

(U)NFV: (UN)SUPERVISED NEURAL FINITE VOLUME METHODS FOR SOLVING HYPERBOLIC PDEs

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
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
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
1

054 The solutions of hyperbolic PDEs are difficult to approximate due to discontinuities such as shocks,
 055 even when starting from smooth initial conditions (Evans, 2022). Consequently, classical (smooth)
 056 solutions typically cease to exist after finite time, and one must instead rely on weak solutions.
 057 Closed-form solutions exist only in rare cases, such as on simple Riemann problems LeVeque (2002)
 058 or through the Lax-Hopf formula (Lax, 1957; Claudel and Bayen, 2010a;b) in specific concave or
 059 convex settings. As a result, practical applications almost always rely on numerical methods for
 060 approximating the PDE solution, with finite volume (FV) methods (LeVeque, 2002) being a popular
 061 choice due to their ability to track conserved quantities across discontinuities and capture shock
 062 dynamics relatively accurately.

063 Classical FV methods involve important trade-offs between accuracy near discontinuities, compu-
 064 tational cost, stencil size, and implementation complexity. In recent years, neural networks have
 065 been explored as flexible and powerful alternatives solvers, showing promise in learning complex
 066 dynamics from data or residuals. Yet, many such methods are designed for non-specific models, often
 067 at the expense of losing physical structure, including conservation laws and entropy behaviors.

068 We introduce the *Neural Finite Volume* (NFV) method, a modular architecture tailored to conserva-
 069 tion laws, that blends the structure-preserving benefits of FV schemes with the expressiveness of neural
 070 networks. Conservation is built into the NFV model, using extended spatial and temporal stencils. We
 071 develop both a supervised version, trained on solution data from simple cases, and an unsupervised
 072 variant (UNFV), which learns directly from the PDE via a weak-form residual loss. This flexibility
 073 allows (U)NFV to adapt to the availability of data, leveraging accurate synthetic or field data when
 074 present, or solving directly from the equation when solutions are inaccurate or expensive to obtain.
 075 **We focus on one-dimensional scalar conservation laws, which are widely used in applications such as**
 076 **traffic flow, pipeline and channel models, and form a standard, well-understood testbed.**

077 **Contributions.** Our main contributions are as follows:

- 078 • We propose (U)NFV, a neural architecture that generalizes the structure of finite volume methods
 079 and thus preserves conservation properties by construction.
- 080 • We introduce two variants: a supervised learning one (NFV) and an unsupervised learning one
 081 (UNFV), depending on data availability, using either solution data or a weak-form residual loss.
- 082 • We demonstrate strong numerical results on several conservation laws, achieving up to 10x
 083 lower error than classical FV solvers, as shown in Figure 1. Additionally, (U)NFV matches the
 084 accuracy of discontinuous Galerkin methods, without their mathematical complexity.
- 085 • We show that NFV can be trained on field data that does not strictly satisfy the conservation law,
 086 and still predicts accurate solutions with more generalizability than classical solvers.

087 The remainder of the article is organized as follows: Section 2 provides a detailed overview of the
 088 related work, Section 3 recalls the FV and introduces necessary notation, Section 4 describes the
 089 proposed (U)NFV method in detail, Section 5 presents the experiments and results on hyperbolic
 090 PDEs, Section 6 extends the NFV to experimental field highway data, and Section 7 concludes
 091 the article. Then, Appendix A provides details about FV schemes, Appendix B illustrates six PDE
 092 variants considered in this work, Appendix C expands on the experimental data handling and results
 093 from Section 6, and Appendix D details the model architecture, dataset, and hyperparameter choices.

094 2 RELATED WORK

095 **Numerical methods.** Classical numerical methods for hyperbolic PDEs, such as FV and *discon-*
 096 *tinuous Galerkin* (DG) (Hu and Shu, 1999) methods, are widely used due to their capabilities in
 097 capturing shocks and discontinuities. First-order schemes such as the *Lax-Friedrichs* (Lax, 1954)
 098 and *Godunov* (Godunov, 1959b) methods provide robustness but suffer from excessive numerical
 099 diffusion, leading to smeared solutions. To address this, higher-order methods like *Essentially*
 100 *Non-Oscillatory* (ENO) (Shu, 1999), *Weighted ENO* (WENO) (Liu et al., 1994), and DG have been
 101 introduced, offering improved accuracy in smooth regions while preserving stability near shocks. DG
 102 as a Finite Element method, further improves accuracy through local polynomial approximations but
 103 incurs high computational costs (Cockburn and Shu, 1998). **In practice, DG and higher-order FV**
 104 **schemes like WENO also demand intricate flux constructions, quadrature rules, and stabilization**
 105 **choices, whereas (U)NFV retains FV-like implementation complexity.** Despite their accuracy, these
 106 methods often require extensive manual effort and careful stabilization, motivating the development
 107 of flexible, data-driven alternatives.

108 **NN approaches for PDEs.** Deep learning has become a powerful alternative for approximating
 109 PDE solutions. In the supervised learning case, neural operators such as *Fourier Neural Operator*
 110 (FNO) (Li et al., 2020) and *Deep Operator Networks* (DeepONet) (Lu et al., 2021) efficiently
 111 approximate solution mappings from parametric inputs, without requiring explicit mesh discretization
 112 in the case of FNO. While successful for general PDEs, these operators have mainly been validated
 113 on elliptic or parabolic PDEs, typically characterized by smooth solutions. Conventional neural
 114 architectures, such as CNNs (LeCun et al., 1995) for structured domains and GNNs (Bronstein et al.,
 115 2017) for irregular geometries, have also been adopted as supervised surrogates. However, supervised
 116 models rely heavily on large, high-quality labeled datasets, and often lack intrinsic enforcement of
 117 physical constraints, leading to limited generalization and poor accuracy on PDEs involving sharp
 118 gradients or shocks (Krishnapriyan et al., 2021).

119 To reduce data reliance, unsupervised approaches like *Physics-Informed Neural Networks* (PINNs)
 120 incorporate PDE residuals directly into training losses (Raissi et al., 2017), proving effective for
 121 elliptic and parabolic equations (Raissi et al., 2019; Jagtap et al., 2020). However, PINNs encounter
 122 significant difficulties with hyperbolic PDEs, especially in capturing discontinuities and shock
 123 dynamics, resulting in unstable optimization, convergence failures, and inaccurate solutions (Wang
 124 and Liu, 2021; Fuks and Tchelepi, 2020). Recent variants, such as *Weak PINNs* (wPINNs) (De Ryck
 125 et al., 2024), *Parareal PINNs* (PPINNs) (Meng et al., 2020), and *Extended PINNs* (XPINNs) (Jagtap
 126 and Karniadakis, 2020), aim to overcome these issues through weak formulations or specialized
 127 training strategies. Nonetheless, these adaptations often introduce considerable complexity and
 128 require extensive hyperparameter tuning, highlighting a persistent need for methods inherently suited
 129 to hyperbolic PDE challenges.

130 **NNs for hyperbolic PDEs and conservation laws.** Neural approaches tailored to hyperbolic PDEs
 131 have introduced innovations to handle shocks. Weak PINNs (wPINNs) (De Ryck et al., 2024)
 132 integrate weak-form residuals or integral constraints to mitigate issues with discontinuities. Others
 133 employ neural networks directly within classical FV schemes to learn improved flux reconstruc-
 134 tions (Kossaczka et al., 2021; Tong et al., 2024). However, these enhancements typically reintroduce
 135 complexity, such as extensive manual parameterization or problem-specific adaptivity, diluting the
 136 key advantage of neural flexibility and generality.

137 Motivated by these limitations, our proposed NFV approach learns local update rules directly from
 138 data or PDE residuals. By preserving the fundamental conservation-law structure of traditional FV
 139 methods while flexibly leveraging neural networks, NFV achieves significantly higher accuracy,
 140 robustness, and scalability with minimal manual intervention.

141 3 PREREQUISITES AND NOTATIONS: FINITE VOLUME METHODS

144 Standard FV methods, such as those presented in LeVeque (2002), solve the integral form of the
 145 conservation law (1) on a mesh of uniform cells $I_i = [x_{i-1/2}, x_{i+1/2}]$, $i = 1, \dots, I_{\max}$, with cell
 146 length Δx . The average of u over cell I_i at time $t_n = n\Delta t$, for $n = 1, \dots, N$ and time discretization
 147 Δt , and the numerical flux through the interface $x_{i+1/2}$ over the time step, are given respectively by

$$149 \quad u_i^n = \frac{1}{\Delta x} \int_{I_i} u(t_n, x) \, dx \quad \text{and} \quad F_{i+1/2}^n = \int_{t_n}^{t_{n+1}} f(u(t, x_{i+1/2})) \, dt. \quad (2)$$

152 A first-order method \mathcal{F} approximates the numerical
 153 flux as $\hat{F}_{i+1/2}^n = \mathcal{F}(u_i^n, u_{i+1}^n)$, while higher-order
 154 methods leverage additional cell averages. Let us
 155 generalize this framework by including cell averages
 156 from previous time steps in order to construct even
 157 better approximations. Let FV_a^b be the class of meth-
 158 ods that use a rectangular stencil of a neighboring
 159 spatial cells times b past time steps to estimate numer-
 160 ical fluxes. Specifically, let $u_{i\pm 1/2}^n(a-1, b)$ be the
 161 left and right $(a-1) \times b$ sub-stencils, as illustrated
 in blue and green in Figure 2.

u_1^4	u_2^4	u_3^4	u_4^4	u_5^4
u_1^3	u_2^3	u_3^3	u_4^3	u_5^3
u_1^2	u_2^2	$F_{2,5}^2 u_3^2$	$F_{3,5}^2 u_4^2$	u_5^2
u_1^1	u_2^1	u_3^1	u_4^1	u_5^1

162 Figure 2: Example stencil for FV_5^2 , taking in a
 163 stencil of 2 time steps times 5 space cells.

Then, an FV_a^b method \mathcal{F} estimates numerical fluxes as $\hat{F}_{i\pm 1/2}^n = \mathcal{F}(\mathbf{u}_{i\pm 1/2}^n(a-1, b))$. Classical first-order methods, such as Godunov, fall under class FV_3^1 ; more details about their computation are provided in Appendix A. To our knowledge, the vast majority of FV methods in the literature use a single time step (i.e., $b = 1$) and a small number of spatial cells. Indeed, designing analytical schemes with larger temporal or spatial stencils becomes exponentially more complex. Finally, the update rule is given by the exact relation

$$u_i^{n+1} = u_i^n - \frac{\Delta t}{\Delta x} (F_{i+1/2}^n - F_{i-1/2}^n), \quad (3)$$

which in practice is approximated using the numerical fluxes $\hat{F}_{i\pm 1/2}^n$, leading to an approximation \hat{u}_i^n of u_i^n . Note that the influx of one cell is the outflux of another, which ensures conservation.

4 OUR METHOD: NEURAL FINITE VOLUME (NFV)

Our method builds upon the FV framework by using neural networks to approximate the numerical flux. Specifically, we define NFV_a^b as a generalization of FV_a^b , where the numerical flux $\hat{F}_{i\pm 1/2}^n$ is predicted by a neural network \mathcal{N} based on a local $a \times b$ spatiotemporal stencil:

$$\hat{F}_{i\pm 1/2}^n = \mathcal{N}(\mathbf{u}_{i\pm 1/2}^n(a-1, b))$$

The prediction of the solution is then updated using the classical FV update rule (3), ensuring mass conservation. We explore NFV models ranging from NFV_3^1 (matching Godunov’s stencil) to NFV_{11}^{11} , using 11 spatial cells and 11 past time steps – configurations that would be exceedingly complex to design manually due to the high-dimensional stencil involved. This extension enables accurate learning even from noisy field data. In practice, we implement NFV as a CNN (LeCun et al., 1995), which allows efficient computation across stencils due to the vectorized nature of CNNs. Since (U)NFV retains the standard finite volume update, boundary conditions such as Dirichlet, Neumann, or open boundaries can be imposed via ghost cells or prescribed interface fluxes exactly as in classical FV schemes, without modifying the neural architecture.

In all experiments we instantiate NFV_a^b as a lightweight two-dimensional CNN applied locally on each cell interface: the first layer uses a kernel of size $a-1$ over the spatial dimension with b input channels (one per time slice), followed by five 1×1 convolutional layers with 15 channels and either ELU or ReLU activations depending on the flux family. This architecture yields $1105 + 16 \cdot ((a-1) \cdot b + 1)$ trainable parameters for NFV_a^b , so even our largest models contain only a few thousand parameters while retaining the exact FV update rule.

