243 train the Condition Net together with DiT.

244 **Local Per-View Condition** A single token list condition can provide limited fine-grained infor-  
 245 mation for geometry and texture generation in detail. We further adopt the Condition Net design to  
 246 provide local per-view tokens as SLAT Flow conditions for fine-grained generation in both geom-  
 247 etric and texture details. A random token list is initialized for each view and fed into the Condition  
 248 Net to produce a view-specific token list  $\mathcal{P}_k, k \in [1, N]$ :

$$249 P_k^{i+1} = \text{CrossAttn}(Q(\mathcal{P}_k^i), K(\phi_k^{\text{vggt}}), V(\phi_k^{\text{vggt}})), \quad i \in \{0, 1, 2, 3\} \text{ and } k \in \{n\}_{n=1}^N, \quad (4)$$

250 where  $\phi_k^{\text{vggt}}$  is VGGT features of the  $k$ -th view. The set of  $\{\mathcal{P}_k\}_{k=1}^N$  is sent into SLAT diffusion  
 251 transformers offering per-view object appearance guidance for fine-grained generation.

### 254 3.3 COARSE-TO-FINE GENERATION

255 The overall generation process consists of three stages: (i) coarse structure generation via SS Flow  
 256 with global geometry condition; (ii) fine detail generation via SLAT Flow with local per-view con-  
 257 dition; (iii) rendering-aware pixel-aligned refinement at the inference stage only.

258 **Reconstruction-conditioned Flow** To integrate the reconstruction prior into generation, the two  
 259 stages of SS and SLAT Flow in TRELLIS take the global geometry condition  $T_g$  and local per-  
 260 view conditions  $\{\mathcal{P}_k\}_{k=1}^N$  for coarse and fine diffusion guidance, respectively. In the first stage, we  
 261 simply compute the cross-attention between the condition  $T_g$  and the noisy SS latent in each SS DiT  
 262 block. In the second stage, as illustrated in Fig. 2, we encourage the cross-attention between the  
 263 noisy SLAT and each view’s condition  $\mathcal{P}_k$  and conduct a weighted fusion in each SLAT DiT block,  
 264 which can be formulated as:

$$266 y_{j+1} = \sum_{k=1}^N \text{CrossAttn}(Q(y'_j), K(\mathcal{P}_k), V(\mathcal{P}_k)) \cdot w_k, \quad j \in \{m\}_{m=1}^M, \quad (5)$$

267 where  $M$  is the number of SLAT DiT blocks,  $y'_j$  is the self-attention layer output of the noisy SLAT  
 268 input  $y_j$ , and  $w_k \in (0, 1)$  is the fusion weight computed via an MLP taking the cross-attention result

270 as input. After the first two stages, the 3D generator can generate multi-view-aware geometry and  
271 texture at both the global and local level.

272 **Rendering-aware Velocity Compensation** To further encourage pixel-aligned consistency be-  
273 tween generation results and input views, we develop a rendering-aware velocity compensation to  
274 constrain the diffusion trajectory according to inputs. In doing so, we first estimate camera pose  
275 with VGGT using the generation results from the second stage, with detailed implementation de-  
276 tails included in the appendix. Inspired by the explicit normal regularization used in Hi3DGen Ye  
277 et al. (2025) to improve the input-output consistency, when  $t < 0.5$ , we decode the SLAT into  $O_t$   
278 (e.g. a textured mesh) and conduct rendering for alignment. The SLAT Flow process initializes and  
279 updates a large number of noisy latents for all voxels simultaneously, which results in a challenging  
280 collaborative optimization problem. To solve this issue, we novelly propose a mechanism called  
281 Rendering-aware Velocity Compensation (RVC) to correct the predicted  $v$  for a more accurate gen-  
282 eration consistent with input views. Specifically, we render images for  $O_t$  from the refined camera  
283 pose estimations  $C$  and calculate the difference between the rendered images and input images as:

$$\mathcal{L}_{\text{RVC}}(v_t) = \mathcal{L}_{\text{SSIM}} + \mathcal{L}_{\text{LPIPS}} + \mathcal{L}_{\text{DreamSim}}, \quad (6)$$

284 where  $\mathcal{L}_{\text{SSIM}}$ ,  $\mathcal{L}_{\text{LPIPS}}$ , and  $\mathcal{L}_{\text{DreamSim}}$  are SSIM Wang et al. (2004), LPIPS Zhang et al. (2018), and  
285 DreamSim losses Fu et al. (2023) (inspired by the practice in V2M4 Chen et al. (2025)), respon-  
286 sible for measuring the structural, perceptual, and semantic similarity, respectively. To exclude the  
287 influence of inaccurate pose estimation, we discard the losses corresponding to some images if their  
288 corresponding losses are higher than 0.8. By minimizing  $\mathcal{L}_{\text{RVC}}$ , we iteratively correct the predicted  
289 velocity in each SLAT denoising step with a compensation term  $\Delta v$ , derived as:

$$\Delta v_t = \frac{\partial \mathcal{L}}{\partial \hat{x}_0} \frac{\partial \hat{x}_0}{\partial v_t} = -t \frac{\partial \mathcal{L}}{\partial \hat{x}_0}, \quad (7)$$

290 where  $\mathcal{L}$  represents  $\mathcal{L}_{\text{RVC}}$  for simplicity and  $\hat{x}_0$  is the predicted target SLAT at current timestep  $t$ ,  
291 computed as  $\hat{x}_0 = x_t - t \cdot v_t$ . The noisy SLAT of next step  $x_{t_{\text{prev}}}$  can be updated as:

$$x_{t_{\text{prev}}} = x_t - (t - t_{\text{prev}})(v + \alpha \cdot \Delta v), \quad (8)$$

292 where  $\alpha$  is a pre-defined hyperparameter that controls the extent of the compensation. In this way,  
293 the input images serve as a strong explicit guidance to find a denoising trajectory for each local  
294 SLAT vector, which leads to more accurate 3D results consistent with all input images in detail.

## 300 4 EXPERIMENTS

### 301 4.1 EXPERIMENT SETUP

302 **Datasets** For LoRA fine-tuning of the VGGT aggregator and Trellis sparse structure transformer,  
303 we employ 390k 3D data from the Objaverse dataset Deitke et al. (2024), a large-scale 3D object  
304 dataset that provides a rich variety of shapes and textures, with 60 views rendered per object for  
305 fine-tuning. For each object mesh, we render 150 view images in a resolution of 512×512 under  
306 uniform lighting conditions following TRELLIS Xiang et al. (2024). For evaluation, we selected  
307 two benchmark datasets to thoroughly assess the performance of our model: (i) Dora-Bench Chen  
308 et al. (2024), a benchmark organized based on 4 levels of complexity, combining 3D data selected  
309 from the Objaverse Deitke et al. (2023), ABO Collins et al. (2022), and GSO Downs et al. (2022)  
310 dataset; and (ii) OmniObject3D, a large-vocabulary 3D object dataset containing 6,000 high-quality  
311 textured meshes scanned from real-world objects, covering 190 daily categories. We randomly  
312 sample 300 objects from Dora-Bench and 200 objects covering 20 categories from OmniObject3D.  
313 We follow He et al. (2024) to render 24 views at different elevations, and randomly chose 4 of  
314 them as multi-view input for evaluation on OmniObject3D. On Dora-Bench, we follow the camera  
315 trajectory of TRELLIS Xiang et al. (2024) to render 40 views and choose 4 views (No.0, 9, 19, and  
316 29) with a uniform interval to adapt to the setting of some baseline methods (LGM Tang et al. (2025)  
317 and InstantMesh Xu et al. (2024b)).

318 **Evaluation Metrics** We employ PSNR, SSIM, and LPIPS to evaluate the accuracy of synthesized  
319 novel views from 3D outputs, Chamfer Distance (CD) and F-score to evaluate the generated geo-  
320 metry accuracy and completeness. PSNR, SSIM, and LPIPS are evaluated on novel views of images  
321 rendered at the resolution of 512×512. CD and F-score are evaluated by sampling 100k points  
322 from the 3D outputs (using the center positions for 3D Gaussian outputs), with all object points  
323 normalized to the range of  $[-1, 1]^3$ . When calculating the F-score, the radius  $r$  is set to 0.1.

324

325

Table 1: Evaluation on the Dora-bench and OminiObject3D dataset. Best results are in **bold**.

Method	Dora-bench					OminiObject3D				
	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
VGGT Wang et al. (2025a)	-	-	-	0.112	0.921	-	-	-	0.091	0.900
TRELLIS-S Xiang et al. (2024)	16.562	0.876	0.103	0.176	0.807	16.021	0.771	0.264	0.102	0.906
TRELLIS-M Xiang et al. (2024)	16.706	0.882	0.111	0.144	0.843	16.861	0.790	0.242	0.072	0.932
Hunyuan3D-2.0-mv Zhao et al. (2025)	20.221	0.896	0.093	0.094	0.937	16.665	0.813	0.165	0.124	0.871
LGM Tang et al. (2025)	17.877	0.869	0.186	0.121	0.839	16.361	0.791	0.193	0.136	0.842
InstantMesh Xu et al. (2024b)	18.922	0.870	0.120	0.110	0.865	17.499	0.818	0.145	0.094	0.907
LucidFusion He et al. (2024)	16.509	0.835	0.144	0.131	0.831	16.254	0.771	0.144	0.114	0.868
<b>ReconViaGen (Ours)</b>	<b>22.632</b>	<b>0.911</b>	<b>0.090</b>	<b>0.090</b>	<b>0.953</b>	<b>19.767</b>	<b>0.847</b>	<b>0.141</b>	<b>0.059</b>	<b>0.959</b>

332

333

**Baseline Methods** Baseline methods for comparisons include (i) TRELLIS-S Xiang et al. (2024): generates 3D meshes from multi-view images using TRELLIS in the stochastic mode, which randomly chooses one input view to condition each step of denoising; (ii) TRELLIS-M Xiang et al. (2024): TRELLIS in the multidiffusion mode, which computes the average denoised results conditioned on all input views; (iii) Hunyuan3D-2.0-mv Zhao et al. (2025): concatenate DINO features of input images from fixed viewpoints as conditions to generate meshes<sup>1</sup>; (iv) InstantMesh Xu et al. (2024b): predicts Triplane for mesh outputs from multiple images with fixed viewpoints; (v) LGM Tang et al. (2025): predicts pixel-aligned 3D Gaussians from multiple images with fixed viewpoints; (vi) LucidFusion He et al. (2024): predicts relative coordinate maps for 3D Gaussian outputs; (vii) VGGT Wang et al. (2025a): reconstructs the point cloud from multi-view inputs in a feed-forward manner. We compare our methods with a wide range of existing SOTA baseline methods: (a) 3D generation models {i, ii, iii}; (b) large reconstruction models with known camera poses {iv, v}; (c) pose-free large reconstruction models with 3DGS or point cloud outputs {vi, vii}. For 3D generation models {i, ii, iii}, we use the same approach as Camera Pose Estimation in fine detail reconstruction to align the generated 3D models to the ground-truth models. Besides, we also compare with closed-source commercial 3D generation models like Hunyuan3D-2.5 and Meshy-5 on in-the-wild testing.

