UPT++: LATENT POINT SET NEURAL OPERATORS FOR MODELING SYSTEM STATE TRANSITIONS

Anonymous authors

004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

Paper under double-blind review

ABSTRACT

Particle methods comprise a wide spectrum of numerical algorithms, ranging from computational fluid dynamics governed by the Navier-Stokes equations to molecular dynamics governed by the many-body Schrödinger equation. At its core, these methods represent the continuum as a collection of discrete particles, on which the respective PDE is solved. We introduce UPT++, a latent point set neural operator for modeling the dynamics of such particle systems by mapping a particle set back to a continuous (latent) representation, instead of operating on the particles directly. We argue via what we call the *discretization paradox* that continuous modeling is advantageous even if the reference numerical discretization scheme comprises particles. Algorithmically, UPT++ extends Universal Physics Transformers - a framework for efficiently scaling neural operators - by novel importance-based encoding and decoding. Furthermore, our encoding and decoding enable outputs that remain consistent across varying input sampling resolutions, i.e., UPT++ is a neural operator. We discuss two types of UPT++ operators: (i) time-evolution operator for fluid dynamics, and (ii) sampling operator for molecular dynamics tasks. Experimentally, we demonstrate that our method reliably models complex physics phenomena of fluid dynamics and exhibits beneficial scaling properties, tested on simulations of up to 200k particles. Furthermore, we showcase on molecular dynamics simulations that UPT++ can effectively explore the metastable conformation states of unseen peptide molecules.

1 INTRODUCTION

In both science and engineering, substantial efforts have led to the formulation of complex mathematical models that accurately represent physical phenomena. Prominent examples include the Navier-Stokes equations, which describe fluid dynamics, and the Schrödinger equation, fundamental to quantum mechanics.



Figure 1: Visualization of a DamBreak3D trajectory. UPT++ successfully captures the characteristics of the evolving fluid, both in terms of predicted occupancy field, and – conditioned on the occupancy field – predicted velocity field. Lighter colors correspond to higher absolute velocities.



Figure 2: The UPT++ modeling paradigm. UPT++ encodes point-cloud information into a continuous latent space representation and decodes this representation at arbitrary query points. This framework enables new ways of simulating particle systems. For example, in our fluid dynamics experiments, particle velocities are sampled at particle positions, whereas in our molecular sampling experiments, densities around atoms are sampled and encoded. The respective latent space operators model the time evolution of the fluid or allow to sample new conformations, respectively. The resulting latent representations are point-wise decoded to occupancies and the corresponding physics information.

073 074

067

068

069

071

072

075

076 The Navier-Stokes equations are the cornerstone of fluid mechanics, and – despite being formulated in the 19th century - continue to present significant challenges to mathematicians and physicists. 077 Most notably, the proof of existence and smoothness of solutions to the Navier-Stokes equations in three dimensions is one of the seven "Millennium Prize Problems" set by the Clay Mathematics 079 Institute (Carlson et al., 2006), with a 1 million prize offered for a solution. What the Navier-Stokes equations are for fluid mechanics, the Schrödinger equation is for quantum mechanics, i.e., the fun-081 damental building block to describe the behavior of particles at atomic and subatomic scales. Unlike 082 classical mechanics, which deals with deterministic paths, the Schrödinger equation (Schrödinger, 083 1926) embraces the probabilistic nature of quantum phenomena, allowing for the superposition of 084 states and the emergence of phenomena like quantum entanglement and tunneling. 085

Solving PDEs numerically is also not straightforward, particularly due to the potential for numer-086 ical instabilities, which can lead to inaccurate results or convergence issues. A successful class of 087 numerical methods to solve certain types of PDEs are particle methods (Pahlke & Sbalzarini, 2023), 880 which represent the underlying continuum media as a collection of discrete particles. For example, 089 for many complex phenomena modeled by the Navier-Stokes equations, e.g., free surface dynamics, or multi-phase flows, Lagrangian discretization schemes are prevalent. Lagrangian methods employ 091 finite material points, often termed particles, whose movement aligns with the local deformation 092 of the continuum (Gingold & Monaghan, 1977; Lucy, 1977; Cundall & Strack, 1979; Brackbill & 093 Ruppel, 1986). Similarly, for molecular dynamics, the Born-Oppenheimer approximation (Born & Oppenheimer, 1927) separates the dynamics of electrons and nuclei, which allows one to move the 094 nuclei – when seen as point particles – according to the laws of classical Newtonian dynamics.