We propose two variants of NFV that share the same architecture but differ in their training objectives: the supervised NFV_a^b , trained on reference solutions, and the unsupervised UNFV_a^b , trained directly from the PDE via a weak-form residual loss. The supervised setting applies when solution data is available, while the unsupervised variant enables training when such data is absent, relying instead on the governing conservation laws. Moreover, supervised NFV can also be applied in cases where the PDE is unknown but observational data is accessible, allowing solvers to be deployed directly on field data with only basic physical constraints, such as mass conservation, imposed, and without extensive hyperparameter tuning (see Section 6).

In all our experiments, we therefore train one (U)NFV model per conservation law, and once trained the same network can be applied to many different initial conditions for that equation, so the one-time training cost is largely amortized and in practice remains very short. At inference time, no optimization is solved: each time step is advanced by a single application of the finite volume update rule (3) with numerical fluxes $\hat{F}_{i\pm 1/2}^n$ given by a forward pass of the neural network, so the overall cost of solving an equation scales linearly with the number of time steps.

4.1 SUPERVISED LEARNING

Supervised learning offers a straightforward framework for training models when reference solutions are available. In this study, we employ supervised learning not only to approximate the solution of known equations but also to predict field data with unknown governing equations. Although solutions to hyperbolic PDEs are typically defined in the L_1 space, we consider their restrictions to

216 compact subsets where the functions are bounded, thereby allowing treatment within the L_2 space.
 217 Accordingly, the loss function is defined as the standard mean square error:

$$218 \quad 219 \quad \mathcal{L}_s = \mathbb{E}_{u_0 \sim \mathcal{R}} \|u - \hat{u}\|_2^2$$

220 where u is the true solution, \hat{u} is the predicted solution, and \mathcal{R} is a distribution over initial conditions.
 221

222 4.2 UNSUPERVISED LEARNING

224 Unsupervised learning for hyperbolic PDEs is particularly challenging because their solutions often
 225 lack closed-form expressions and classical (strong) solutions may not exist. Instead, these equations
 226 are typically defined through weak formulations. Although weak solutions are not unique: multiple
 227 functions can satisfy the PDE, but only one corresponds to the physically relevant *entropy solution*,
 228 which enforces admissibility conditions across shocks and discontinuities.

229 The unsupervised loss function is defined to minimize the residuals of the weak formulation, in order
 230 to approximate the entropy solution. While imposing this loss does not guarantee convergence to the
 231 entropy solution, empirical results indicate that our method successfully converges to the entropy
 232 solution across various equations and numerous trials. To enhance learning efficiency, we optimize
 233 the weak formulation independently at each time step by minimizing the squared residuals. The
 234 collection of test functions Φ consists of 250 randomly sampled, compactly supported polynomials
 235 of degree 50 over the spatial domain. The unsupervised loss reads:

$$236 \quad 237 \quad \mathcal{L}_w = \mathbb{E}_{\varphi \in \Phi} \left[\sum_{n=1}^N \left(\sum_{i=1}^{I_{\max}} \left((\Delta t)^{-1} (\hat{u}_i^n - \hat{u}_i^{n-1}) \int_{I_i} \varphi + f(\hat{u}_i^n) [\varphi]_{x_{i-1/2}}^{x_{i+1/2}} \right) \right)^2 \right]$$

238 where \hat{u}_i^n denotes the predicted solution at spatial index i and time step n , and \mathcal{R} is a distribution
 239 over initial conditions. Note that for the scalar conservation laws considered here, integration by
 240 parts removes spatial derivatives from the weak-form loss, and time derivatives are handled via finite
 241 differences in the FV update, so no explicit spatial derivatives of the primal variables are required
 242 during training.
 243

244 5 EXPERIMENTS

247 Experiments have been designed to answer four main questions:

- 248 • Is (U)NFV a compelling alternative to classical finite volume methods?
- 249 • Does UNFV converge to an entropy solution despite being trained on the weak formulation?
- 250 • How does (U)NFV compare to much more complicated finite element methods?
- 251 • Can NFV perform well on field data that contains noise and may not be conservative?

253 5.1 BASELINES

254 Selecting appropriate baselines for PDE solvers poses challenges due to the diversity in computational
 255 frameworks: methods vary by mesh dependency (mesh-free versus mesh-based), solution generation
 256 (autoregressive versus single-pass), and generalizability (operator-based versus retrained per initial
 257 condition). Therefore, we adopt classical numerical schemes, the foundation of our NFV method,
 258 as baselines, ensuring a fair comparison. Given the fact that NFV is developed based on traditional
 259 first-order FV methods, the present work provides a compelling case for replacing standard FV
 260 solvers with the simpler yet effective NFV method whenever FV methods are typically employed.
 261 We consider all the numerical schemes introduced in Section 2 as baselines: first-order FV methods
 262 (Godunov, Lax-Friedrichs, and Engquist-Osher), higher-order ones (ENO, WENO), and DG, a finite-
 263 element method that is well-known for superior accuracy but suffers from computational burden.
 264 More details can be found in Appendix A.

265 5.2 EQUATIONS

267 The **Lighthill-Whitham-Richards** model (Lighthill and Whitham, 1955; Richards, 1956), known as
 268 LWR, is a first-order hyperbolic conservation law used to model traffic flow. It is expressed as
 269

$$270 \quad \partial_t \rho + \partial_x (\rho v(\rho)) = 0 \quad (4)$$

270 Table 1: Performance comparison between neural network models and classical numerical schemes. Results
 271 are computed over the evaluation set of 1000 piecewise constant initial conditions. For each method, we report
 272 mean and standard deviation in L_2 norm (mean($(u - \hat{u})^2$)).

	1 st order FV				Higher order FV			FEM
	NFV ₃ ¹	UNFV ₃ ¹	GD	LF	EO	ENO	WENO	DG
G.shields	1.3e⁻⁴_{±4e⁻⁵}	2.0e ⁻⁴ _{±6e⁻⁵}	4.5e ⁻⁴ _{±2e⁻⁴}	1.3e ⁻² _{±4e⁻³}	4.5e ⁻⁴ _{±2e⁻⁴}	6.4e ⁻⁴ _{±4e⁻⁴}	6.4e ⁻⁴ _{±4e⁻⁴}	3.1e ⁻⁵ _{±1e⁻⁵}
Tri. 1	1.4e⁻³_{±6e⁻⁴}	1.9e ⁻³ _{±9e⁻⁴}	2.3e ⁻³ _{±1e⁻³}	9.6e ⁻³ _{±4e⁻³}	2.3e ⁻³ _{±1e⁻³}	2.0e ⁻³ _{±2e⁻³}	1.9e ⁻³ _{±2e⁻³}	2.6e ⁻⁴ _{±1e⁻⁴}
Tri. 2	2.4e⁻³_{±1e⁻³}	3.1e ⁻³ _{±2e⁻³}	3.8e ⁻³ _{±2e⁻³}	1.4e ⁻² _{±8e⁻³}	3.8e ⁻³ _{±2e⁻³}	5.8e ⁻³ _{±4e⁻³}	5.8e ⁻³ _{±4e⁻³}	4.1e ⁻⁴ _{±2e⁻⁴}
Trapez.	1.1e⁻³_{±4e⁻⁴}	1.6e ⁻³ _{±7e⁻⁴}	2.1e ⁻³ _{±8e⁻⁴}	2.5e ⁻² _{±1e⁻²}	2.1e ⁻³ _{±8e⁻⁴}	6.2e ⁻⁴ _{±2e⁻⁴}	5.3e ⁻⁴ _{±2e⁻⁴}	2.9e ⁻⁴ _{±1e⁻⁴}
G.berg	1.4e⁻⁴_{±9e⁻⁵}	3.8e ⁻⁴ _{±2e⁻⁴}	4.9e ⁻⁴ _{±2e⁻⁴}	5.3e ⁻³ _{±2e⁻³}	4.9e ⁻⁴ _{±2e⁻⁴}	1.1e ⁻³ _{±6e⁻⁴}	1.2e ⁻³ _{±9e⁻⁴}	3.4e ⁻⁴ _{±2e⁻³}
U.wood	3.8e⁻⁴_{±1e⁻⁴}	6.9e ⁻⁴ _{±2e⁻⁴}	9.2e ⁻⁴ _{±3e⁻⁴}	2.7e ⁻² _{±1e⁻²}	9.2e ⁻⁴ _{±3e⁻⁴}	1.1e ⁻⁴ _{±3e⁻⁵}	9.8e ⁻⁵ _{±2e⁻⁵}	5.9e ⁻⁵ _{±2e⁻⁵}
Burgers	8.5e⁻⁴_{±3e⁻⁴}	1.3e ⁻³ _{±6e⁻⁴}	1.9e ⁻³ _{±7e⁻⁴}		2.6e ⁻³ _{±1e⁻³}	2.7e ⁻³ _{±1e⁻³}	2.8e ⁻³ _{±1e⁻³}	1.0e ⁻⁴ _{±4e⁻⁵}

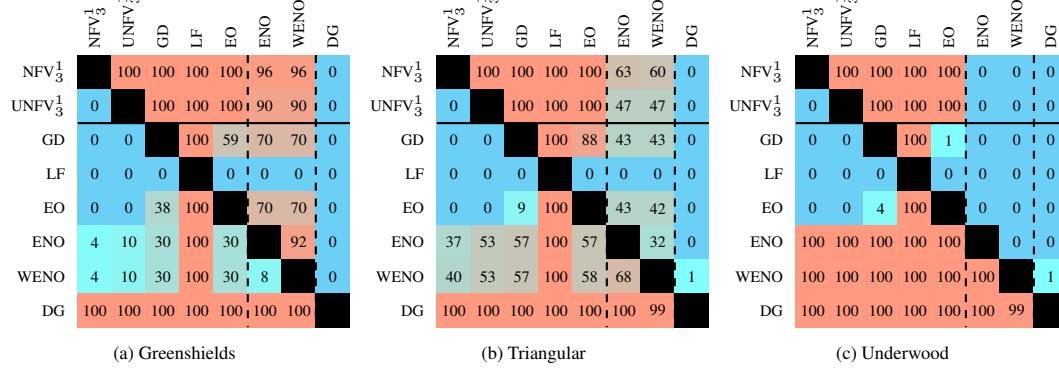
284
 285 where ρ is the density of the traffic, $f : \rho \mapsto \rho v(\rho)$ is the flux function and v is the velocity. The
 286 flux function is typically modeled as a concave function of the density. Variations in the underlying
 287 velocity function give rise to different traffic flow models. In this work, six different models have
 288 been considered: Greenshields' (Greenshields et al., 1935), Triangular (Geroliminis and Daganzo,
 289 2008), Triangular skewed (Geroliminis and Daganzo, 2008), Trapezoidal (Geroliminis and Sun,
 290 2011), Greenberg (Greenberg, 1959) and Underwood (Underwood, 1961). These models behave *very*
 291 differently and should be considered as different equations, as shown in Figure 1. Formulations and
 292 illustrations of those six models are given in Appendix B.

293 The **inviscid Burgers' equation** is a well-known hyperbolic conservation law used in various
 294 domains such as fluid mechanics (Burgers, 1939), non-linear acoustics (Lombard et al., 2013), gas
 295 dynamics (Panayotounakos and Drikakis, 1995), and traffic flow (Musha and Higuchi, 1978). We
 296 refer the reader to Cameron (2011) for a thorough introduction. It is expressed as

$$\partial_t u + \frac{1}{2} \partial_x u^2 = 0. \quad (5)$$

300 This equation can be written in the classical form of a conservation law using the flux function
 301 $f : u \mapsto \frac{1}{2} u^2$. Exact solutions to Riemann initial conditions are also known for this problem.
 302 Visualization of some solutions, including videos, are available on our webpage and in Figure 1.

303 Our experiments focus on one-dimensional conservation laws in this work. Demonstrating that NFV
 304 can consistently outperform classical schemes in 1D establishes a strong foundation before tackling
 305 more complex systems. NFV architecture is, in principle, extendable to higher dimensions, since
 306 neural networks naturally scale to higher-dimensional inputs. Extending NFV to multi-dimensional
 307



321 Figure 3: **Comparison of numerical schemes across flow functions.** Each cell shows the proportion of the
 322 evaluation set on which the row scheme outperforms the column scheme. DG, the only FEM tested, is rarely
 323 beaten. NFV₃¹ and UNFV₃¹ outperform other first-order schemes and rival higher-order ones, making them strong
 324 choices depending on the equation.