349

**Implementation Details** For LoRA fine-tuning of VGGT aggregator and TRELLIS transformer, we set the rank as 64, the alpha parameter for LoRA scaling as 128, and the dropout probability for LoRA layers as 0. We only apply the adapter to qkv mapping layer and the projectors of each attention layer. During fine-tuning VGGT aggregator, we randomly sample 1 ~ 4 views from 150 images and use the AdamW optimizer with a fixed learning rate of  $1 \times 10^{-4}$ . For the fine-tuning of SS Flow and Slat-Flow transformer, we build upon TRELLIS Xiang et al. (2024), incorporating classifier-free guidance (CFG) with a drop rate of 0.3 and an AdamW optimizer with a fixed learning rate of  $1 \times 10^{-4}$ . We fine-tune the SS-Flow transformer using 8 NVIDIA A800 GPUs (80GB memory) for 40k steps with a batch size of 192. **Differently, we finetune the SLat-Flow transformer with a batch size of 128.** During inference, we set the CFG strengths in SS generation and SLAT generation to 7.5 and 3.0, and use 30 and 12 sampling steps to achieve optimal results. The  $\alpha$  in rendering-aware velocity compensation is set to 0.1 in our practice.

360

361

## 4.2 EXPERIMENT RESULTS

362

363

**Quantitative Results** We present the quantitative comparisons between our ReconViaGen and other baseline methods in Tab. 1 for evaluation on the Dora-bench and OminiObject3D dataset. The proposed method achieves consistently superior performance to other methods on both image-reconstruction consistency (PSNR, SSIM, and LPIPS), geometry accuracy (CD), and shape completeness (F-score). Impressively, our ReconViaGen seamlessly integrate the generation and reconstruction priors from TRELLIS Xiang et al. (2024) and VGGT Wang et al. (2025a), whose performance surpasses both of them. Note that VGGT performs better on Dora-bench than on OminiObject3D because uniformly-distributed views can capture richer visual cues than random views. Besides, our method also gets better results than previous SOTA pose-free multi-view reconstruction methods that integrate regression-based generation priors by a large margin, especially on PSNR, CD, and F-score, which validates the superiority of ReconViaGen. On the settings of more input views, we separately evaluate VGGT that can accept an arbitrary number of inputs for comparison, which is included in the appendix. We also present the camera pose estimation accuracy in the appendix.

377

<sup>1</sup>The fresh version, Hunyuan3D-2.5, has not been open-sourced, which is unsuitable for large-scale evaluation on benchmarks, so we use the open-sourced version, Hunyuan3D-2.0.

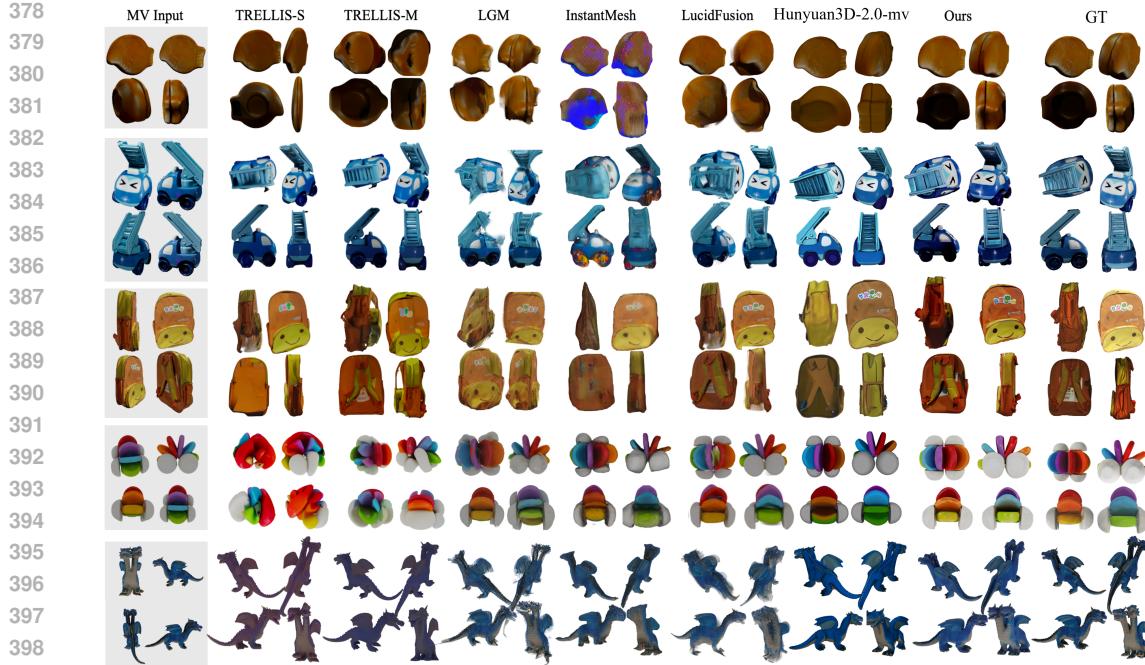


Figure 3: Reconstruction result comparisons between our ReconViaGen and other baseline methods on samples from the Dora-bench and OminiObject3D datasets. Zoom in for better visualization.



Figure 4: Reconstruction results on in-the-wild samples. Note that commercial 3D generators require input images from orthogonal viewpoints, while ours can accept views from arbitrary camera poses for robust outputs. Zoom in for better visualization in detail.

**Qualitative Results** We further present extensive qualitative comparisons to demonstrate the superiority of our ReconViaGen. We first select some examples from the OminiObject3D and Dora-bench dataset for visualization, as shown in Fig. 3. The reconstruction results of ReconViaGen have the most accurate geometry and textures compared to other methods. We further evaluate several baseline methods on in-the-wild multi-view images. As shown in Fig. 4, our ReconViaGen exhibits strong robustness even in comparison with the multi-view version of closed-source commercial 3D

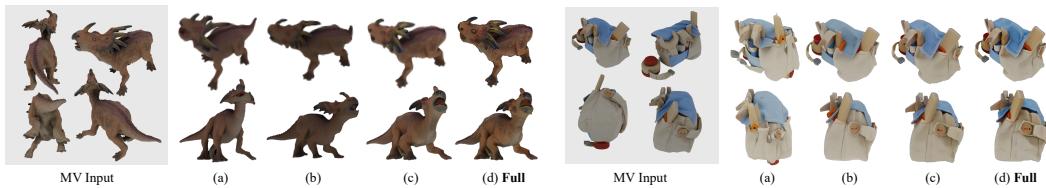


Figure 5: Qualitative comparisons for different variants of ReconViaGen for ablative study. Zoom in for better visualization in detail.

generation models like Hunyuan3D-2.5 and Meshy-5. More qualitative results are included in the appendix.

#### 4.3 ABLATION STUDY

The proposed ReconViaGen framework comprises three novel designs to integrate reconstruction priors into the diffusion-based 3D generation: (i) the global geometry condition (GGC); (ii) the per-view condition (PVC); and (iii) the rendering-aware velocity compensation (RVC). We conduct ablation studies to validate the individual effectiveness of each component. On the Dora-bench dataset, we start from a basic TRELLIS-M baseline (ReconViaGen without all designs, Tab. 2a) and progressively add one component, leading to 3 variants (b,c,d). As shown in Tab. 2, integrating GGC, which strongly improves the prediction accuracy of coarse structure, brings a large performance gain on almost all metrics. Further integrating PVC can lead to extra improvement, especially on PSNR, which proves the effectiveness in improving local per-view alignment. Finally, adopting RVC, though in the inference stage only, brings additional increments in both shape completeness and fine-grained accuracy in geometry and texture. Qualitative comparisons in Fig. 5 visualize the positive effect of each component: global geometry conditioning greatly corrects the global shape, per-view conditioning produces local details in geometry and texture of high consistency with each view, and rendering-aware velocity compensating impressively refines the fine-grained appearance, leading to high-quality results. More ablation results on the detailed designs, including the number of image inputs and the choice of condition form for SS and SLAT Flow, can be seen in the appendix.

Table 2: Quantitative ablation results on the Dora-bench dataset.

	GGC	PVC	RVC	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
(a)	✗	✗	✗	16.706	0.882	0.111	0.144	0.843
(b)	✓	✗	✗	20.462	0.894	0.102	0.093	0.941
(c)	✓	✓	✗	21.045	0.905	0.093	0.093	0.937
(d)	✓	✓	✓	<b>22.632</b>	<b>0.911</b>	<b>0.090</b>	<b>0.090</b>	<b>0.953</b>

## 5 CONCLUSION

In this paper, we have presented ReconViaGen, a novel coarse-to-fine framework that effectively integrates strong reconstruction priors with diffusion-based 3D generative priors for accurate and complete multi-view 3D object reconstruction. We first analyze the inherent reasons leading to the challenge of leveraging diffusion-based 3D generative priors into reconstruction: insufficient cross-view correlation modeling and stochastic denoising process with weak constraint from input images. Therefore, we effectively use powerful reconstruction priors with three novelly designed mechanisms to enhance the multi-view correlation awareness in 3D diffusion learning and establish strong constraints for a reliable denoising process. Extensive experiments have demonstrated that ReconViaGen achieves SOTA performance in both global shape accuracy and completeness as well as local details in geometry and textures. As future work, with the development of 3D reconstruction and 3D generation, stronger reconstruction or generation priors can be integrated into our framework to further improve reconstruction quality via generation.