096 The discretization paradox. Contrasting numerical discretization schemes and recent advances in deep learning reveal a subtle paradox. While particle-based discretization schemes are often used to 098 model the most complex natural phenomena, deep learning shines at learning continuous representation, e.g., neural fields (Sitzmann et al., 2020; Mildenhall et al., 2021; Xie et al., 2022), at modeling 099 continuous transformation, e.g., diffusion models (Ho et al., 2020) and flow matching (Lipman et al., 100 2022), and at continuous modulation (Perez et al., 2018; Peebles & Xie, 2023). Naturally, the ques-101 tion arises, why - if the underlying nature is continuous anyway - we don't leverage the power of 102 deep learning on continuous representations? 103

104 Towards this end, we introduce UPT++, which builds on Universal Physics Transformers (Alkin 105 et al., 2024), a latent space neural operator framework for efficiently scaling neural operators to larger physics systems. UPT++ extends Universal Physics Transformers by importance-based en-106 coding and importance-based decoding schemes, which allows us to convert point cloud represen-107 tations of particle methods into continuous latent space representations. Most importantly, encoding and decoding enable outputs that remain consistent across varying input sampling resolutions, which qualifies UPT++ as a neural operator. The UPT++ operator types we are discussing are time-evolution operators – as used for the time-evolution of the Navier-Stokes equations, and sampling operators – applicable to molecular dynamics simulations.

112 113 Our contributions are summarized as follows:

- We connect particle-based numerical discretization schemes and latent space modeling. UPT++ allows us to suggest new modeling paradigms – demonstrated on particle-based fluid dynamics and molecular dynamics simulations.
 - We introduce novel importance-based encoding and importance-based decoding schemes to switch between discretized physics space and continuous latent space.
 - We demonstrate the efficacy and scaling properties of UPT++ on fluid simulations of up to 200k particles, and strong sampling performance on molecular dynamics data.
- 120 121 122

114

115

116

117

118

119

123

2 BACKGROUND: PARTICLE METHODS AND NEURAL OPERATORS

124 Example: SPH discretization in fluid dynamics. The incompressible Navier-Stokes equa-125 tions (Temam, 2001) are defined for the velocity flow field $\boldsymbol{u}: \mathcal{X} \times [0,T] \to \mathbb{R}^3, \mathcal{X} \subset \mathbb{R}^3$, and entail 126 momentum and mass conservation, i.e., $\rho \frac{\mathrm{d} \boldsymbol{u}}{\mathrm{d} t} = \mu \nabla^2 \boldsymbol{u} - \nabla p + \rho \boldsymbol{f}$, and $\nabla \cdot \boldsymbol{u} = 0$, respectively. 127 128 Here, ρ is the density, du/dt is the material derivative, i.e., the rate of change of u of a material 129 element, $\mu \nabla^2 u$ is the viscosity, i.e., the diffusion of u modulated by the viscosity parameter μ , p is 130 the pressure, and f an external force. Lagrangian discretization schemes discretize the continuum 131 via finite material points that move according to the local deformation of the continuum. A famous 132 example is given by smoothed particle hydrodynamics (SPH) proposed by Lucy (1977) and Gingold & Monaghan (1977). SPH approximates the field properties using radial kernel interpolations over 133 adjacent particles at the location of each particle. The strength of the SPH method is that it does not 134 require connectivity constraints, e.g., meshes, which is particularly useful for simulating systems 135 with large deformations, complex fluid-boundary interactions, or free surface flows. 136