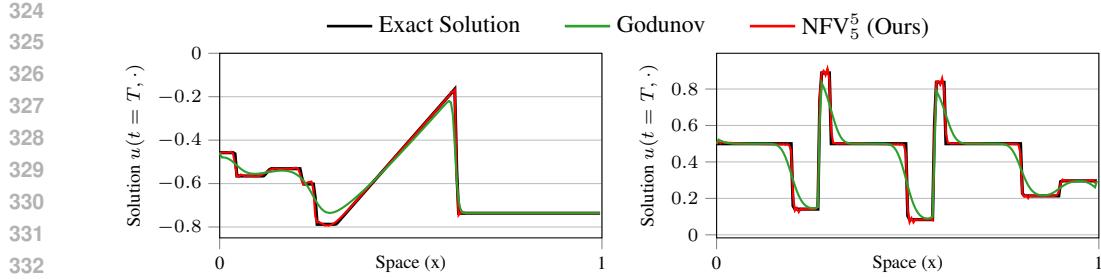


Figure 4: Comparison of the final density of the Burgers’ equation (left) and LWR triangular equation (right) for NFV_5^5 and the Godunov Scheme. The proposed method displays an excellent approximation of the exact solution, capturing sharp features such as discontinuities and points of non-differentiability. It contains some minor oscillations in the solution, which are not present in the Godunov scheme. The latter, however, fails to capture the discontinuities and points of non-differentiability, offering a very smoothed solution.

Table 2: Evaluation of NFV_5^5 using piecewise constant initial conditions. Error is reported in L_2 norm. NFV_5^5 achieve outstanding performance, gaining up to an order of magnitude improvement compared to Godunov and WENO. Its performance is close to DG, while keeping the implementation simplicity of a finite volume method and the computational complexity of NFV.

	Godunov	WENO	NFV_3^1	NFV_5^5	DG
Burgers’	$1.8e^{-3}_{\pm 6e^{-4}}$	$2.6e^{-3}_{\pm 1e^{-3}}$	$8.3e^{-4}_{\pm 3e^{-4}}$	$2.2e^{-4}_{\pm 1e^{-4}}$	$1.0e^{-4}_{\pm 4e^{-5}}$
Greenshields	$4.1e^{-4}_{\pm 1e^{-4}}$	$6.9e^{-4}_{\pm 4e^{-4}}$	$1.2e^{-4}_{\pm 4e^{-5}}$	$4.6e^{-5}_{\pm 3e^{-5}}$	$4.2e^{-5}_{\pm 2e^{-5}}$
Triangular	$2.2e^{-3}_{\pm 1e^{-3}}$	$2.0e^{-3}_{\pm 2e^{-3}}$	$1.3e^{-3}_{\pm 6e^{-4}}$	$2.9e^{-4}_{\pm 2e^{-4}}$	$2.7e^{-4}_{\pm 1e^{-4}}$

will introduce additional challenges (e.g., numerical stability, computational complexity, and coupled variables), which we identify as important avenues for future work.

5.3 DATASETS

Training is performed using solutions derived from Riemann problems, which are initial value problems characterized by piecewise constant initial conditions with a single discontinuity (see Figure 8 for examples). These problems are fundamental in the study of hyperbolic PDEs and serve as essential test cases for numerical methods. For the scenarios considered in this work, analytical solutions to Riemann problems are available, making supervised learning possible. Evaluation is performed on a more complicated set of several hundred complex initial conditions to assess the model’s generalization capabilities. These conditions consist of piecewise constant functions with ten discontinuities, giving rise to entropy solutions with multiple interacting shocks and rarefactions. Exact solutions for these test cases are computed using the Lax-Hopf algorithm (Lax, 1957; Claudel and Bayen, 2010a;b) on a finer grid (see Appendix D).

For the LWR benchmarks we train NFV autoregressively on 2048 randomly sampled Riemann problems with a single discontinuity, using discretization parameters $\Delta t = 5 \cdot 10^{-3}$, $\Delta x = 10^{-2}$, 100 spatial cells, and prediction horizons that are progressively increased from 10 to 250 steps under a robust CFL ratio of 0.5. Evaluation uses several hundred more complex piecewise-constant initial conditions whose exact solutions are computed on a finer grid with $\Delta t = 10^{-4}$, $\Delta x = 10^{-3}$, 200 cells, and 1000 time steps via the Lax–Hopf algorithm. Unsupervised UNFV models minimize a weak-form residual loss using 250 compactly supported polynomials of degree 50 as test functions over the spatial domain. Note that we use uniform space–time grids for simplicity and fair comparison to FV baselines, but the (U)NFV update depends only on cell volumes and interface fluxes and is therefore compatible with non-uniform or adaptive discretizations.

5.4 RESULTS AND DISCUSSION

Table 1 reports L_2 error for NFV_3^1 , UNFV_3^1 , and baseline methods across the seven benchmark equations. Our models consistently outperform all first-order FV methods, and surpass ENO/WENO schemes on about half of the equations. As expected, the higher-order DG method achieves signif-

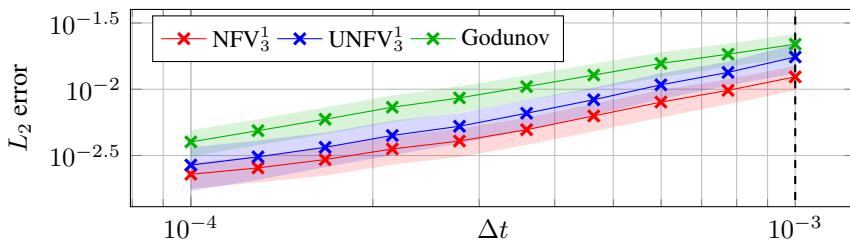


Figure 5: **Convergence plots on Greenshields’ flux.** The L_2 error is computed against the exact solution on the evaluation set for different mesh discretizations. We report both error average and standard deviation, on a log-log scale. The dashed vertical line illustrates the discretization on which NFV_3^1 and $UNFV_3^1$ were trained; the models generalize to smaller discretizations. The ratio $\Delta t / \Delta x = 0.1$ remains constant as the mesh is refined.

significantly lower errors. Table 2 shows that NFV_5^5 , while as simple to implement as standard NFV_3^1 , achieves up to 10x better accuracy, approaching the performance of DG. In this sense, (U)NFV offers DG-level accuracy but with much smaller FV-like implementation complexity, substantially faster inference, and training that typically completes within fifteen minutes, while using memory comparable to Godunov and significantly lower than DG, since only the numerical flux is learned while the rest of the finite volume solver remains unchanged.

Figure 3 shows the fraction of test cases each method wins. NFV_3^1 and $UNFV_3^1$ consistently surpass first-order FV methods. Against ENO/WENO, performance varies: our models outperform on some equations, match on others, and underperform in a few, highlighting the complexity of benchmarking across diverse problem settings. Still, the fact that NFV_3^1 and $UNFV_3^1$ consistently do better than first-order methods is seen as a sign that the approach appears to converge well. Specifically, NFV_3^1 and $UNFV_3^1$ consistently produce errors bounded by those of Godunov, emphasizing their robustness.

Since all methods use autoregressive prediction, evaluating performance at the final time step provides a good proxy for cumulative error. Figure 4 shows that the prediction of NFV_5^5 closely aligns with the exact solution, with only minor oscillations observed. Notably, NFV_5^5 effectively captures sharp discontinuities with high accuracy without relying on smoothing techniques, which are commonly employed in traditional FV methods to mitigate numerical artifacts. Qualitatively, across the large set of complex test initial conditions that span weak to strong shocks and rarefaction patterns, NFV predictions remain visibly less diffusive than the others FV schemes, with sharp features dissipating more slowly while preserving stability.

Ablation on discretization size: Classical numerical schemes are known to converge as the grid is refined. Figure 5 shows that NFV_3^1 and $UNFV_3^1$ consistently achieve lower error than Godunov, a scheme proven to converge, across discretizations, suggesting that (U)NFV also converges to the entropy solution; the approximately linear trend in the log–log plot further indicates a polynomial convergence rate.

Ablation on CFL ratio: To further assess stability under different time step choices, we vary the CFL ratio while keeping the spatial grid fixed on Greenshields’ LWR and report the resulting errors in Table 3. Across this range of CFL values, NFV_3^1 consistently attains lower mean error and substantially reduced variance compared to first-order FV baselines, and remains competitive with higher-order ENO and WENO schemes. The only exception is DG at very small CFL, which achieves the lowest error but becomes unstable and fails to run at higher CFL ratios, whereas NFV_3^1 remains robust.

Overall, the results support our hypothesis that training on simple Riemann problems is sufficient to generalize to complex piecewise-constant initial conditions. In particular, the ability of (U)NFV trained only on these analytically tractable Riemann building blocks to generalize reliably to much richer piecewise-constant and real-world configurations turns this seemingly strong assumption into a practical strength rather than a limitation. Additional dataset and training details, along with the heuristic exploration and hyperparameter tuning that led the method to work, are provided in Appendix D. In the next section, we show that NFV also generalizes to experimental highway data, where conservation is often violated and traditional methods typically fail.

432 Table 3: Mean and standard deviation of final-time L_2 error on the standard LWR benchmark with Greenshields' 433 flux for different CFL ratios, comparing NFV_3^1 with classical finite volume baselines and DG.

CFL	NFV_3^1	GD	LF	EO	ENO	WENO	DG
0.2	$1.6e^{-4}_{\pm 3e^{-5}}$	$3.8e^{-4}_{\pm 1e^{-4}}$	$7.6e^{-3}_{\pm 2e^{-3}}$	$3.8e^{-4}_{\pm 1e^{-4}}$	$6.0e^{-4}_{\pm 4e^{-4}}$	$6.2e^{-4}_{\pm 4e^{-4}}$	$3.0e^{-5}_{\pm 1e^{-5}}$
0.4	$1.3e^{-4}_{\pm 2e^{-5}}$	$3.3e^{-4}_{\pm 1e^{-4}}$	$4.1e^{-3}_{\pm 1e^{-3}}$	$3.3e^{-4}_{\pm 1e^{-4}}$	$6.0e^{-4}_{\pm 4e^{-4}}$	$6.4e^{-4}_{\pm 4e^{-4}}$	fail
0.6	$1.2e^{-4}_{\pm 5e^{-5}}$	$2.1e^{-4}_{\pm 2e^{-4}}$	$1.3e^{-3}_{\pm 4e^{-4}}$	$2.2e^{-4}_{\pm 2e^{-4}}$	$1.5e^{-2}_{\pm 1e^{-2}}$	$1.5e^{-3}_{\pm 1e^{-3}}$	fail
0.8	$1.0e^{-4}_{\pm 2e^{-5}}$	$2.2e^{-4}_{\pm 7e^{-5}}$	$2.0e^{-3}_{\pm 6e^{-4}}$	$2.3e^{-4}_{\pm 7e^{-5}}$	$1.6e^{-3}_{\pm 2e^{-3}}$	$7.2e^{-4}_{\pm 4e^{-4}}$	fail
1.0	$9.1e^{-5}_{\pm 2e^{-5}}$	$1.7e^{-4}_{\pm 5e^{-5}}$	$1.5e^{-3}_{\pm 5e^{-4}}$	$1.8e^{-4}_{\pm 5e^{-5}}$	$5.6e^{-3}_{\pm 6e^{-3}}$	$9.6e^{-4}_{\pm 7e^{-4}}$	fail
1.2	$1.2e^{-4}_{\pm 5e^{-5}}$	$2.1e^{-4}_{\pm 2e^{-4}}$	$1.3e^{-3}_{\pm 4e^{-4}}$	$2.2e^{-4}_{\pm 2e^{-4}}$	$1.5e^{-2}_{\pm 1e^{-2}}$	$1.5e^{-3}_{\pm 1e^{-3}}$	fail