486 REFERENCES  
487

488 Ziang Cao, Zhaoxi Chen, Liang Pan, and Ziwei Liu. Physx-3d: Physical-grounded 3d asset genera-  
489 tion. *arXiv preprint arXiv:2507.12465*, 2025.

490 Hansheng Chen, Jiatao Gu, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Hao Su. Single-  
491 stage diffusion nerf: A unified approach to 3d generation and reconstruction. In *ICCV*, pp. 2416–  
492 2425, 2023.

493 Jianqi Chen, Biao Zhang, Xiangjun Tang, and Peter Wonka. V2m4: 4d mesh animation reconstruc-  
494 tion from a single monocular video. *arXiv preprint arXiv:2503.09631*, 2025.

495 Rui Chen, Songfang Han, Jing Xu, and Hao Su. Point-based multi-view stereo network. In *ICCV*,  
496 pp. 1538–1547, 2019.

497 Rui Chen, Jianfeng Zhang, Yixun Liang, Guan Luo, Weiyu Li, Jiarui Liu, Xiu Li, Xiaoxiao Long,  
498 Jiashi Feng, and Ping Tan. Dora: Sampling and benchmarking for 3d shape variational auto-  
499 encoders. *arXiv preprint arXiv:2412.17808*, 2024.

500 Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, and Hao Su.  
501 Deep stereo using adaptive thin volume representation with uncertainty awareness. In *CVPR*, pp.  
502 2524–2534, 2020.

503 Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu,  
504 Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, et al. Abo: Dataset and  
505 benchmarks for real-world 3d object understanding. In *CVPR*, pp. 21126–21136, 2022.

506 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig  
507 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-  
508 tated 3d objects. In *ICCV*, pp. 13142–13153, 2023.

509 Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan  
510 Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of  
511 10m+ 3d objects. *NeurIPS*, 36, 2024.

512 Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann,  
513 Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of  
514 3d scanned household items. In *ICRA*, pp. 2553–2560, 2022.

515 Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting  
516 with applications to image analysis and automated cartography. *Communications of the ACM*, 24  
517 (6):381–395, 1981. doi: 10.1145/358669.358692.

518 Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and  
519 Phillip Isola. Dreamsim: Learning new dimensions of human visual similarity using synthetic  
520 data. volume 36, pp. 50742–50768, 2023.

521 Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and*  
522 *trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.

523 Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by  
524 surface normal diffusion. In *ICCV*, pp. 873–881, 2015.

525 Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost  
526 volume for high-resolution multi-view stereo and stereo matching. In *CVPR*, pp. 2495–2504,  
527 2020.

528 Hao He, Yixun Liang, Luozhou Wang, Yuanhao Cai, Xinli Xu, Hao-Xiang Guo, Xiang Wen, and  
529 Yingcong Chen. Lucidfusion: Generating 3d gaussians with arbitrary unposed images. *arXiv*  
530 *preprint arXiv:2410.15636*, 2024.

531 Xianglong He, Zi-Xin Zou, Chia-Hao Chen, Yuan-Chen Guo, Ding Liang, Chun Yuan, Wanli  
532 Ouyang, Yan-Pei Cao, and Yangguang Li. Sparseflex: High-resolution and arbitrary-topology  
533 3d shape modeling. *arXiv preprint arXiv:2503.21732*, 2025.

540 Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli,  
 541 Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. In *ICLR*, 2023.  
 542

543 Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural wavelet-domain diffusion for 3d shape  
 544 generation. In *SIGGRAPH Asia*, pp. 1–9, 2022.

545 Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-  
 546 view stereopsis evaluation. In *2014 IEEE Conference on Computer Vision and Pattern Recogni-  
 547 tion*, pp. 406–413. IEEE, 2014.

548 Hanwen Jiang, Zhenyu Jiang, Yue Zhao, and Qixing Huang. Leap: Liberate sparse-view 3d model-  
 549 ing from camera poses. *arXiv preprint arXiv:2310.01410*, 2023.

550 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-  
 551 ting for real-time radiance field rendering. *TOG*, 42(4):139–1, 2023.

552 Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnp: An accurate o(n) solution to the  
 553 pnp problem. *IJCV*, 81(2):155–166, 2009. doi: 10.1007/s11263-008-0152-6.

554 Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r.  
 555 In *European Conference on Computer Vision*, pp. 71–91. Springer, 2024.

556 Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan  
 557 Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view gen-  
 558 eration and large reconstruction model. *arXiv preprint arXiv:2311.06214*, 2023.

559 Peng Li, Yuan Liu, Xiaoxiao Long, Feihu Zhang, Cheng Lin, Mengfei Li, Xingqun Qi, Shanghang  
 560 Zhang, Wenhan Luo, Ping Tan, et al. Era3d: High-resolution multiview diffusion using efficient  
 561 row-wise attention. *arXiv preprint arXiv:2405.11616*, 2024a.

562 Weiyu Li, Jiarui Liu, Rui Chen, Yixun Liang, Xuelin Chen, Ping Tan, and Xiaoxiao Long. Crafts-  
 563 man: High-fidelity mesh generation with 3d native generation and interactive geometry refiner.  
 564 *arXiv preprint arXiv:2405.14979*, 2024b.

565 Yangguang Li, Zi-Xin Zou, Zexiang Liu, Dehu Wang, Yuan Liang, Zhipeng Yu, Xingchao Liu,  
 566 Yuan-Chen Guo, Ding Liang, Wanli Ouyang, et al. Triposg: High-fidelity 3d shape synthesis  
 567 using large-scale rectified flow models. *arXiv preprint arXiv:2502.06608*, 2025a.

568 Zhihao Li, Yufei Wang, Heliang Zheng, Yihao Luo, and Bihan Wen. Sparc3d: Sparse representa-  
 569 tion and construction for high-resolution 3d shapes modeling. *arXiv preprint arXiv:2505.14521*,  
 570 2025b.

571 Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural  
 572 radiance fields. In *ICCV*, pp. 5741–5751, 2021.

573 Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten  
 574 Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content  
 575 creation. In *CVPR*, pp. 300–309, 2023.

576 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow match-  
 577 ing for generative modeling. In *ICLR*, 2023.

578 Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-  
 579 2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *NeurIPS*,  
 580 36, 2023a.

581 Minghua Liu, Ruoxi Shi, Linghao Chen, Zhuoyang Zhang, Chao Xu, Xinyue Wei, Hansheng Chen,  
 582 Chong Zeng, Jiayuan Gu, and Hao Su. One-2-3-45++: Fast single image to 3d objects with  
 583 consistent multi-view generation and 3d diffusion. In *CVPR*, pp. 10072–10083, 2024.

584 Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick.  
 585 Zero-1-to-3: Zero-shot one image to 3d object. In *ICCV*, pp. 9298–9309, 2023b.

586 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and  
 587 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

594 Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *CVPR*,  
 595 pp. 2837–2845, 2021.

596

597 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and  
 598 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pp.  
 599 405–421, 2020.

600 Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulo, Peter Kontschieder, and  
 601 Matthias Nießner. Diffff: Rendering-guided 3d radiance field diffusion. In *CVPR*, pp. 4328–  
 602 4338, 2023.

603

604 Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system  
 605 for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.

606

607 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,  
 608 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning  
 609 robust visual features without supervision. *Transactions on Machine Learning Research Journal*,  
 pp. 1–31, 2024.

610

611 Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d  
 612 diffusion. *arXiv preprint arXiv:2209.14988*, 2022.

613

614 Lingteng Qiu, Guanying Chen, Xiaodong Gu, Qi Zuo, Mutian Xu, Yushuang Wu, Weihao Yuan,  
 615 Zilong Dong, Liefeng Bo, and Xiaoguang Han. Richdreamer: A generalizable normal-depth  
 616 diffusion model for detail richness in text-to-3d. In *CVPR*, pp. 9914–9925, 2024.

617

618 Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view  
 619 selection for unstructured multi-view stereo. In *ECCV*, pp. 501–518. Springer, 2016.

620

621 J Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 3d  
 622 neural field generation using triplane diffusion. In *CVPR*, pp. 20875–20886, 2023.

623

624 Brandon Smart, Chuanxia Zheng, Iro Laina, and Victor Adrian Prisacariu. Splatt3r: Zero-shot  
 625 gaussian splatting from uncalibrated image pairs. *arXiv preprint arXiv:2408.13912*, 2024.

626

627 Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative  
 628 gaussian splatting for efficient 3d content creation. In *ICLR*, 2024.

629

630 Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm:  
 631 Large multi-view gaussian model for high-resolution 3d content creation. In *ECCV*, pp. 1–18.  
 632 Springer, 2025.

633

634 Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-  
 635 it-3d: High-fidelity 3d creation from a single image with diffusion prior. In *ICCV*, pp. 22819–  
 22829, 2023a.

636

637 Zhicong Tang, Shuyang Gu, Chunyu Wang, Ting Zhang, Jianmin Bao, Dong Chen, and Baining  
 638 Guo. Volumediffusion: Flexible text-to-3d generation with efficient volumetric encoder. *arXiv  
 639 preprint arXiv:2312.11459*, 2023b.

640

641 Tencent Hunyuan3D Team. Hunyuan3d 2.0: High-resolution 3d asset generation. <https://3d-models.hunyuan.tencent.com>, 2025.

642

643 Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patch-  
 644 matchnet: Learned multi-view patchmatch stereo. In *CVPR*, pp. 14194–14203, 2021a.

645

646 Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David  
 647 Novotny. Vggt: Visual geometry grounded transformer. *arXiv preprint arXiv:2503.11651*, 2025a.

648

649 Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexi-  
 650 ang Xu, and Kai Zhang. Pf-lrm: Pose-free large reconstruction model for joint pose and shape  
 651 prediction. *arXiv preprint arXiv:2311.12024*, 2023a.