434
435
436
437
438
439
440
441
442
443
444 Table 4: **Improvements of NFV at different scales against numerical methods with fitted flow functions on**
445 **field data.** The reported metrics include L1 error ($\text{mean}(|u - \hat{u}|)$), L2 error ($\text{mean}((u - \hat{u})^2)$), and relative
446 error ($\text{mean}(|u - \hat{u}| / |\max\{\varepsilon, u\}|)$). The larger the input size of NFV, the better the performance. NFV_3^1
447 outperforms all calibrated Godunov fits, despite having the same input size and underlying structure.
448
449
450
451
452
453
454
455
456
457

	Calibrated numerical schemes (Godunov)					NFV (Ours)		
	Greenshields	Triangular	Trapezoidal	Greenberg	Underwood	NFV_3^1	NFV_5^5	NFV_{11}^{11}
L1	$6.05e^{-2}$	$2.77e^{-2}$	$2.73e^{-2}$	$2.79e^{-2}$	$4.98e^{-2}$	$2.37e^{-2}$	$2.31e^{-2}$	$2.02e^{-2}$
L2	$1.93e^{-1}$	$1.31e^{-1}$	$1.30e^{-1}$	$1.33e^{-1}$	$1.81e^{-1}$	$1.23e^{-1}$	$1.21e^{-1}$	$1.09e^{-1}$
Rel.	$5.04e^{-1}$	$3.83e^{-1}$	$3.74e^{-1}$	$3.75e^{-1}$	$5.45e^{-1}$	$3.57e^{-1}$	$3.51e^{-1}$	$2.83e^{-1}$

458 6 MODELING LARGE-SCALE EXPERIMENTAL FIELD DATA USING NFV

459 We apply the proposed NFV method to large-scale traffic field data collected on Interstate 24 (I-24)
460 in Tennessee, USA, using the I-24 MOTION infrastructure (Gloudemans et al., 2023a;b), which
461 enables high-resolution vehicle trajectory collection and constitutes the most extensive publicly
462 available traffic dataset to date. Rather than predicting traffic speed, we focus on modeling traffic
463 density, which is more directly tied to conservation laws and often exhibits sharp transitions that are
464 challenging to capture. Although conservation of mass is not strictly satisfied in highway traffic data
465 due to merges, exits, and incidents, it serves as a strong inductive bias. We show that NFV achieves
466 superior predictive accuracy compared to classical numerical schemes. Moreover, incorporating the
467 PDE structure leads to substantially more stable training, particularly in data-scarce regimes. These
468 findings suggest that our approach can enhance the accuracy and efficiency of traffic simulations,
469 thereby contributing to better-informed decision-making in urban planning and traffic management.
470

471 6.1 DATASET AND TRAINING

472 We evaluate our method on the I-24 MOTION dataset (Gloudemans et al., 2023a), which provides
473 high-resolution vehicle trajectories collected on a four-mile stretch of Interstate 24 (mile markers
474 58.7 to 62.7) near Nashville, Tennessee. The data is captured by a network of high-definition cameras
475 mounted along the highway as part of the I-24 MOTION infrastructure, leading to intricate trajectory
476 data as illustrated in Figure 9. Vehicle trajectories are reconstructed using a computer vision and data
477 association pipeline (Wang et al., 2022), resulting in high-fidelity, though inherently noisy, field data.

478 The dataset consists of 10 days of vehicle trajectory data, collected during the morning rush hour
479 (6:00 AM to 10:00 AM) over the 4-mile segment. From the raw trajectory data, we construct
480 spatiotemporal vehicle density fields by aggregating vehicle counts over fixed spatial cells. Details of
481 the data cleaning, processing, and preparation are provided in Appendix C.1. Visualization of the
482 resulting density fields is shown in Figure 10. Further training details are available in Appendix D.

483 Concretely, all NFV models and tuned finite-volume baselines are trained on the first hour of data
484 from November 29, 2022 using a single boundary cell on each side; the autoregressive prediction
485 horizon is increased from 10 to 100 steps during training, while the learning rate decays from 10^{-3}
486 to 10^{-4} over roughly 3000–5000 epochs, leading to convergence within 15–30 minutes on a single
487 NVIDIA RTX A5000 GPU.

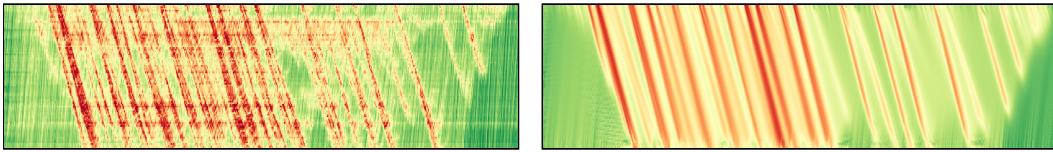


Figure 6: **Autoregressive prediction of NFV_{11}^{11} (right) compared to the ground truth (left).** Full results are shown in Figure 13. See Appendix C.3 for how to read the heatmaps.

Table 5: **Generalization of NFV against Godunov on 7 days of I-24 data never seen during training.** As in Table 4, we report mean and standard deviation of L1, L2 and relative errors.

	L1 error	L2 error	Relative error
Godunov	$1.56e^{-1} \pm 2.02e^{-2}$	$3.74e^{-2} \pm 8.25e^{-3}$	$6.26e^{-1} \pm 2.58e^{-1}$
NFV_{11}^{11}	$1.12e^{-1} \pm 7.39e^{-3}$	$2.20e^{-2} \pm 2.59e^{-3}$	$3.59e^{-1} \pm 7.58e^{-2}$

6.2 RESULTS AND DISCUSSION

We compare NFV to numerical schemes using the flux functions from Appendix B. These functions, each defined by a few parameters, were calibrated via optimization to minimize the Godunov scheme’s prediction error on the training set. The search ranges were intentionally broad, prioritizing predictive performance over physical plausibility to ensure a fair comparison. We chose the Godunov scheme as the representative baseline to compare with since we observed only a marginal performance difference (up to 5%) between Godunov and other FV schemes on this dataset, and the Godunov scheme is known to converge to the entropy solution. We evaluate three NFV variants of increasing capacity: NFV_3^1 , NFV_5^5 , and NFV_{11}^{11} (training details can be found in Appendix D), to assess how well they generalize and capture complex field dynamics.

Table 4 shows that all NFV models outperform the five tuned Godunov schemes, with performance improving as input size increases. This trend matches what was seen on synthetic data (Section 5). Despite training on just one hour of data, NFV predicts nearly four hours of traffic evolution autoregressively (Figure 6). While performance degrades in out-of-distribution zones (e.g., dark green regions unseen during training), the models still capture key wave patterns with high fidelity. Larger stencils help smooth out noise and improve accuracy, as seen in Figure 13.

We further evaluate generalization on 7 other days. As shown in Figure 14 and Table 5, NFV_{11}^{11} consistently outperforms the best Godunov scheme on the evaluation set, even though both perform similarly on the training day. Indeed, although far from perfect, it is able to capture the evolution of free-flow traffic (dark green) with much greater accuracy, allowing it to successfully capture the end of congestion waves (red). NFV scales naturally with capacity: NFV_{11}^{11} adds only 1728 parameters over NFV_3^1 but achieves significantly better accuracy with similar runtime and memory usage, unlike hand-crafted schemes, which significantly grow in complexity (see for example Appendix A).

7 CONCLUSION

We introduced (U)NFV, a neural network-based framework that extends finite volume methods for hyperbolic conservation laws by learning numerical fluxes over extended spatio-temporal stencils while preserving conservation. (U)NFV achieves high accuracy and efficiency, capturing complex wave dynamics with high fidelity, outperforming classical baselines on standard PDE benchmarks and large-scale field traffic data. Its modular design scales to large spatial and temporal stencils, matching the accuracy of state-of-the-art methods such as DG with much lower implementation complexity and significantly faster inference. In parallel work we have established convergence guarantees by controlling error propagation and deriving bounds on network size and training set requirements. These theoretical results will appear in a forthcoming journal publication. Future directions include applying (U)NFV to velocity-based formulations to learn speed-flux relationships, which aligns with GPS observations and avoids the need for a closed-form velocity PDE. Architecturally, NFV is dimension agnostic, but a systematic multi-dimensional empirical study remains for future work.

540 REPRODUCIBILITY STATEMENT
541

542 To support reproducibility, we provide detailed descriptions of the NFV and UNFV architectures
543 and training objectives in Sections 3 and 4, along with experimental setups in Sections 5 and 6.
544 The formulations of all classical numerical baselines are presented in A, while benchmark equations
545 are introduced in Section 5 and expanded in Appendix B. Additional implementation details,
546 including model architecture, training procedures, hyperparameters, and dataset processing, are
547 provided in Appendix D. Finally, code, datasets, benchmarks, and trained models are released at
548 gregarious-dusk-f2b3d2.netlify.app.

550 ETHICS STATEMENT
551

552 The submission does not have any ethics issues.
553

554 REFERENCES
555

556 Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric
557 deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42,
558 2017.

560 JM Burgers. Math. examples illustrating relations occurring in the theory of turbulent fluid motion.
561 akademie van wetenschappen. *Amsterdam. 1. Sect.*, 17:2, 1939.

562 MARIA Cameron. Notes on the burgers equation. *University of Maryland*, 2011.

564 Christian G Claudel and Alexandre M Bayen. Lax–hopf based incorporation of internal boundary
565 conditions into hamilton–jacobi equation. part i: Theory. *IEEE Transactions on Automatic Control*,
566 55(5):1142–1157, 2010a.

568 Christian G Claudel and Alexandre M Bayen. Lax–hopf based incorporation of internal boundary
569 conditions into hamilton–jacobi equation. part ii: Computational methods. *IEEE Transactions on*
570 *Automatic Control*, 55(5):1158–1174, 2010b.

571 Bernardo Cockburn and Chi-Wang Shu. The local discontinuous galerkin method for time-dependent
572 convection-diffusion systems. *SIAM Journal on Numerical Analysis*, 35(6):2440–2463, 1998.

574 Tim De Ryck, Siddhartha Mishra, and Roberto Molinaro. wpinns: Weak physics informed neural
575 networks for approximating entropy solutions of hyperbolic conservation laws. *SIAM Journal on*
576 *Numerical Analysis*, 62(2):811–841, 2024.

578 Björn Engquist and Stanley Osher. One-sided difference approximations for nonlinear conservation
579 laws. *Mathematics of Computation*, 36(154):321–351, 1981.

580 Lawrence C Evans. *Partial differential equations*, volume 19. American Mathematical Society, 2022.

582 Olga Fuks and Hamdi A Tchelepi. Limitations of physics informed machine learning for nonlinear
583 two-phase transport in porous media. *Journal of Machine Learning for Modeling and Computing*,
584 1(1), 2020.

586 Nikolas Geroliminis and Carlos F Daganzo. Existence of urban-scale macroscopic fundamental
587 diagrams: Some experimental findings. *Transportation Research Part B: Methodological*, 42(9):
588 759–770, 2008.

589 Nikolas Geroliminis and Jie Sun. Properties of a well-defined macroscopic fundamental diagram for
590 urban traffic. *Transportation Research Part B: Methodological*, 45(3):605–617, 2011.

592 Derek Gloudemans, Yanbing Wang, Junyi Ji, Gergely Zachar, William Barbour, Eric Hall, Meredith
593 Cebelak, Lee Smith, and Daniel B Work. I-24 motion: An instrument for freeway traffic science.
594 *Transportation Research Part C: Emerging Technologies*, 155:104311, 2023a.

594 Derek Gloudemans, Gergely Zachár, Yanbing Wang, Junyi Ji, Matt Nice, Matt Bunting, William
 595 Barbour, Jonathan Sprinkle, Benedetto Piccoli, Maria Laura Monache, Alexandre Bayen, Benjamin
 596 Seibold, and Daniel B. Work. So you think you can track? *arXiv preprint arXiv:2309.07268*,
 597 2023b.

598 S. K. Godunov. A finite difference method for the numerical computation of discontinuous solutions
 599 of the equations of fluid dynamics. *Mathematicheckii Sbornik*, 47:271–290, 1959a.

601 SK Godunov. A finite difference method for the computation of discontinuous solutions of the
 602 equations of fluid dynamics. *Sbornik: Mathematics*, 47(8-9):357–393, 1959b.

603 Harold Greenberg. An analysis of traffic flow. *Operations research*, 7(1):79–85, 1959.

604 Bruce D Greenshields, J Rowland Bibbins, WS Channing, and Harvey H Miller. A study of traffic
 605 capacity. In *Highway research board proceedings*, volume 14, pages 448–477. Washington, DC,
 606 1935.

607 Changqing Hu and Chi-Wang Shu. A discontinuous galerkin finite element method for hamilton–
 608 jacobi equations. *SIAM Journal on Scientific computing*, 21(2):666–690, 1999.

609 Ameya D Jagtap and George Em Karniadakis. Extended physics-informed neural networks (xpinns):
 610 A generalized space-time domain decomposition based deep learning framework for nonlinear
 611 partial differential equations. *Communications in Computational Physics*, 28(5), 2020.

612 Ameya D Jagtap, Ehsan Kharazmi, and George Em Karniadakis. Conservative physics-informed
 613 neural networks on discrete domains for conservation laws: Applications to forward and inverse
 614 problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020.

615 Tatiana Kossaczká, Matthias Ehrhardt, and Michael Günther. Enhanced fifth order weno shock-
 616 capturing schemes with deep learning. *Results in Applied Mathematics*, 12:100201, 2021.

617 Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Char-
 618 acterizing possible failure modes in physics-informed neural networks. *Advances in neural
 619 information processing systems*, 34:26548–26560, 2021.

620 Geoff Langdale and Daniel Lemire. Parsing gigabytes of json per second. *The VLDB Journal*, 28(6):
 621 941–960, 2019.

622 PD Lax. The initial value problem for nonlinear hyperbolic equations in two independent variables.
 623 *Ann. Math. Studies*, 33(21):1–229, 1954.

624 PD Lax. Hyperbolic systems of conservation laws ii. *Communications on Pure and Applied
 625 Mathematics*, 10(4):537–566, 1957.

626 Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The
 627 handbook of brain theory and neural networks*, 3361(10):1995, 1995.

628 Randall J LeVeque. *Finite volume methods for hyperbolic problems*, volume 31. Cambridge university
 629 press, 2002.

630 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew
 631 Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations.
 632 *arXiv preprint arXiv:2010.08895*, 2020.

633 Michael James Lighthill and G Be Whitham. On kinematic waves i. flood movement in long rivers.
 634 *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 229
 635 (1178):281–316, 1955.

636 Xu-Dong Liu, Stanley Osher, and Tony Chan. Weighted essentially non-oscillatory schemes. *Journal
 637 of computational physics*, 115(1):200–212, 1994.

638 Bruno Lombard, Denis Matignon, and Yann Le Gorrec. A fractional burgers equation arising in
 639 nonlinear acoustics: theory and numerics. *IFAC Proceedings Volumes*, 46(23):406–411, 2013.

648 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning
 649 nonlinear operators via deeponet based on the universal approximation theorem of operators.
 650 *Nature machine intelligence*, 3(3):218–229, 2021.

651 Xuhui Meng, Zhen Li, Dongkun Zhang, and George Em Karniadakis. Ppinn: Parareal physics-
 652 informed neural network for time-dependent pdes. *Computer Methods in Applied Mechanics and*
 653 *Engineering*, 370:113250, 2020.

654 Toshimitsu Musha and Hideyo Higuchi. Traffic current fluctuation and the burgers equation. *Japanese*
 655 *journal of applied physics*, 17(5):811, 1978.

656 DE Panayotounakos and D Drikakis. On the closed-form solutions of the wave, diffusion and burgers
 657 equations in fluid mechanics. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift*
 658 *für Angewandte Mathematik und Mechanik*, 75(6):437–447, 1995.

659 Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics informed deep learning (part i):
 660 Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*,
 661 2017.

662 Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A
 663 deep learning framework for solving forward and inverse problems involving nonlinear partial
 664 differential equations. *Journal of Computational physics*, 378:686–707, 2019.

665 P. I. Richards. Shock waves on the highway. *Operations Research*, 4:42–51, 1956.

666 Chi-Wang Shu. High order eno and weno schemes for computational fluid dynamics. In *High-order*
 667 *methods for computational physics*, pages 439–582. Springer, 1999.

668 Yunjin Tong, Shiying Xiong, Xingzhe He, Shuqi Yang, Zhecheng Wang, Rui Tao, Runze Liu,
 669 and Bo Zhu. Roenet: Predicting discontinuity of hyperbolic systems from continuous data.
 670 *International Journal for Numerical Methods in Engineering*, 125(6):e7406, 2024.

671 Robin T Underwood. Speed, volume and density relationships. *Quality and theory of traffic flow*,
 672 1961.

673 Feng Wang and Huaping Liu. Understanding the behaviour of contrastive loss. In *Proceedings of the*
 674 *IEEE/CVF conference on computer vision and pattern recognition*, pages 2495–2504, 2021.

675 Yanbing Wang, Derek Gloudemans, Zi Nean Teoh, Lisa Liu, Gergely Zachár, William Barbour,
 676 and Daniel Work. Automatic vehicle trajectory data reconstruction at scale. *arXiv preprint*
 677 *arXiv:2212.07907*, 2022.

678
 679
 680
 681
 682
 683
 684
 685
 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701

702 **A FINITE VOLUME METHODS**
 703

704 Several finite volume-based numerical schemes are studied in this work. They include the following
 705 common classical first-order schemes:
 706

707 **The Godunov method (Godunov, 1959a):**

$$708 \quad 709 \quad 710 \quad 711 \quad 712 \quad \forall i, n \quad \hat{F}_{i-1/2}^n = \begin{cases} \min_{[u_{i-1}^n, u_i^n]} f & \text{if } u_{i-1}^n \leq u_i^n \\ \max_{[u_i^n, u_{i-1}^n]} f & \text{if } u_{i-1}^n > u_i^n \end{cases}$$

713 **The Lax-Friedrichs method (Lax, 1954):**

$$714 \quad 715 \quad 716 \quad \forall i, n \quad \hat{F}_{i-1/2}^n = \frac{1}{2} (f(u_i^n) + f(u_{i-1}^n)) - \frac{1}{2} \frac{\Delta x}{\Delta t} \times |u_i^n - u_{i-1}^n|.$$

717 **The Engquist-Osher method (Engquist and Osher, 1981):**

$$718 \quad 719 \quad 720 \quad \forall i, n \quad \hat{F}_{i-1/2}^n(u_{i-1}^n, u_i^n) = \frac{1}{2} (f(u_i^n) + f(u_{i-1}^n)) - \frac{1}{2} \int_{u_{i-1}^n}^{u_i^n} |f'|.$$

721 Additionally, higher-order schemes such as the **Essentially Non-Oscillatory (ENO) method (Shu, 722 1999)** and the **Weighted Essentially Non-Oscillatory (WENO) method (Liu et al., 1994)** are 723 considered. The main idea in these methods is that by considering more stencils, one can expect to 724 increase the accuracy of approximation of the solution.

725 For the ENO scheme, we consider the semi-discrete form of

$$726 \quad \partial_t u_i = -\frac{1}{\Delta x} (\hat{F}_{i+1/2} - \hat{F}_{i-1/2}). \quad (6)$$

727 Using the Lax-Friedrichs Splitting technique, we have

$$728 \quad 729 \quad f(u) = f^+(u) + f^-(u), \quad f^\pm(u) = \frac{1}{2}(f(u) \pm \alpha u), \quad (7)$$

730 where $\alpha = \max|f'(u)|$ is the maximum wave speed. The key point in the ENO scheme is the 731 high-order upwind interpolation of f^+ and f^- based on the smoothest stencils. For instance, for the 732 2-stencil ENO scheme, the procedure is as follows:

733 1. Evaluate the smoothness indicators:

$$734 \quad 735 \quad \delta_- = |f_i^+ - f_{i-1}^+|, \quad \delta_+ = |f_{i+1}^+ - f_i^+|$$

736 2. Select the stencil that minimizes the smoothness indicator:

- 737 • If $\delta_+ < \delta_-$, choose the stencil $\{f_i^+, f_{i+1}^+\}$.
- 738 • Otherwise, choose the stencil $\{f_{i-1}^+, f_i^+\}$.

739 3. Perform linear interpolation to compute the numerical flux:

$$740 \quad 741 \quad 742 \quad 743 \quad \hat{f}_{i+\frac{1}{2}}^+ = f_i^+ + \frac{1}{2} \delta^+$$

744 where δ^+ is the difference between the selected stencil elements.

745 A similar approach is applied to compute $\hat{f}_{i+\frac{1}{2}}^-$ using the right-biased stencil.

746 The final numerical flux at the interface is obtained by combining the positive and negative parts:

$$747 \quad 748 \quad 749 \quad \hat{f}_{i+\frac{1}{2}} = \hat{f}_{i+\frac{1}{2}}^+ + \hat{f}_{i+\frac{1}{2}}^-$$

750 In this work, we have used a 3-stencil scheme for ENO.

751 The WENO scheme follows the same idea as ENO by using specific weights in defining $\hat{f}_{i+1/2}^+$,
 752 rather than explicit conditions. In this work, we use the 5-stencil WENO scheme.

756 **B VARIANTS OF LWR**
757

758 We consider six different LWR PDEs variants, each consisting of a different fundamental diagram,
759 illustrated in Figure 7. All of the considered flows are concave continuous mappings from $[0, \rho_{\max}]$
760 to \mathbb{R}_+ , where ρ_{\max} is the maximum density, with the exception of the Greenberg flow whose
761 domain is $(0, \rho_{\max}]$. The critical density ρ_c denotes the density at which the flow is maximized, i.e.
762 $\rho_c = \arg \max_{\rho \in [0, \rho_{\max}]} f(\rho)$. The following introduces the six flow models we consider in this work,
763 each time detailing the flow’s parameters, the parameter values we use in Section 5 (in parentheses),
764 and the flow’s definition. Note that we consider normalized parameter values lying between 0 and 1
765 for the most part.

766 **Greenshields** Parameters: free-flow speed v_{\max} (1 m/s), maximum density ρ_{\max} (1 veh/m).
767

$$768 \quad 769 \quad 770 \quad f(\rho) = v_{\max} \rho \left(1 - \frac{\rho}{\rho_{\max}}\right)$$

771 **Triangular 1 (symmetrical)** Parameters: free-flow speed v_{\max} (1 m/s), critical density ρ_c (0.5
772 veh/m), maximum density ρ_{\max} (1 veh/m), wave propagation speed (-1 m/s).
773

$$774 \quad 775 \quad 776 \quad f(\rho) = \begin{cases} v_{\max} \rho & \text{if } \rho < \rho_c \\ w(\rho - \rho_{\max}) & \text{if } \rho \geq \rho_c \end{cases}$$

777 **Triangular 2 (skewed)** A non-symmetric variant of the Triangular flow, with parameters $v_{\max} = 2$
778 m/s, $\rho_c = 1/3$ veh/m, and $w = -1$ m/s.
779

780 **Trapezoidal** Parameters: free-flow speed v_{\max} (1 m/s), first density cusp ρ_1 (0.2 veh/m), second
781 density cusp ρ_2 (0.8 veh/m), maximum density ρ_{\max} (1 veh/m), wave propagation speed (-1.5 m/s).
782

$$783 \quad 784 \quad 785 \quad 786 \quad f(\rho) = \begin{cases} v_{\max} \rho & \text{if } \rho < \rho_1 \\ (w(\rho_2 - \rho_{\max}) - v_{\max} \rho_1) \frac{\rho - \rho_1}{\rho_2 - \rho_1} + v_{\max} \rho_1 & \text{if } \rho_1 \leq \rho \leq \rho_2 \\ w(\rho - \rho_{\max}) & \text{if } \rho > \rho_2 \end{cases}$$

787 **Greenberg** Parameters: maximum density ρ_{\max} (1 veh/m), coefficient c_0 (2).
788

$$789 \quad f(\rho) = c_0 \rho \log(\rho_{\max}/\rho)$$

790 **Underwood** Parameters: maximum density ρ_{\max} (1 veh/m), coefficients c_1 (0.25) and c_2 (1).
791

$$792 \quad f(\rho) = c_1 \rho \exp(1 - c_2 \rho)$$

793 Example solutions of the Greenshields LWR are shown on various initial conditions in Figure 8.
794

795 **C I-24 EXPERIMENTAL DATASET: DENSITY EXTRACTION, TRAINING, AND
796 EVALUATION**
797800 **C.1 DENSITY FIELDS EXTRACTION FROM I-24 MOTION DATASET**
801

802 I-24 MOTION is a large-scale traffic monitoring system installed along a section of Interstate 24
803 near Nashville, Tennessee. It uses a dense network of high-resolution cameras and computer vision
804 algorithms to capture detailed, real-time vehicle trajectories across multiple lanes and miles of
805 highway. The data collection network and resulting trajectory data are illustrated in Figure 9.