652

653 Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Ge-  
 654 ometric 3d vision made easy. In *CVPR*, pp. 20697–20709, 2024a.

648 Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen,  
 649 Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3d digital  
 650 avatars using diffusion. In *CVPR*, pp. 4563–4573, 2023b.

651

652 Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Pro-  
 653 lificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation.  
 654 *NeurIPS*, 36, 2024b.

655 Zhengyi Wang, Yikai Wang, Yifei Chen, Chendong Xiang, Shuo Chen, Dajiang Yu, Chongxuan Li,  
 656 Hang Su, and Jun Zhu. Crm: Single image to 3d textured mesh with convolutional reconstruction  
 657 model. In *ECCV*, pp. 57–74. Springer, 2025b.

658

659 Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment:  
 660 From error visibility to structural similarity. *TIP*, 13(4):600–612, 2004. doi: 10.1109/TIP.2004.  
 661 827628.

662 Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf–: Neural  
 663 radiance fields without known camera parameters. 2021b.

664

665 Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli,  
 666 Hao Su, and Zexiang Xu. Meshlrm: Large reconstruction model for high-quality mesh. *arXiv*  
 667 preprint *arXiv:2404.12385*, 2024.

668 Chin-Hsuan Wu, Yen-Chun Chen, Bolivar Solarte, Lu Yuan, and Min Sun. ifusion: Inverting diffu-  
 669 sion for pose-free reconstruction from sparse views. *arXiv preprint arXiv:2312.17250*, 2023a.

670

671 Kailu Wu, Fangfu Liu, Zhihan Cai, Runjie Yan, Hanyang Wang, Yating Hu, Yueqi Duan, and  
 672 Kaisheng Ma. Unique3d: High-quality and efficient 3d mesh generation from a single image.  
 673 *arXiv preprint arXiv:2405.20343*, 2024a.

674 Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Jingxi Xu, Philip Torr, Xun Cao, and Yao Yao.  
 675 Direct3d: Scalable image-to-3d generation via 3d latent diffusion transformer. *arXiv preprint*  
 676 *arXiv:2405.14832*, 2024b.

677

678 Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Yikang Yang, Yajie Bao, Jiachen Qian, Siyu  
 679 Zhu, Xun Cao, Philip Torr, et al. Direct3d-s2: Gigascale 3d generation made easy with spatial  
 680 sparse attention. *arXiv preprint arXiv:2505.17412*, 2025.

681 Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi  
 682 Wang, Chen Qian, et al. Omniobject3d: Large-vocabulary 3d object dataset for realistic percep-  
 683 tion, reconstruction and generation. In *CVPR*, pp. 803–814, 2023b.

684

685 Zhennan Wu, Yang Li, Han Yan, Taizhang Shang, Weixuan Sun, Senbo Wang, Ruihai Cui, Weizhe  
 686 Liu, Hiroyuki Sato, Hongdong Li, et al. Blockfusion: Expandable 3d scene generation using  
 687 latent tri-plane extrapolation. *TOG*, 43(4):1–17, 2024c.

688 Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin  
 689 Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv*  
 690 preprint *arXiv:2412.01506*, 2024.

691

692 Chao Xu, Ang Li, Linghao Chen, Yulin Liu, Ruoxi Shi, Hao Su, and Minghua Liu. Sparp: Fast 3d  
 693 object reconstruction and pose estimation from sparse views. In *ECCV*, pp. 143–163. Springer,  
 694 2024a.

695

696 Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh:  
 697 Efficient 3d mesh generation from a single image with sparse-view large reconstruction models.  
 698 *arXiv preprint arXiv:2404.07191*, 2024b.

699

700 Jiale Xu, Shenghua Gao, and Ying Shan. Freesplatter: Pose-free gaussian splatting for sparse-view  
 701 3d reconstruction. *arXiv preprint arXiv:2412.09573*, 2024c.

701 Qingshan Xu and Wenbing Tao. Multi-scale geometric consistency guided multi-view stereo. In  
 702 *CVPR*, pp. 5483–5492, 2019.

702 Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and  
 703 Gordon Wetzstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and  
 704 generation. *arXiv preprint arXiv:2403.14621*, 2024d.

705 Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli,  
 706 Gordon Wetzstein, Zexiang Xu, et al. Dmv3d: Denoising multi-view diffusion using 3d large  
 707 reconstruction model. In *ICLR*, 2024e.

708 Jiayu Yang, Wei Mao, Jose M Alvarez, and MiaoMiao Liu. Cost volume pyramid based depth  
 709 inference for multi-view stereo. In *CVPR*, pp. 4877–4886, 2020.

710 Xianghui Yang, Huiwen Shi, Bowen Zhang, Fan Yang, Jiacheng Wang, Hongxu Zhao, Xinhai Liu,  
 711 Xinzhou Wang, Qingxiang Lin, Jiaao Yu, et al. Hunyuan3d 1.0: A unified framework for text-to-  
 712 3d and image-to-3d generation. *arXiv preprint arXiv:2411.02293*, 2024.

713 Yunhan Yang, Yufan Zhou, Yuan-Chen Guo, Zi-Xin Zou, Yukun Huang, Ying-Tian Liu, Hao Xu,  
 714 Ding Liang, Yan-Pei Cao, and Xihui Liu. Omnipart: Part-aware 3d generation with semantic  
 715 decoupling and structural cohesion. *arXiv preprint arXiv:2507.06165*, 2025.

716 Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstruc-  
 717 tured multi-view stereo. In *ECCV*, pp. 767–783, 2018.

718 Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for  
 719 high-resolution multi-view stereo depth inference. In *CVPR*, pp. 5525–5534, 2019.

720 Chongjie Ye, Yushuang Wu, Ziteng Lu, Jiahao Chang, Xiaoyang Guo, Jiaqing Zhou, Hao Zhao,  
 721 and Xiaoguang Han. Hi3dgen: High-fidelity 3d geometry generation from images via normal  
 722 bridging. *arXiv preprint arXiv:2503.22236*, 3, 2025.

723 Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape  
 724 representation for neural fields and generative diffusion models. *TOG*, 42(4):1–16, 2023.

725 Bowen Zhang, Yiji Cheng, Jiaolong Yang, Chunyu Wang, Feng Zhao, Yansong Tang, Dong Chen,  
 726 and Baining Guo. Gaussiancube: Structuring gaussian splatting using optimal transport for 3d  
 727 generative modeling. *arXiv e-prints*, pp. arXiv–2403, 2024a.

728 Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan  
 729 Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d  
 730 assets. *TOG*, 43(4):1–20, 2024b.

731 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable  
 732 effectiveness of deep features as a perceptual metric. In *CVPR*, pp. 586–595, 2018.

733 Zibo Zhao, Wen Liu, Xin Chen, Xianfang Zeng, Rui Wang, Pei Cheng, Bin Fu, Tao Chen, Gang Yu,  
 734 and Shenghua Gao. Michelangelo: Conditional 3d shape generation based on shape-image-text  
 735 aligned latent representation. *NeurIPS*, 36, 2024.

736 Zibo Zhao, Zeqiang Lai, Qingxiang Lin, Yunfei Zhao, Haolin Liu, Shuhui Yang, Yifei Feng,  
 737 Mingxin Yang, Sheng Zhang, Xianghui Yang, et al. Hunyuan3d 2.0: Scaling diffusion models for  
 738 high resolution textured 3d assets generation. *arXiv preprint arXiv:2501.12202*, 2025.

739 Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel  
 740 diffusion. In *ICCV*, pp. 5826–5835, 2021.

741 Qi Zuo, Xiaodong Gu, Lingteng Qiu, Yuan Dong, Zhengyi Zhao, Weihao Yuan, Rui Peng, Siyu Zhu,  
 742 Zilong Dong, Liefeng Bo, et al. Videomv: Consistent multi-view generation based on large video  
 743 generative model. *arXiv preprint arXiv:2403.12010*, 2024.

744

745

746

747

748

749

750

751

752

753

754

755

756 **A APPENDIX**  
757758 **A.1 DETAILS ON CAMERA POSE ESTIMATION**  
759760 For better alignment with the input images, we register them into the TRELLIS generation space.  
761 Specifically, we first render 30 images from randomly sampled camera views on a sphere, concate-  
762 nate them with the input images, and feed them into VGGT for pose estimation. Since the camera  
763 poses of the rendered views are known, we can recover coarse camera poses for the input images  
764 in the TRELLIS space. While VGGT provides robust pose predictions, they remain insufficiently  
765 accurate for constructing pixel-level rendering constraints.  
766767 To refine the results, we render images and depth maps using the coarse poses, then apply an image  
768 matching method to establish 2D-2D correspondences between rendered and input images. Lever-  
769 aging the depth maps and camera parameters of rendered views, we further obtain 2D-3D corre-  
770 spondences between each input image and the generated object. By aggregating multi-view corre-  
771 spondences, we solve for refined camera poses  $C$  using a PnP Lepetit et al. (2009) solver with  
772 RANSAC Fischler & Bolles (1981). This image-matching-based refinement effectively corrects the  
773 initial pose predictions from TRELLIS’s generative priors, yielding higher accuracy. The refined  
774 poses enable pixel-wise constraints from the input views, thereby supporting finer detail alignment  
775 in generation.  
776777 **A.2 EVALUATION WITH MORE INPUT IMAGES**  
778779 **Table 3: Evaluation with more input images on the Dora-bench dataset. Best results are in bold.**780 

Method	Uniform (PSNR↑ / LPIPS↓)			Limited View (PSNR↑ / LPIPS↓)		
	6 views	8 views	10 views	6 views	8 views	10 views
Object VGGT + 3DGS	18.476/0.123	19.890/0.109	21.363/0.102	16.498/0.139	16.774/0.135	17.121/0.133
<b>ReconViaGen (Ours)</b>	<b>22.823/0.089</b>	<b>23.067/0.090</b>	<b>23.193/0.087</b>	<b>21.427/0.098</b>	<b>21.782/0.099</b>	<b>21.866/0.103</b>