806 For our experiments, we use the INCEPTION dataset¹ (Gloudemans et al., 2023a) from I-24 MOTION,
807 consisting of ten days of data, each covering the morning rush hour (6:00 AM to 10:00 AM). The
808 dataset for each day comprises 15-20 GB stored as a single JSON file. We first split each file into
809

¹Available at i24motion.org as part of the INCEPTION data release.

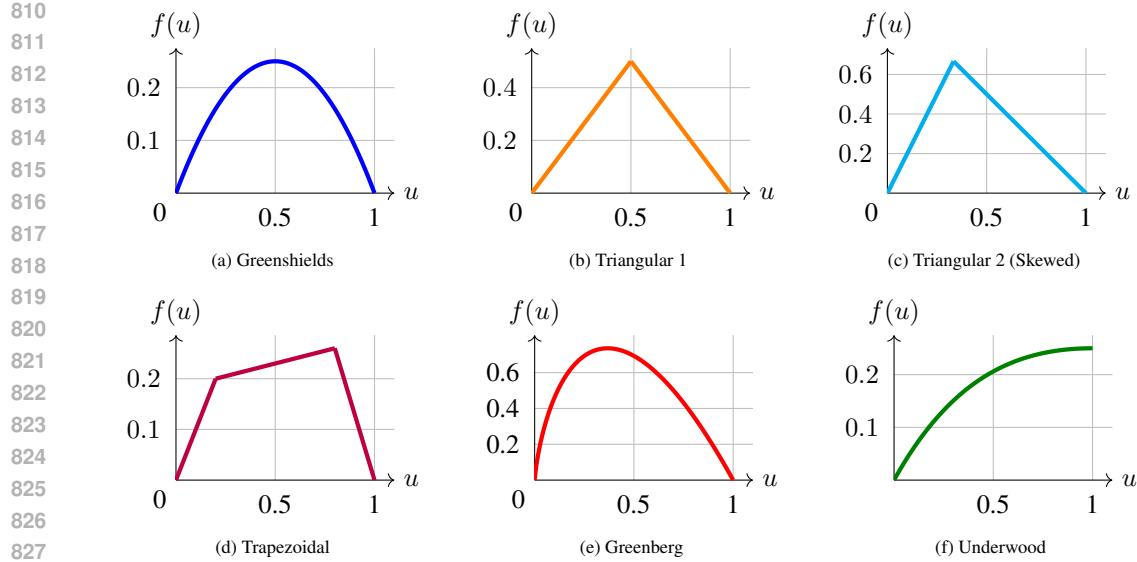


Figure 7: **Flow models for LWR.** We consider six different variants of the LWR PDE with the flows illustrated here, each mapping road density (veh/m) to traffic flow (veh/s).

manageable 1 GB chunks and parse them efficiently using `simdjson` (Langdale and Lemire, 2019), which enables extraction of density fields in approximately 3-5 minutes per 20 GB file.

To construct the density fields, we discretize the spatiotemporal domain into cells of size 0.02 miles (≈ 32 meters) in space and 0.1 seconds in time, aggregating data across all four lanes. Vehicle counts in each cell are normalized to obtain densities in vehicles per kilometer per lane. To reduce noise, we average over 100 consecutive time steps (i.e., 10 seconds) and over 2 adjacent spatial cells (i.e., 0.04 miles or ≈ 64 meters). This results in a grid of 100 spatial cells (4 miles / 0.04 miles) and approximately 1440 time steps (4 hours / 10 seconds). We clip the first and last segments of each day to exclude low-density, free-flow regimes with incomplete data, retaining 1300 time steps per day depending on data quality. To avoid extreme outliers, we cap densities at 140 vehicles/km/lane. For all training and evaluation purposes, we then normalize densities so that the maximum density is 1.

Due to occasional sensor failures, such as malfunctioning camera poles or occlusions by bridges, there are seven spatial locations with missing data. We fill these gaps by linear interpolation between the adjacent upstream and downstream cells.

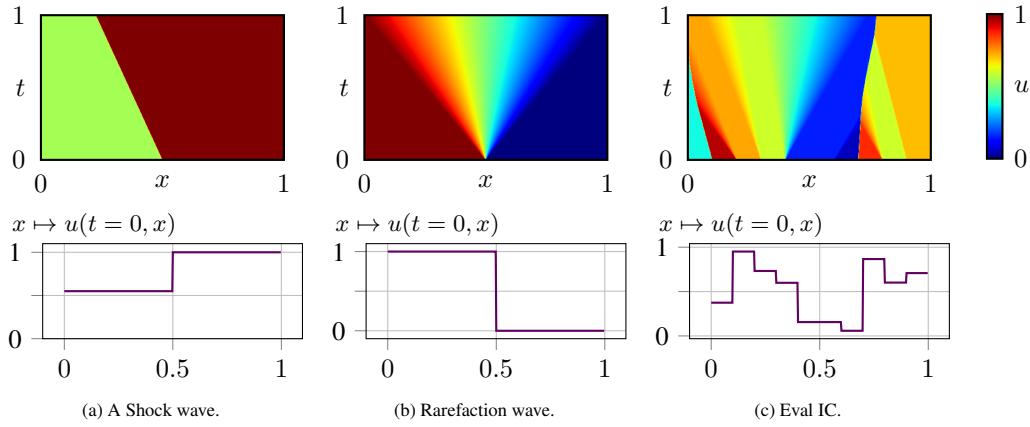


Figure 8: Exact solution for two Riemann problems (left, middle) and one piecewise-constant initial condition (right) from the evaluation set.

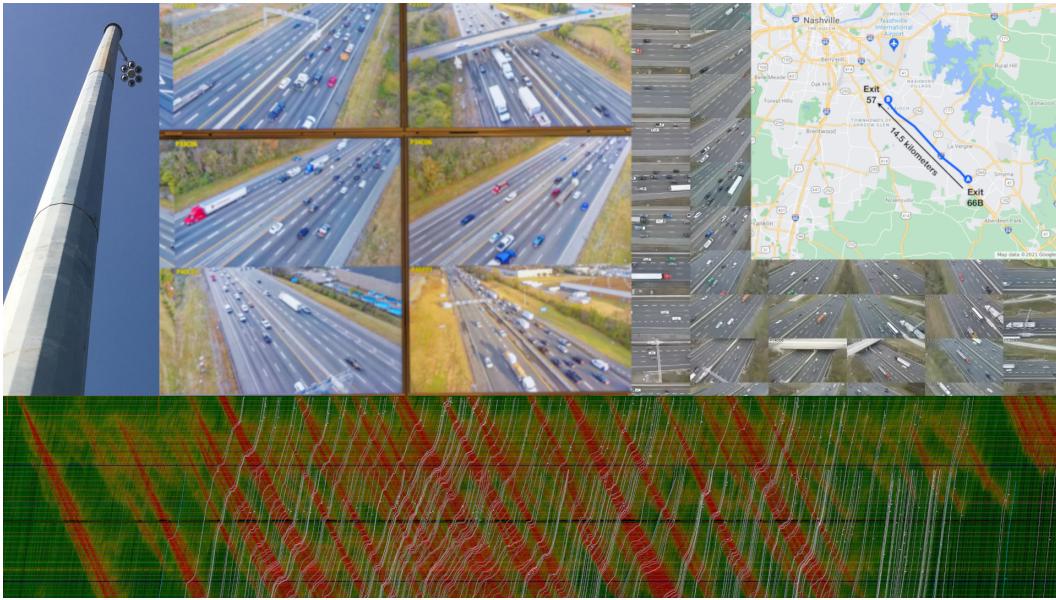


Figure 9: **I-24 MOTION illustration.** High-definition camera poles are mounted along a portion of I-24 at regular intervals. This generates massive amounts of video data, which is processed through a software stack. The resulting data for a single day is shown in the time-space diagram, displaying thousands of individual vehicle trajectories color-coded by speed (red for low speeds, green for high speeds), illustrating the complexity of the dataset.

Figure 10 shows the density fields we extracted from I-24 MOTION data. Higher densities (in red) correspond to stop-and-go waves and congestion, while lower densities (green) correspond to free-flow traffic. The processing code and resulting data are available in our codebase.

We exclude data from November 24 and 25, 2022, from our analysis, as both days correspond to holiday periods with purely free-flow, low-traffic conditions and no observable stop-and-go waves. These days are therefore not relevant to our study, which focuses on modeling traffic dynamics in the presence of congestion. The remaining days still include sufficient free-flow segments to evaluate model robustness in those regimes. Nevertheless, we include the excluded days in the released dataset for completeness.

C.2 BOUNDARY CONDITIONS

For both training and evaluation of NFV on the I-24 dataset, we initialize the model using a single time step of real data and provide one boundary cell at each end of the road, using the corresponding real values. While it is possible to use additional ground truth data to improve accuracy, we deliberately restrict ourselves for two main reasons: (1) to allow fair comparison with the Godunov scheme, which uses a single boundary cell per side, and (2) to reflect realistic deployment scenarios, where boundary densities might only be measured at a few fixed points (e.g., at the road extremities), or predicted using a separate model.

For models that require a wider input stencil (e.g., those larger than NFV_3^1), we pad the boundaries by duplicating the available single-cell values. This ensures that all models, no matter their size: Godunov, NFV_3^1 , or NFV_{11}^{11} , receive the same amount of boundary information. Figure 11 illustrates the boundary setup in both cases, showing which values are provided as input and which are left to be predicted.

Finally, we emphasize that initial and boundary conditions are not included when computing metrics, whether in the training loss or at evaluation.

918 C.3 READING THE HEATMAPS
919920 This section provides a brief explanation and intuition for interpreting the heatmaps displaying I-24
921 MOTION data. The horizontal axis represents time, increasing from left to right, while the vertical
922 axis represents space along the road, increasing from bottom to top. The color encodes traffic density,
923 normalized between 0 and 1, according to the colormap shown in Figure 12, where green indicates
924 low density traffic (free flow) and red indicates high density traffic (congestion). Unless otherwise
925 specified, only the model predictions are shown, while initial and boundary conditions are omitted
926 for clarity. Stop-and-go waves appear as high-density (red) bands that propagate upstream, i.e., move
927 backward through traffic.
928929 C.4 PREDICTIONS
930931 Predictions on the training day from Section 6.2 are displayed in Figure 13. Predictions on evaluation
932 days are displayed in Figure 14.
933934 D EXPERIMENT DETAILS
935936 D.1 MODEL ARCHITECTURE
937938 The model is applied locally on each cell to estimate the corresponding numerical flux. It is
939 implemented as a two-dimensional CNN. The first layer uses a kernel of size $a - 1$, followed by five
940 convolutional layers with 15 channels and kernel size 1. Using a CNN enables efficient vectorized
941 computation over all stencils, which is equivalent to sliding a fully connected network along the input
942 but significantly faster. Each time step is represented as a separate input channel, for a total of b input
943 channels. Note that when $b = 1$, a one-dimensional CNN can be used. The output is a single channel
944 providing the estimated flux. Each NFV model consists of 6 hidden layers of width 15, totaling
945 $1105 + 16 \cdot ((a - 1) \cdot b + 1)$ parameters for NFV_a^b . This remains quite small, with around 1200
946 trainable parameters for the smallest variant, which is intentionally chosen as the smallest architecture
947 that achieves competitive performance [while consistently outperforming first-order FV baselines](#)
948 [across our benchmarks](#).
949950 We also found that activation functions have a modest effect: ELU activations perform slightly better
951 on smooth flow functions (Greenshields, Greenberg, Underwood), while ReLU is preferable for
952 piecewise-linear flows (Triangular, Trapezoidal). However, the difference in performance is minor.
953954 D.2 TRAINING ON SYNTHETIC DATA
955956 For the LWR model, training is performed autoregressively: the model predicts future time steps by
957 feeding its own outputs as inputs. We use 2048 randomly sampled Riemann problems (ρ_1, ρ_2) for
958 training, which proved more effective than uniformly spaced samples. To encourage generalization,
959 small perturbations are added to the discontinuity location. [Empirically, increasing the number of](#)
960 [training Riemann problems beyond this scale did not yield noticeable accuracy gains and mainly](#)
961 [increased training time, indicating that performance in this regime is not limited by data size](#).
962963 The discretization parameters are $\Delta t = 5 \cdot 10^{-3}$, $\Delta x = 10^{-2}$, 100 space cells, and up to 250 time
964 steps. Evaluation is done on 100 test cases generated with a finer grid and the Lax–Hopf algorithm:
965 $\Delta t = 10^{-4}$, $\Delta x = 10^{-3}$, with 200 cells and 1000 time steps. A CFL ratio of 0.5 (e.g. $dx = 10^{-3}$
966 and $dt = 5 \cdot 10^{-4}$) was robust across different flow functions. Higher CFLs sometimes work but
967 were less reliable. For instance, a CFL of 1.0 (e.g. $dx = dt$) is effective for the Greenshield flux but
968 leads to poor performance on the Triangular flux.
969970 The prediction horizon is progressively increased from 10 to 250 steps during training. Most progress
971 occurs at short horizons (10 steps already outperform Godunov on average), while longer horizons
972 provide additional fine-tuning and stability. The learning schedule that proved robust is summarized
973 below:
974975 For unsupervised experiments, we compute the weak loss function using test functions consisting
976 of 250 randomly sampled compactly supported polynomials of degree 50. This proved to work
977 well across PDEs, and details on test function generation are provided in the released codebase.
978

972 Empirically, we observed that performance is largely insensitive to the specific random draws, number,
 973 degree, or family of test functions once these values are moderately large, and that similar results
 974 are obtained with trigonometric test functions; degradation only appears for very small numbers or
 975 degrees (around 1-2), where the weak loss becomes poorly conditioned, so we found that the chosen
 976 setting offers a stable yet memory-efficient default.

Stage	Training steps	Learning rate	n_x	n_t
1	10,000	$1 \cdot 10^{-4}$	10	10
2	20,000	$1 \cdot 10^{-5}$	50	50
3	20,000	$5 \cdot 10^{-6}$	100	100
4	20,000	$1 \cdot 10^{-6}$	200	200

984 Here, (n_x, n_t) denotes the size of the space–time window predicted autoregressively. Most learning
 985 occurs during the first stage, with the later stages serving as progressive fine-tuning. Training uses
 986 the Adam optimizer with a decaying learning rate, from 10^{-4} to 10^{-6} , and the largest batch size that
 987 fits in memory (ideally the entire dataset). Training on an RTX A5000 GPU takes about 30 minutes.
 988

989 D.3 TRAINING ON EXPERIMENTAL DATA

991 All models and fitted finite volume schemes are trained on the first hour of data from November 29,
 992 2022, and evaluated on the full morning period (nearly four hours) and the remaining days of data.
 993 To ensure fairness and reflect practical deployment constraints, each model only receives a single
 994 boundary cell on each side, as described in Appendix C.2, even though larger models could benefit
 995 from additional context.

996 Each NFV model consists of 6 hidden layers of width 15, totaling $1105 + 16 \cdot ((a - 1) \cdot b + 1)$
 997 parameters for NFV_a^b . Training takes 15–30 minutes on an RTX A5000 GPU. The prediction horizon
 998 increases from 10 to 100 steps during training, while the learning rate decays from 10^{-3} to 10^{-4}
 999 over 3000–5000 epochs depending on model size, until convergence.

1000 D.4 HARDWARE AND RUNTIME

1001 All experiments were run on a single NVIDIA A5000 GPU with 24GB of VRAM. The codebase
 1002 (including (U)NFV and baselines) is fully vectorized so that all solutions are computed in parallel.
 1003 Among classical schemes, Lax–Friedrichs is fastest due to its simplicity. Godunov, Engquist–Osher,
 1004 and (U)NFV are 2–3× slower, ENO and WENO are 6× slower, and DG is up to 20× slower, with
 1005 significantly larger memory requirements. Equivalently, NFV runs within a small constant factor of
 1006 Godunov, is roughly twice as fast as ENO and WENO, and more than an order of magnitude faster
 1007 than DG in our benchmarks, while still benefiting from GPU batching across hundreds of solutions
 1008 where DG must be evaluated in smaller batches. These relative runtimes are hardware dependent but
 1009 give a representative picture.

1010 Training is also relatively fast on GPU. The model generally surpasses the Godunov baseline after
 1011 only a few minutes and reaches most of its final performance within 15 minutes. In most runs
 1012 we extended training to one hour or more, though the remaining time typically yields only minor
 1013 fine-tuning. Figure 15 illustrates a typical training curve. To quantify progress, we also compute a
 1014 *winrate* metric, defined as the percentage of evaluation initial conditions on which the model achieves
 1015 lower L_2 error than a baseline such as Godunov. Under this metric, the model usually remains at
 1016 0% for the first few minutes, then rapidly increases to above 95%, and often reaches 100% winrate
 1017 against Godunov within the first 15 minutes of training. A typical winrate curve is also depicted in
 1018 Figure 15.

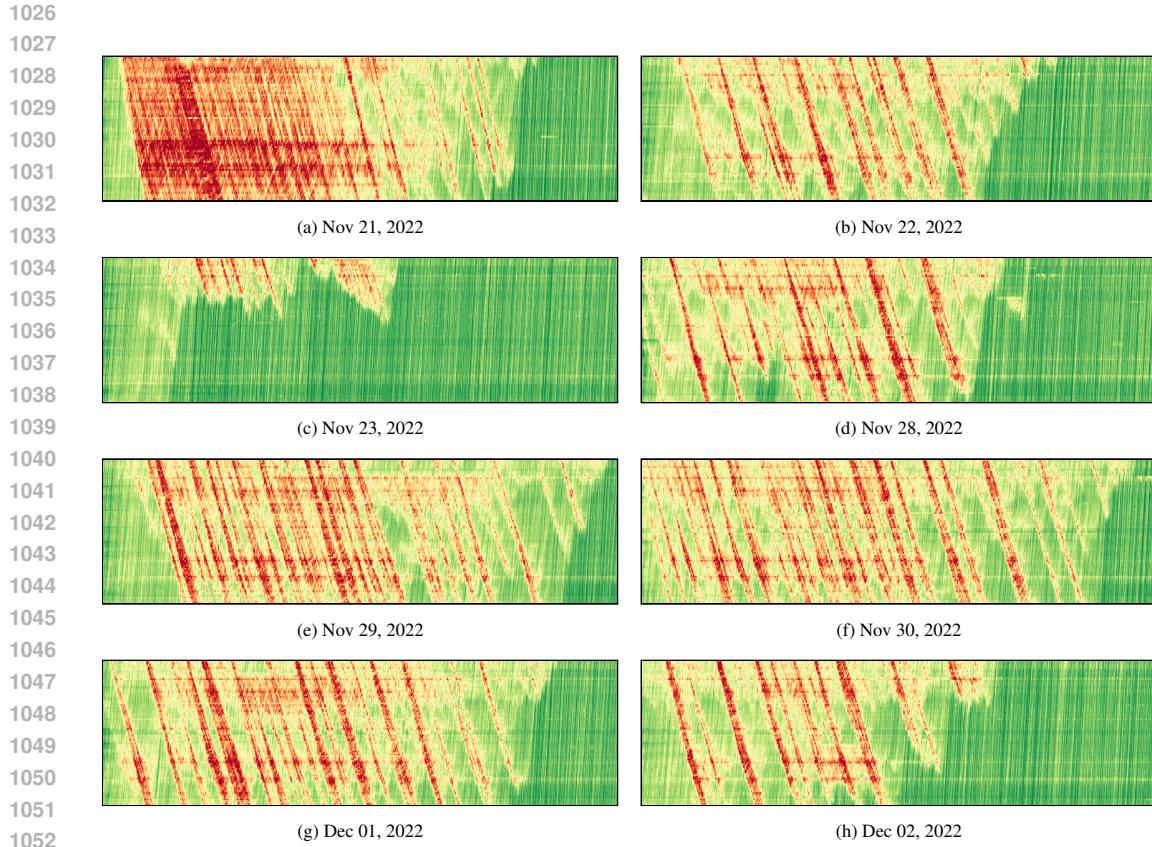


Figure 10: Time-space diagrams of car trajectories extracted from the video, colour-coded by density, for different dates. See Appendix C.3 for how to read the heatmaps.

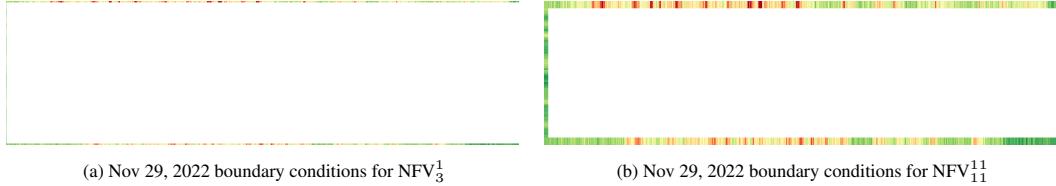


Figure 11: **Boundary conditions used by NFV during training and evaluation on the I-24 dataset.** The figures show the input provided to the model: the initial condition at $t = 0$ on the left, and boundary conditions at $x = 0$ (bottom) and $x = x_{\max}$ (top). The model must then predict the interior (i.e., the region shown in Figure 10) autoregressively: it uses its own output at time t to predict the state at time $t + dt$, without receiving any additional data beyond the fixed boundaries. Note that both figures use the same underlying data; for NFV_{11}^{11} , the boundary values are duplicated to provide the required input padding.

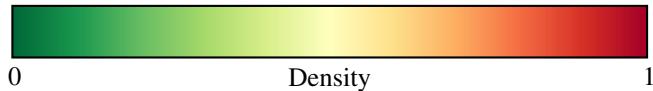


Figure 12: Colorbar showing density scale for all I-24 data heatmaps.