781 To thoroughly evaluate and validate the effectiveness of ReconViaGen, we compare it against 3DGS  
782 reconstruction initialized with point clouds and camera poses from object VGGT (denoted as object  
783 VGGT + 3DGS) on the Dora-Bench dataset. We conduct experiments under two input scenar-  
784 os: uniformly and limited-view sampled views. As shown in Tab. 3, ReconViaGen consistently  
785 outperforms object VGGT+3DGS at 6/8/10 input views, regardless of the sampling strategy. This  
786 advantage arises because the generative prior in ReconViaGen plays a crucial role in completing  
787 invisible regions of the object.  
788789 **A.3 EVALUATION OF CAMERA POSE ESTIMATION**  
790791 **Table 4: Evaluation of camera pose estimation on the Dora-bench dataset. Best results are in bold.**792 

Method	RRE↓	Acc. @ 15° ↑	Acc. @ 30° ↑	TE↓
VGGT Wang et al. (2025a)	8.575	90.67	92.00	0.066
Object VGGT	<b>7.257</b>	93.44	94.11	0.055
Ours	7.925	<b>93.89</b>	<b>96.11</b>	<b>0.046</b>

793 To assess the performance of our finetuned object VGGT and the effectiveness of our proposed  
794 camera pose estimation strategy, we evaluate pose prediction quality on the Dora-Bench dataset. We  
795 adopt both rotation and translation metrics: relative rotation error (RRE, in degrees), the proportion  
796 of RRE values below 15° and 30°, and translation error (TE), measured as the distance between pre-  
797 dicted and ground-truth camera centers. For evaluation, we use four input images and transform both  
798 predicted and ground-truth poses into the coordinate system of the first image, which is excluded  
799 from the metric computation. To address translation scale ambiguity, we compute relative transla-  
800 tions between views for both predictions and ground truth and normalize them by their respective  
801 mean L2-norm. As reported in Tab. 4, the finetuned object VGGT achieves clear improvements over  
802 the original VGGT. Our method further delivers the best overall performance, as the generative prior  
803 effectively ‘densifies’ sparse views. However, our RRE is slightly higher than that of object VGGT,  
804 likely due to minor discrepancies between the generated 3D model and the ground-truth geometry.  
805

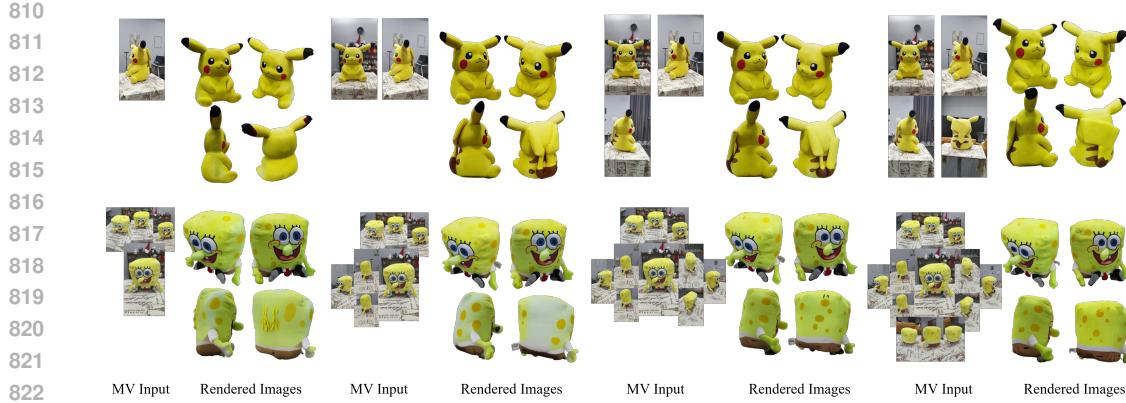


Figure 6: Qualitative comparisons for different numbers of input images with ReconViaGen. Zoom in for better visualization in detail.

#### A.4 ABLATION STUDY ON THE NUMBER OF INPUT IMAGES

Table 5: Quantitative ablation results of the number of input images on the Dora-bench dataset.

Number of Images	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
1	18.438	0.887	0.106	0.135	0.838
2	19.568	0.894	0.099	0.131	0.867
4	22.632	0.911	0.090	0.090	0.953
6	22.823	0.912	<b>0.089</b>	0.084	0.958
8	<b>23.067</b>	<b>0.914</b>	0.090	<b>0.081</b>	<b>0.961</b>

Since ReconViaGen can take an arbitrary number of input images, a natural question is how reconstruction quality scales with the number of images. To investigate this, we conducted an ablation study varying the number of input views on Dora-Bench, with results summarized in Tab. 5. We observe that reconstruction performance consistently improves as more images are provided. However, the marginal gains gradually diminish, indicating a saturation effect when the number of views becomes large. The visualization results are shown in Fig. 6, which also shows that our ReconViaGen can process any number of input images from any viewpoint.

#### A.5 ABLATION STUDY ON THE FORM OF CONDITION

Table 6: Quantitative ablation results of condition at SS Flow on the Dora-bench dataset.

Form of Condition	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
(i) Feature Volume	16.229	0.858	0.126	0.172	0.814
(ii) Concatenation	19.749	0.871	0.137	0.121	0.873
(iii) PVC	19.878	0.882	0.135	0.120	0.870
(iv) GGC	<b>20.462</b>	<b>0.894</b>	<b>0.102</b>	<b>0.093</b>	<b>0.941</b>

For SS Flow, as described in the method section, we explored several strategies to leverage VGGT features for sparse structure generation on Dora-Bench: (i) Downsampling the point cloud from VGGT to a  $64^3$  resolution occupancy volume, projecting DINO features from each view into the volume, and averaging them to form a feature-volume condition; (ii) Fusing VGGT features with DINO features for each view through several linear layers, then concatenating all input-view tokens as condition; (iii) adopting the same local per-view condition (PVC) used in our SLAT Flow; (iv) employing the proposed global geometry condition (GGC). For fair comparison, we use the original SLat Flow in TRELLIS and train all models for 40k steps. As shown in Tab. 6, our GGC achieves the best performance among all strategies. We attribute this to the limitations of the alternative designs: for (i), inaccurate predicted poses or point clouds lead to erroneous projections, introducing noise into the condition and harming generation; for (ii) and (iii), view-level features are not effectively aggregated, resulting in redundancy and making the model overly dependent on the accuracy of VGGT outputs.

864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917

Table 7: Quantitative ablation results of condition at SLAT Flow on the Dora-bench dataset.

Form of Condition	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
(i) GGC	17.784	0.858	0.120	0.097	0.939
(ii) PVC	<b>22.632</b>	<b>0.911</b>	<b>0.090</b>	<b>0.090</b>	<b>0.953</b>

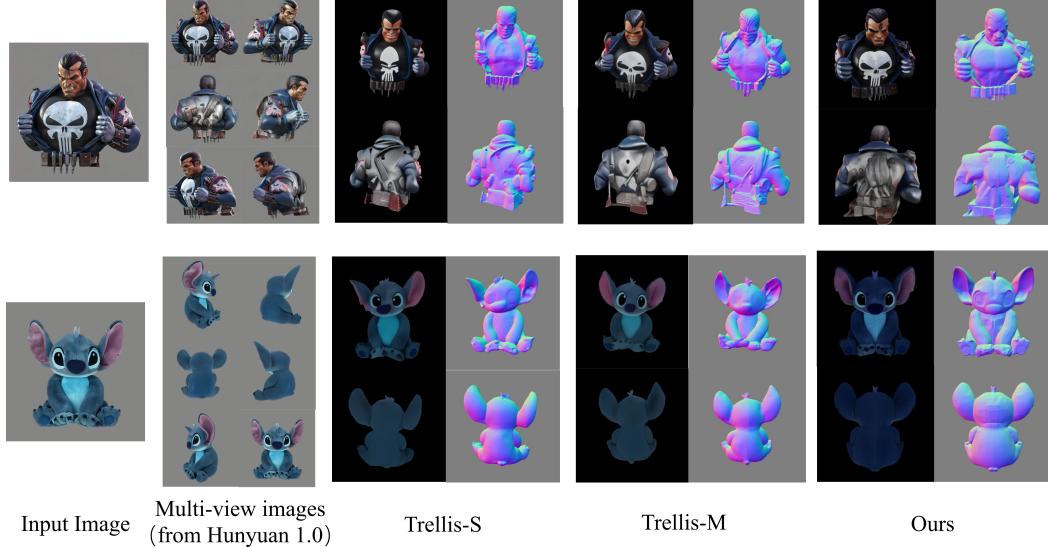


Figure 7: Reconstruction result comparisons between TRELLIS-M, TRELLIS-S, and our ReconVi-aGen on samples produced by the multi-view image generator.

For SLat Flow, we conduct an ablation study on Dora-Bench with two conditioning strategies: (i) the same global geometry condition (GGC) used in SS Flow, and (ii) the local per-view condition (PVC). For fairness, we pair both variants with SS Flow conditioned on GGC and train all models for 40k steps. As shown in Tab. 7, PVC substantially outperforms GGC in SLat Flow. We attribute this to the information compression in GGC, which leads to a loss of fine-grained details in the condition and degrades performance. This observation also explains why we adopt PVC instead of GGC for SLat Flow.

The following description provides an interpretability of GGC and PVC. GGC performs global multi-view aggregation, which strengthens overall structural consistency—this is exactly what SS needs, since SS focuses on coarse geometry where global cues dominate and multi-view compression is acceptable. In contrast, SLAT requires fine-grained appearance and local geometry refinement on a fixed coarse shape. PVC’s per-view feature interaction preserves high-frequency cues that GGC would smooth out. Therefore, GGC naturally aligns with global structure formation, while PVC is better suited for detail-aware refinement.

#### A.6 RECONSTRUCTION ON GENERATED MULTI-VIEW IMAGES OR VIDEOS

Table 8: Quantitative comparison of generated multi-view images on the Dora-bench dataset.

Method	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
ours (generated 6-view)	14.379	0.808	0.226	0.190	0.723
ours (single-view)	18.438	0.887	0.106	0.135	0.838
ours (real 6-view)	<b>22.823</b>	<b>0.912</b>	<b>0.089</b>	<b>0.084</b>	<b>0.958</b>

Given the growing interest in multi-view image generation using large image and video generative models, we further evaluate the robustness of our approach on such generated data. These multi-view images are hallucinated from a single view and often suffer from cross-view inconsistencies in fine details. Specifically, we generate 6-view samples using the open-sourced multi-view generator Hunyuan3D-1.0 Yang et al. (2024). The quantitative comparisons on Dora-Bench are summarized in Tab. 8. In our observations, the results of generated 6 views as input are much worse than those

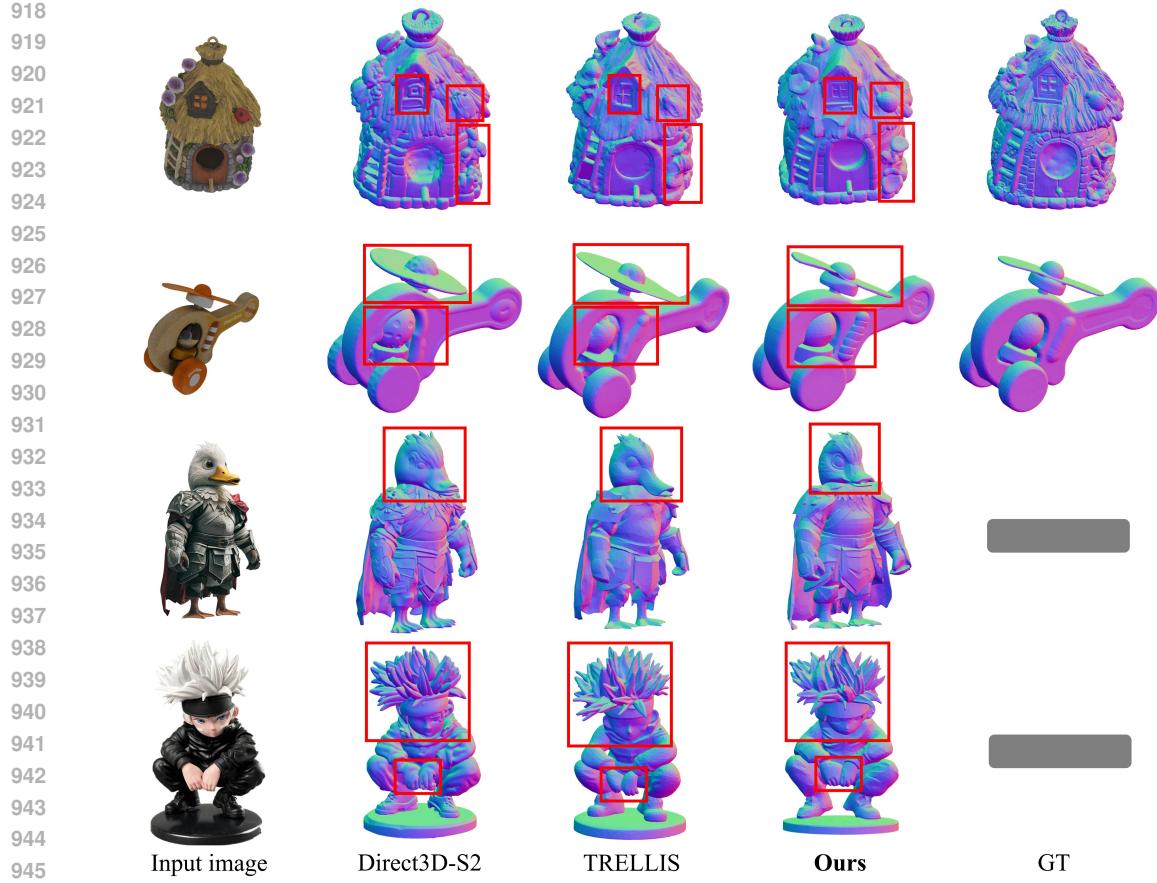


Figure 8: Qualitative result of our ReconViaGen and other baseline methods with single-view input on samples from the Dora-bench and in-the-wild scenarios. Since in-the-wild cases lack corresponding ground truth meshes, these ground truth meshes are not displayed in the figure. Zoom in for better visualization.

of single-view and real 6 view input due to much severe inconsistency in multi-view generation. However, when the inconsistency in generated multi-view images is not severe, visualizations in Fig. 7 show that ReconViaGen exhibits strong robustness to moderate cross-view inconsistency. Please refer to the supplementary video for additional results on generated videos.

#### A.7 MORE RECONSTRUCTION RESULTS

We further showcase our method on in-the-wild data, including not only multiple objects but also scenes, even from generated dynamic object videos. For scene reconstruction, we segment individual objects, reconstruct them separately, and then register the reconstructed 3D objects back into the scene using our predicted camera poses. Please refer to the supplementary video for qualitative results.

#### A.8 COMPARISON EXPERIMENT OF SINGLE-VIEW INPUT

Table 9: Quantitative results of single view input on the Dora-bench dataset.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CD $\downarrow$	F-score $\uparrow$
TRELLIS Xiang et al. (2024)	15.264	0.858	0.182	0.162	0.781
Direct3D-S2 Wu et al. (2025)	-	-	-	0.165	0.805
<b>Ours</b>	<b>18.438</b>	<b>0.887</b>	<b>0.106</b>	<b>0.135</b>	<b>0.838</b>

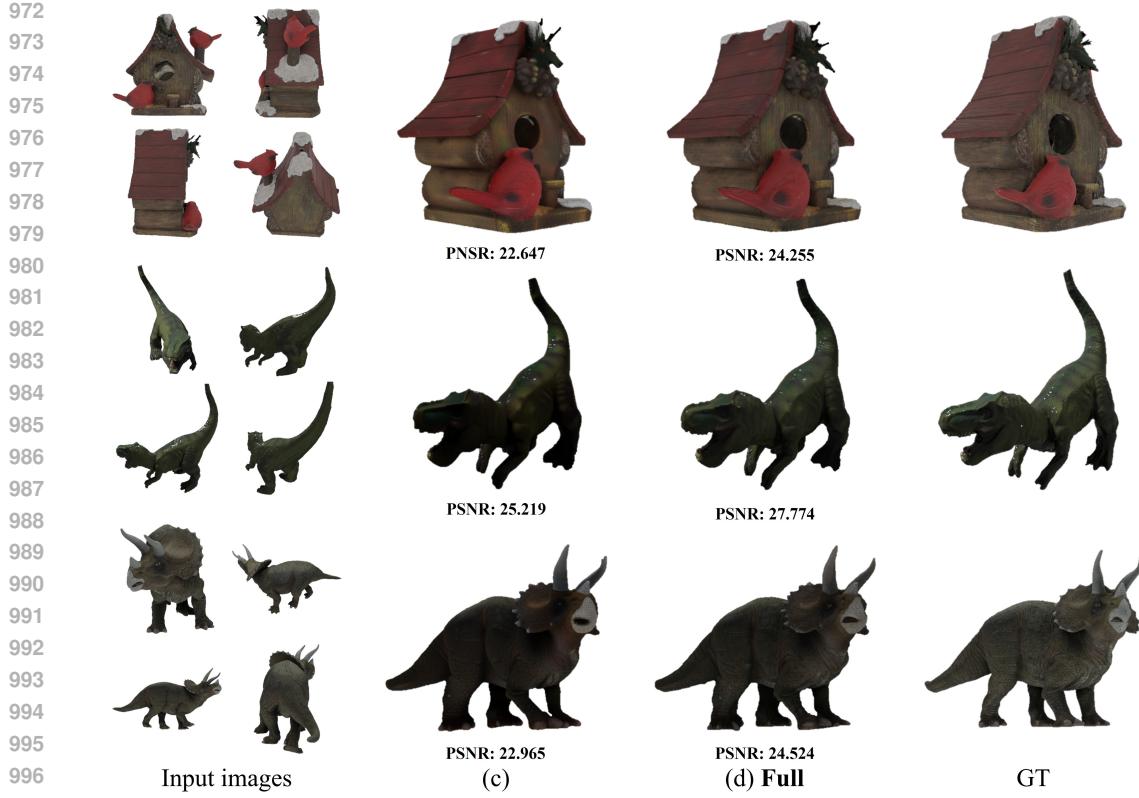


Figure 9: More qualitative comparisons for (c) vs (d) ablation study. Each case is labeled with a corresponding PSNR value. Zoom in for better visualization in detail.

To further demonstrate the superiority of our ReconViaGen, we conduct comparison experiments with single-view input. We select TRELLIS Xiang et al. (2024) and Direct3D-S2 Wu et al. (2025) as representative single-view generation baselines for comparison. For evaluation on the Dora-Bench dataset, we use a single randomly selected view as input. Quantitative comparisons between our ReconViaGen and the baseline methods are summarized in Tab. 9. Note that since Direct3D-S2 Wu et al. (2025) outputs geometry without color, we do not report its visual metrics. The quantitative data demonstrate that ReconViaGen consistently achieves superior performance across both visual and geometric evaluation criteria. We further provide extensive qualitative comparisons on both Dora-Bench and in-the-wild scenarios, as shown in Fig. 8. The visual evidence confirms that our proposed method consistently outperforms the baseline approaches, yielding reconstructed geometry and textures that exhibit the strongest fidelity and alignment with the input view.

#### A.9 MORE QUALITATIVE EXAMPLES FOR THE (C) VS (D) ABLATION STUDY

We present more qualitative examples and label the PSNR of each case for the (c) vs (d) ablation study in Fig. 9. The qualitative analysis confirms that the 1.587 dB improvement is primarily due to RVC’s ability to correct color and texture drift and enforcing high-frequency aligned with the input views.