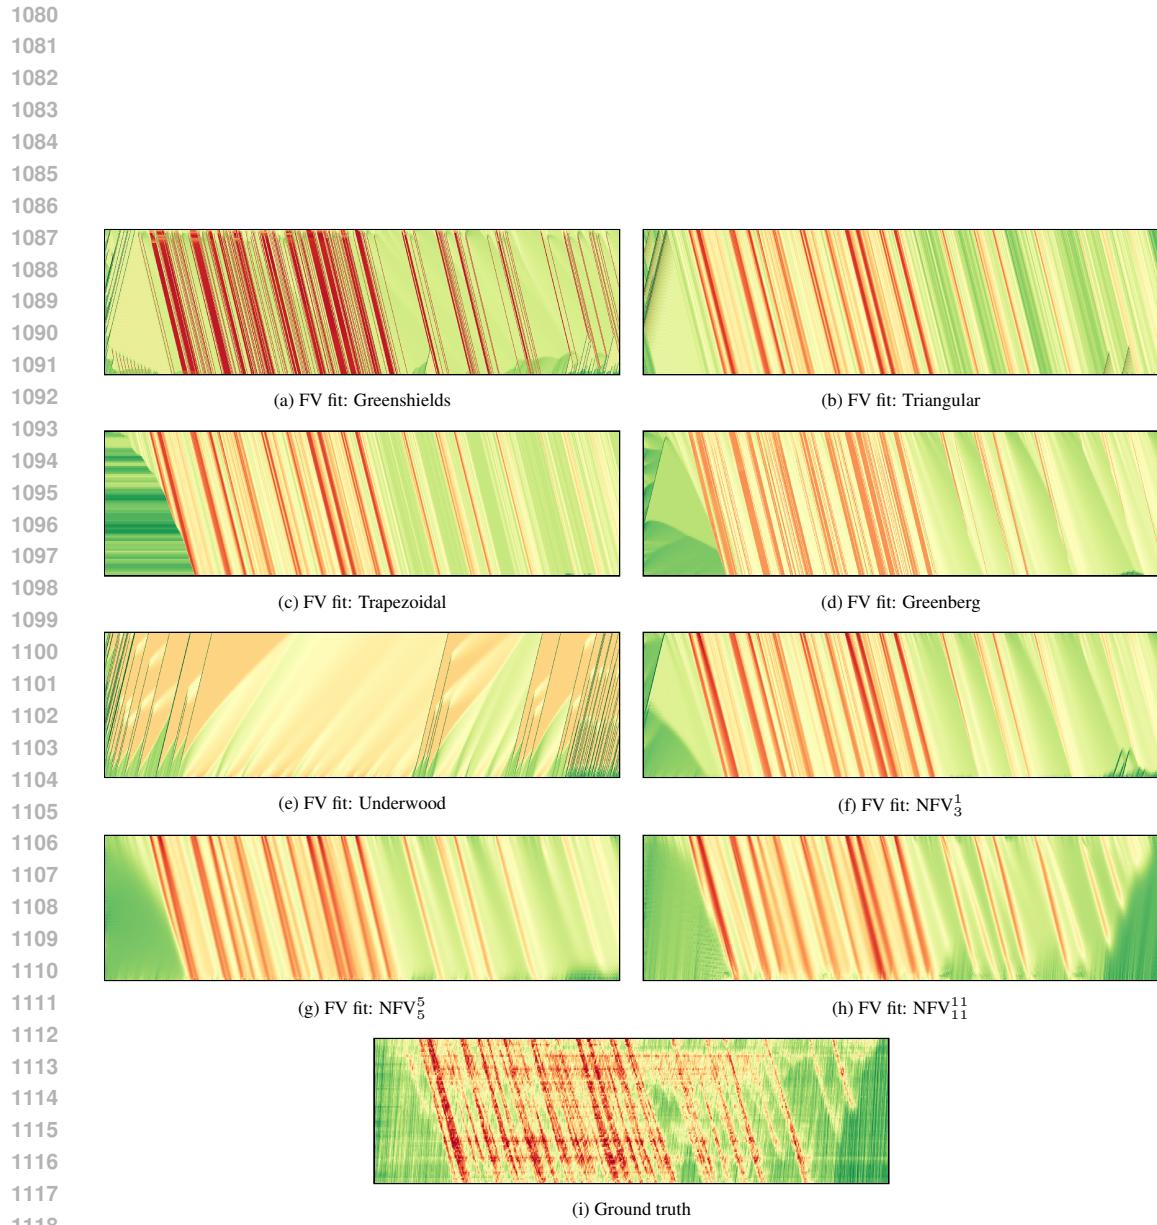
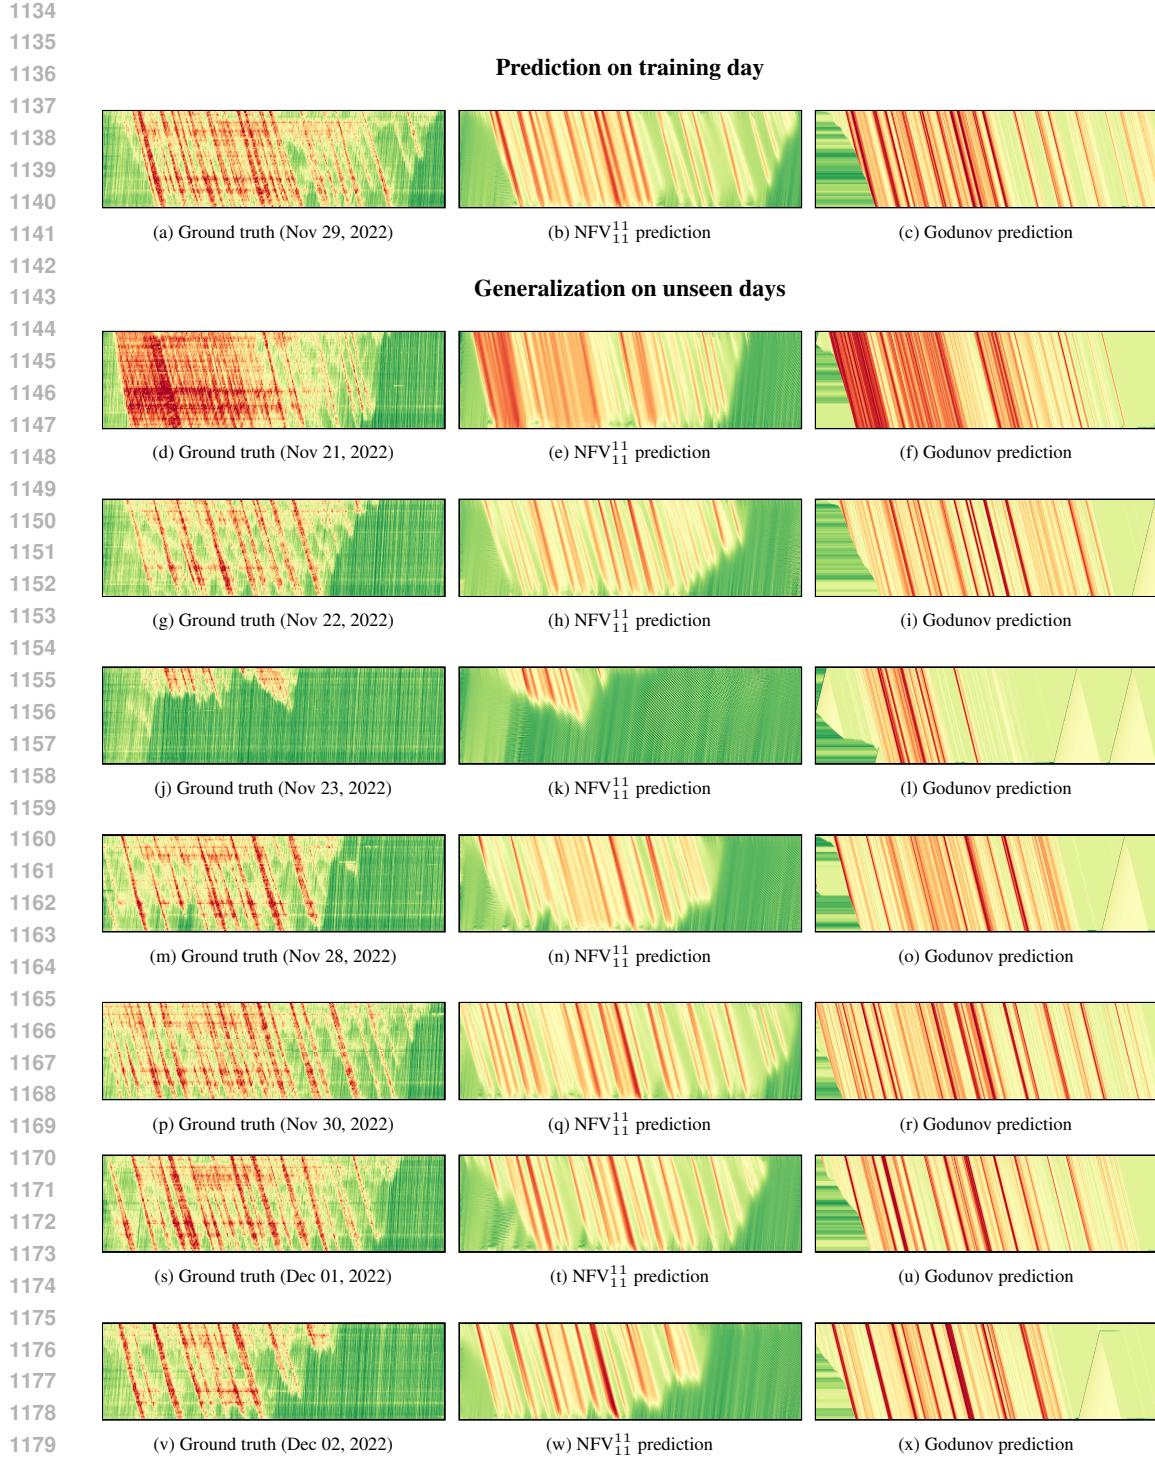


Figure 13: **Predictions of FV methods and trained NFV.** Corresponding metrics are reported in Table 4. Among the FV methods, only the Triangular, Trapezoidal, and Greenberg flows provide a reasonable fit to the I-24 MOTION data. In contrast, NFV models show increasing predictive accuracy with model complexity. For example, NFV₁₁¹¹ captures significantly more stop-and-go waves (in red) than NFV₅⁵ or NFV₃¹, as well as fast low-density waves (in green), enabling it to correctly predict the early dissipation of the final two waves. However, it exhibits oscillations toward the end of the prediction window, likely due to limited generalization caused by the scarcity of low-density (dark green) patterns in the training data; nevertheless, the primary objective when modeling experimental data is to accurately capture the evolution of congestion waves, whereas free-flow traffic is of lesser interest. All models were trained on only the first 25% of the ground truth sequence, and the predictions are generated fully autoregressively. See Appendix C.3 for how to read the heatmaps.



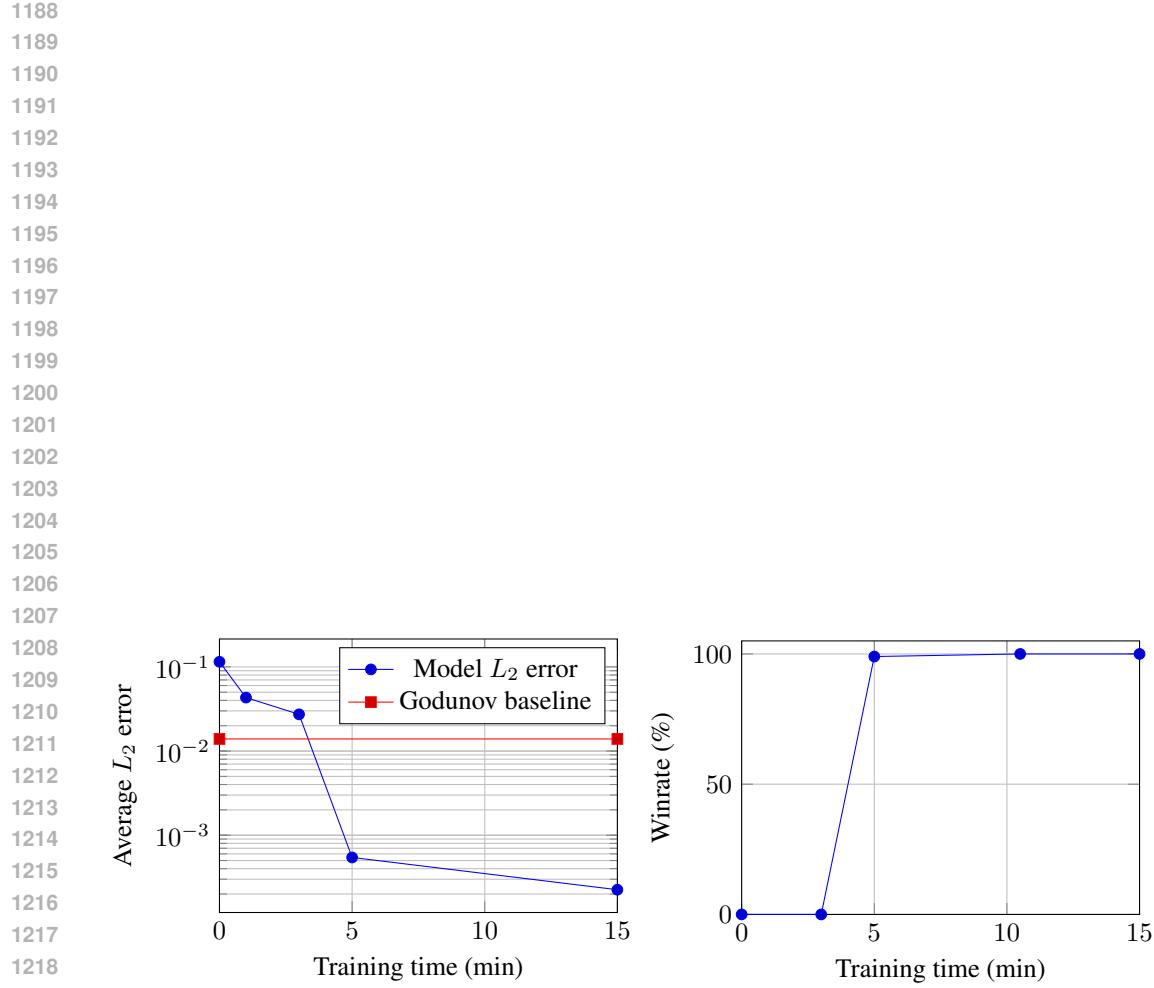


Figure 15: Training dynamics. **Left:** Average L_2 error of the model and of a Godunov baseline. **Right:** Winrate of the model against the Godunov baseline. All the metrics are computed periodically on an evaluation dataset of 100 random complex initial conditions, with prediction over 1000 timesteps.