#### A.10 ABLATION STUDY ON RVC HYPERPARAMETERS

The RVC stage in ReconViaGen involves three main hyperparameters: the extent of RVC  $\alpha$ , timestep  $t$ , and outlier rejection threshold  $t_o$ . Then we conduct extensive ablation experiments on these three hyperparameters on the Dora-Bench dataset. As shown in Tab. 10, for  $\alpha$ , varying  $\alpha$  from 1 to 0.01, we observe stable performance across a wide range, with the best results around  $\alpha = 0.1$ . Extremely small values slightly reduce accuracy, but no instability was observed. Importantly,  $\alpha$

1026

1027

Table 10: **Ablation study on the extent of RVC  $\alpha$  on the Dora-bench dataset.**

1028

1029

1030

1031

1032

1033

1034

1035

1036

Table 11: **Ablation study on timestep  $t$  in RVC on the Dora-bench dataset.**

1037

1038

1039

1040

1041

1042

1043

1044

introduces no runtime overhead, as it only changes the update weight of a vector. As shown in Tab. 11, applying RVC at different timesteps shows that  $t = 0.5$  offers the best trade-off between fidelity and cost. As expected, runtime scales linearly with the fraction of steps with RVC (due to decoding/rendering), ranging from 4.8s ( $t = 0.3$ ) to 9.3s ( $t = 0.7$ ). As shown in Tab. 12, for outlier rejection threshold  $t_o$ , results are insensitive across a broad range. We choose  $t_o = 0.8$  as it yields strong PSNR/F-score. Therefore,  $\alpha = 0.1$ ,  $t = 0.5$ , and  $t_o = 0.8$  is the optimal combination of hyperparameters. Importantly, the best-performing configuration ( $\alpha = 0.1$ ,  $t = 0.5$ ,  $t_o = 0.8$ ) is shared across all assets and datasets, and no per-asset tuning is required. This confirms that the final RVC step generalizes well and does not rely on instance-specific adjustments.

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

### A.11 THE DETAILED LATENCY OF EACH COMPONENT IN RECONVIAGEN

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

To demonstrate the efficiency of our method in detail, we report the inference time for each component on an NVIDIA L20 GPU (46GB memory) with different numbers of input images in Tab. 13. As shown in the table, the dominant runtime contributors are pose estimation and RVC. Nevertheless, the total inference time remains reasonable. The RVC is a justified trade-off for accuracy and efficiency. RVC provides a significant quality boost of approximately 1.587 dB PSNR compared to the pipeline without it. This gain is crucial for achieving high-fidelity 3D reconstruction and eliminating visual inconsistencies to input images. Importantly, even without pose estimation and RVC, the results still surpass current SOTA reconstruction methods, while RVC further improves fine geometry and texture alignment.

1067

### A.12 THEORETICAL EXPLANATION ON HOW RVC INTERACTS WITH RECTIFIED FLOW OBJECTIVES

1068

1069

1070

1071

1072

1073

RVC does not modify the underlying rectified-flow training objective or the learned vector field  $v_\theta$ . Instead, it is an inference-time correction mechanism that adjusts the denoising trajectory according to rendering errors from the input views. Concretely, the rectified flow predicts the direction of the straightened transport path as:

$$x_{t-\Delta t} = x_t - \Delta t v_\theta(x_t, t). \quad (9)$$

RVC introduces a small correction:

$$x_{t-\Delta t} = x_t - \Delta t(v_\theta + \alpha \Delta v), \quad (10)$$

1074

1075

1076

1077

1078

1079

where  $\Delta v = -t \frac{\partial L_{\text{RVC}}}{\partial x_0}$ . This adjustment is orthogonal to the rectified-flow objective. The flow model is fully trained under the standard CFM loss. And RVC only guides the denoising trajectory toward a solution that is better aligned with input images. The correction does not alter the learned velocity field or its theoretical properties. In summary, RVC provides input-conditioned guidance during inference, but does not change the theoretical formulation or training of rectified flow.

1080

1081 Table 12: **Ablation study on the outlier rejection threshold  $t_o$  in RVC on the Dora-bench dataset.**

$t_o$	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
1	21.618	0.912	0.094	0.091	0.938
0.9	22.075	0.906	0.090	0.091	0.946
0.8	<b>22.632</b>	<b>0.911</b>	<b>0.090</b>	0.090	<b>0.953</b>
0.7	22.450	0.909	0.094	<b>0.089</b>	0.941
0.6	22.312	0.906	0.093	0.094	0.944

1088

1089

1090 Table 13: **Inference time (in seconds) of each component under varying numbers of input views.**

Component \Number of images	1	3	5	7	9
DINO feature + VGGT feature	0.2157s	0.4393s	0.8789s	1.3788s	2.0745s
GGC	0.0169s	0.0175s	0.0184s	0.0199s	0.0224s
SS Flow + SS Decoder	4.4193s	4.4039s	4.4105s	4.4141s	4.4267s
PVC	0.0017s	0.0019s	0.0022s	0.0023s	0.0025s
SLat Flow + SLat Decoder	2.7178s	3.0591s	3.7154s	3.9062s	4.4952s
Pose estimation	20.6085s	24.0013s	27.4166s	30.8672s	35.3407s
RVC	5.9726s	6.4080s	7.5483s	10.3359s	12.2985s
Overall	33.9525s	38.3310s	43.9903s	50.9244s	58.6605s

1099

1100

## 1101 A.13 ABLATION STUDY ON DECODING TO DIFFERENT 3D REPRESENTATIONS IN RVC

1102

1103 Table 14: **Ablation study on decoding to different 3D representations in RVC on the Dora-bench 1104 dataset.**

3d representation	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
Radiance Field (RF)	21.899	0.908	0.0910	0.0917	0.937
Mesh	21.561	0.906	0.0924	<b>0.0879</b>	<b>0.954</b>
3DGS	<b>22.632</b>	<b>0.911</b>	<b>0.0901</b>	0.0895	0.953

1109

1110 To investigate the influence of 3D representation in RVC, we decode the output SLAT to different  
 1111 3D representations and report their metrics on the Dora-Bench dataset. In Tab. 14, mesh shows  
 1112 slightly better geometric accuracy due to its structured nature. 3DGS provides superior rendering  
 1113 quality, leading to the best overall visual fidelity. Based on its superior rendering quality and overall  
 1114 balanced performance, choosing 3DGS as the representation is more appropriate.

1115

## 1116 A.14 ATTENTION VISUALIZATIONS IN GGC AND CROSS-VIEW TOKEN SIMILARITY IN PVC

1117

1118 To fully understand the proposed GGC and PVC, we provide attention visualization in GGC and  
 1119 cross-view token similarity in PVC. Specifically, for attention visualization in GGC, we first average  
 1120 attention maps along the dimension of multi-head in each attention layer, and then max pool the  
 1121 attention maps along the dimension of learnable tokens. Finally, we resize the attention map as  
 1122 the same size as input images. For cross-view token similarity in PVC, we calculate the similarity  
 1123 between tokens of two different input images and select a token from one image to visualize its  
 1124 similarity to tokens of another image. As shown in Fig. 10, for GGC’s global aggregation in SS,  
 1125 attention visualizations of the first three layers all show a strong correlation between the global  
 1126 learnable tokens and object features in images. In Fig. 11, for PVC’s local correspondence in SLAT,  
 1127 we find there is high similarity between corresponding points across different viewpoints.

1128

## A.15 ABLATION STUDY ON THE LENGTH OF LEARNABLE TOKENS IN GGC

1129

1130

1131 We conduct the ablation study on the length of learnable tokens in GGC. Since the length of SS latent  
 1132 is  $16^3 = 4096$ , we believe 4096 is enough for the length of learnable tokens in GGC to represent the  
 1133 visible structure. As shown in Tab. 15, the significant performance drop shows that 2048 tokens are  
 1134 insufficient to represent information of input. Although 8192 tokens offer a slight improvement in  
 1135 PSNR and CD, they introduce additional computational overhead. Therefore, 4096 learnable tokens  
 1136 provide the balance between performance and efficiency.

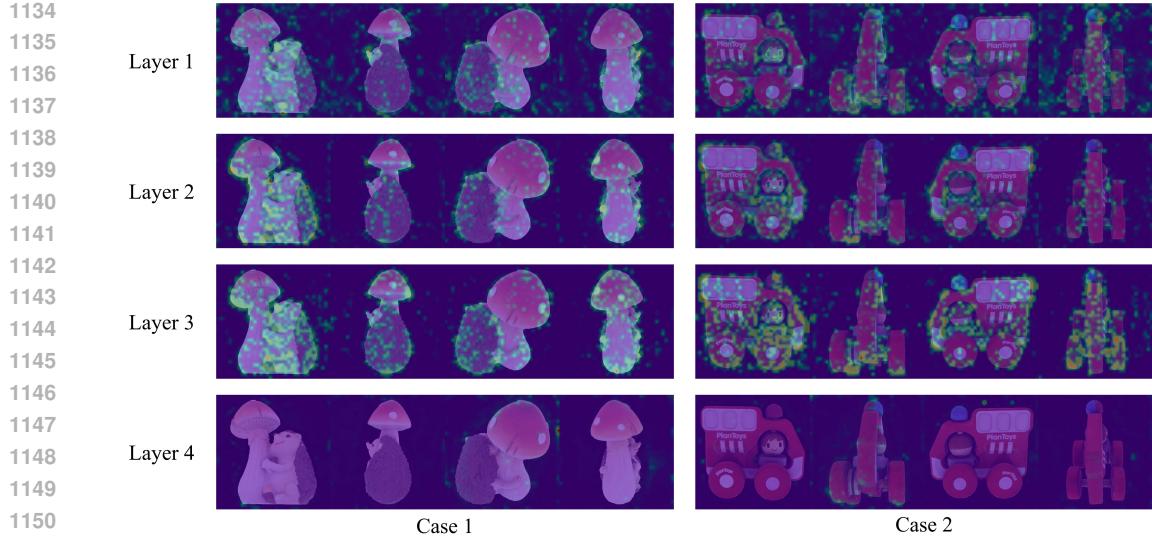


Figure 10: Attention visualization in GGC.

Table 15: Ablation study on the length of learnable tokens in GGC on the Dora-bench dataset.

Length of tokens	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
2048	18.366	0.890	0.107	0.114	0.884
4096	20.462	<b>0.894</b>	0.102	0.093	<b>0.941</b>
8192	<b>20.527</b>	0.887	<b>0.105</b>	<b>0.092</b>	0.938

## A.16 EVALUATION ON THE DTU DATASET

Table 16: Evaluation on the DTU dataset.

Method	PSNR↑	SSIM↑	LPIPS↓
TRELLIS-M Xiang et al. (2024)	15.8675	0.6010	0.3237
Hunyuan3D-2.0-mv Team (2025)	19.0951	0.6763	0.2654
Instantmesh Xu et al. (2024b)	18.8966	0.6712	0.2527
Ours	<b>21.7639</b>	<b>0.7576</b>	<b>0.2175</b>

DTU Jensen et al. (2014) is a real-world multi-view dataset. To further verify the generalizability of our method on real multi-view captures, we evaluate three SOTA multi-view reconstruction baselines (TRELLIS-M, InstantMesh, and Hunyuan3D-2.0-mv) and our method on the DTU dataset, and the results are summarized in Tab. 16. Our approach substantially improves reconstruction fidelity on real-world multi-view data (+2.7 PSNR over the strongest baseline). These results confirm that ReconViaGen generalizes well to real multi-view captures, without any dataset-specific tuning. Qualitative comparisons in Fig. 12 further demonstrate the superiority of our ReconViaGen in real multi-view captures.

## A.17 ABLATION STUDY ON RVC ONLY BASELINE WITH OUR FULL-VERSION MODEL

We conduct experiments of TRELLIS-M with RVC on Dora-Bench. As shown in the Tab. 17, RVC alone consistently improves the baseline, demonstrating that RVC is indeed effective. However, its performance remains far below our full model, since RVC mainly provides local refinement and cannot replace the global structure aggregation (GGC) and per-view detail modeling (PVC). This confirms that RVC is beneficial but complements rather than substitutes the other components.

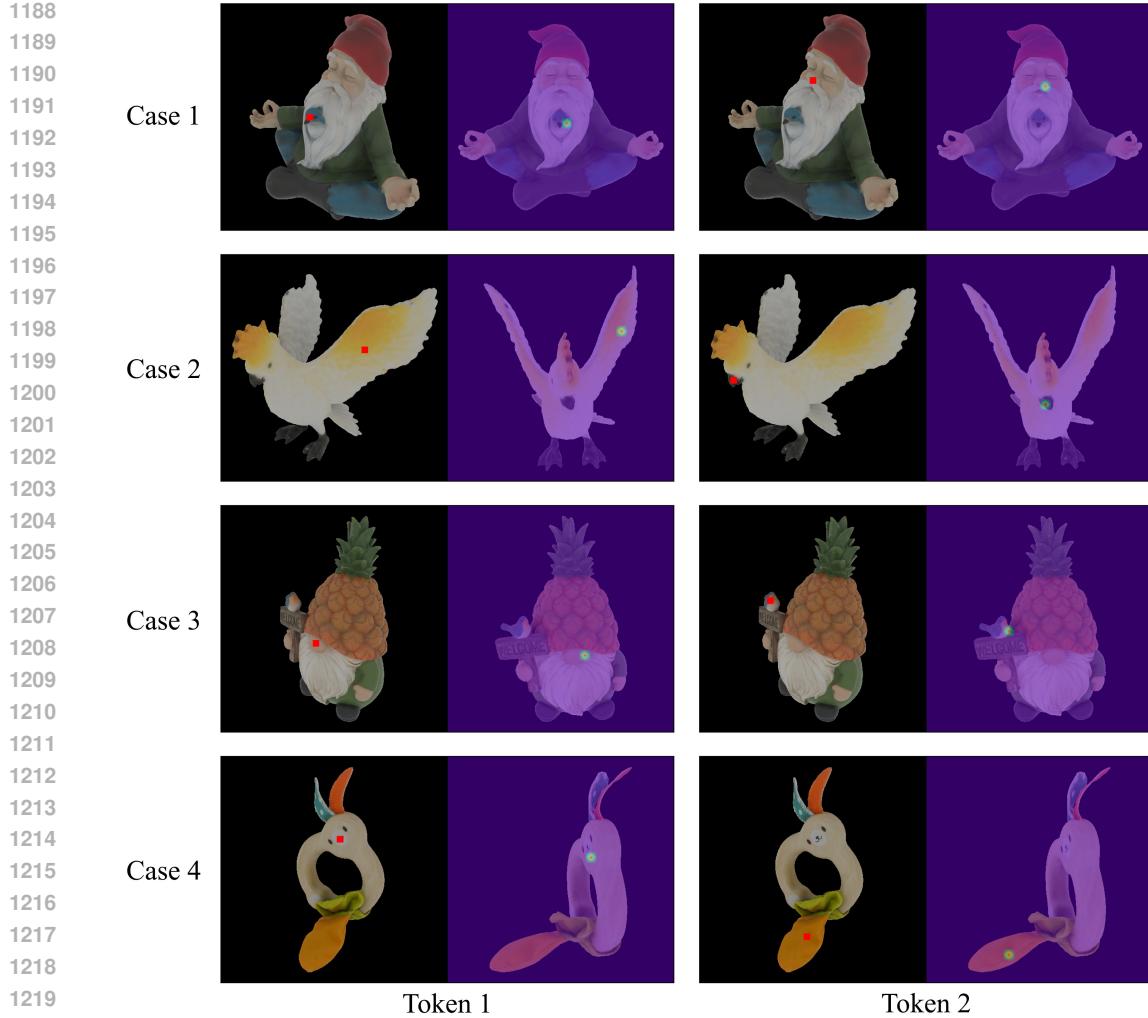


Figure 11: Cross-view token similarity in PVC.

Table 17: Ablation on RVC only baseline with our full-version model.

Method	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
TRELLIS-M Xiang et al. (2024)	16.706	0.882	0.111	0.144	0.843
TRELLIS-M Xiang et al. (2024) w. RVC	17.728	0.857	0.101	0.139	0.851
Ours	22.632	0.911	0.090	0.090	0.953

## A.18 ABLATION STUDY ON ALTERNATIVE DESIGN FOR WEIGHTED FUSION IN SLAT-FLOW

The weighted fusion design aims to trade off multi-view specificity and computational efficiency. The SLAT latent already encodes global geometry and acts as a “volume memory”. Therefore, the result of cross attention between SLAT latent and each individual view already reflects how much that view contributes to the current latent state. The MLP is then used to map this interaction to a scalar weight, enabling view-dependent importance estimation without requiring joint attention over all views. Another alternative design is to add all the cross attention results from all views together and then apply a MLP. We ablation on this alternative design with ours in Tab. 18. The alternative yields worse performance. We think direct summation reduces view-specific differences, confirming the need for a view-aware weighting scheme. A direct way is applying a cross-attention with the KV cache coming from all the views. We experimented with this design, but the KV cache grows linearly with the number of views and quickly exceeds GPU memory. Thus, per-view cross-attention with a learned weight provides the balance of accuracy, interpretability, and memory efficiency.



Figure 12: Qualitative comparisons on the DTU dataset.

Table 18: Ablation on alternative design for weighted fusion in SLat-Flow on the Dora-bench dataset.

Method	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
Add up + MLP	20.809	0.894	<b>0.091</b>	0.098	0.924
Our weighted fusion (c)	<b>21.045</b>	<b>0.905</b>	0.093	<b>0.093</b>	<b>0.937</b>

### A.19 CLARIFICATION FOR OMITTING THE VISUAL METRICS OF VGGT IN THE DORA-BENCH EVALUATION

Using point clouds from VGGT directly for image rendering does indeed produce poor results. For example, due to the discrete representation of point clouds, when rendering an object from the front, rays may pass through the front point clouds and hit the back point clouds instead. This is due to deficiencies in point cloud representation, therefore we believe that reporting visual metrics in this way is unfair to VGGT. To use the output of VGGT for image rendering, a direct way is to combine the VGGT with 3DGS optimization. We specifically compared ours with VGGT+3DGS in Sec.A.2. However, 3DGS optimization introduces many artifacts and floaters, which negatively impacts geometric metrics like CD and F-score, which is unfair to VGGT. 3D reconstruction often focuses more on geometry, so we choose to report only VGGT’s geometric metrics and omit its visual metrics.

### A.20 SENSITIVITY ANALYSIS OF THE QUALITY TO CAMERA POSE ACCURACY

Table 19: Ablation study on camera poses under different extent of rotational perturbation on Dora Bench.

Method	PSNR↑	SSIM↑	LPIPS↓	CD↓	F-score↑
Ours w/o RVC	21.045	0.905	0.093	0.093	0.937
Ours w. GT pose	<b>23.957</b>	<b>0.917</b>	0.092	<b>0.086</b>	0.955
Ours w. 3° perturbation	23.606	0.914	<b>0.089</b>	0.088	<b>0.957</b>
Ours w. 5° perturbation	22.788	0.916	0.093	0.092	0.951
Ours w. 10° perturbation	21.221	0.909	0.094	0.095	0.950
Ours w. 30° perturbation	20.939	0.903	0.094	0.093	0.941
Ours w. 50° perturbation	20.906	0.903	0.097	0.093	0.938
Ours	22.632	0.911	0.090	0.090	0.953

1296 To investigate the robustness and sensitivity of RVC to camera pose accuracy, we intentionally in-  
1297 troduce severe rotational perturbation to the ground-truth camera pose in RVC. As shown in Tab. 19,  
1298 RVC improves reconstruction significantly when accurate poses are available. With moderate errors  
1299 ( $\leq 10$ ), the degradation is minimal, showing high tolerance. Even with extreme perturbations  $30^\circ$   
1300 and  $50^\circ$  (far beyond practical estimation errors), the performance remains comparable to the baseline  
1301 "Ours w/o RVC", demonstrating that RVC does not destabilize or corrupt the denoising trajectory.  
1302 The key to this robustness lies in our outlier mechanism within the rendering-based loss.

1303

## 1304 A.21 THE USE OF LARGE LANGUAGE MODELS

1305

1306 We only utilize LLMs to refine the writing style and enhance the clarity of exposition. The LLMs  
1307 are not involved in research ideation, experimental design or data analysis.

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