137 **Example:** Nuclei discretization in molecular dynamics. The temporal evolution of quantum systems is governed by the many-body Schrödinger equation, for which analytic solutions are hardly 138 139 known and are only available for simple systems like free particles or hydrogen atoms. Classical molecular dynamics (MD) enables the simulation of large atomic systems by various approxima-140 tions reducing the degrees of freedom used to describe the system. A well-known example is the 141 Born-Oppenheimer approximation, where the time-evolving solution is separated into components 142 for the heavier atomic nuclei and the lighter, faster-moving electrons. Further, the positions of the 143 nuclei are updated using numerical integration, where forces on the nuclei are obtained via the nega-144 tive gradient of the potential energy given by an approximate solution of the electronic Schrödinger 145 equation. The potential energy is a parameterized sum of intra- and intermolecular interaction terms; 146 we refer to Appendix B.1 and Frenkel & Smit (2002) for more details on MD. A common applica-147 tion of MD is for sampling conformation states of biomolecules from the Boltzmann distribution, 148 which is also an active area of research in machine learning (Noé et al., 2019). From ergodic theory, 149 we know that, in most cases, MD generates samples from the Boltzmann distribution when the simulation is long enough (Frenkel & Smit, 2002), which, unfortunately, for biomolecules often means 150 simulating for 10^{15} integration steps. Data-driven approaches could accelerate this sampling process 151 by either simulating with larger integration steps or directly sampling from the target distribution. 152

153 **Operator learning**. The operator learning paradigm (Lu et al., 2019; 2021; Li et al., 2020b;a; Ko-154 vachki et al., 2021) targets the approximation of mappings between function spaces. Such function 155 spaces can e.g., comprise the solutions of partial differential equations (PDEs). Following Kovachki et al. (2021), we assume \mathcal{U}, \mathcal{V} to be Banach spaces of functions on compact domains $\mathcal{X} \subset \mathbb{R}^{d_x}$ or 156 $\mathcal{Y} \subset \mathbb{R}^{d_y}$, mapping into \mathbb{R}^{d_u} or \mathbb{R}^{d_v} , respectively. Operator learning aims to approximate the ground 157 truth operator $G: \mathcal{U} \to \mathcal{V}$ via $\hat{\mathcal{G}}: \mathcal{U} \to \mathcal{V}$, typically framed as a supervised learning problem, where 158 input-output pairs are i.i.d. sampled. However, the notable difference is that in operator learning, 159 the space sampled from is not finite-dimensional. More precisely, with a given data set consisting 160 of N function pairs $(u_i, v_i) = (u_i, \mathcal{G}(u_i)) \subset \mathcal{U} \times \mathcal{V}, i = 1, ...N$, we aim to learn $\hat{\mathcal{G}} : \mathcal{U} \to \mathcal{V}$, so 161 that \mathcal{G} can be approximated in a suitably chosen norm.

A widely adopted approach is to approximate \mathcal{G} via three maps (Seidman et al., 2022; Alkin et al., 2024): $\mathcal{G} \approx \hat{\mathcal{G}} := \mathcal{D} \circ \mathcal{A} \circ \mathcal{E}$, comprising encoder \mathcal{E} , approximator \mathcal{A} , and decoder \mathcal{D} . In recent works (Wang et al., 2024; Alkin et al., 2024), the decoder \mathcal{D} has been formulated via point-wise queries at the output grid or mesh. In this work, we focus on two instances of the approximator \mathcal{A} :

166 167

168 169

170

171

186

187 188

189

190

191

192

193

194

196

197

199

200

201

• Time-evolution operators (Seidman et al., 2022; Alkin et al., 2024; Wang et al., 2024), which approximate the deterministic time evolution of a system.

• Sampling operators, which are trained to represent the molecular conformation space (Xu et al., 2022; Jing et al., 2022).

Neural operators for particle systems. Whereas most state-of-the-art neural operator methods are 172 tailored for grid-based, predominantly regular domains, neural operator formulations for particle- or 173 mesh-based dynamics remain limited. In such cases, graph neural networks (GNNs) (Scarselli et al., 174 2008; Kipf & Welling, 2017) with graph-based latent space representations offer a promising alter-175 native. Often, predicted node accelerations are numerically integrated to simulate the time evolution 176 of multi-particle systems (Sanchez-Gonzalez et al., 2020; Mayr et al., 2023; Toshev et al., 2023a). 177 GNNs inherently possess a strong inductive bias for Lagrangian dynamics, which, however, also 178 presents a significant downside since the number of nodes, and thus the computational complexity 179 grows with the number of Lagrangian particles. Thus, computational complexity becomes quickly infeasible for an increasing number of particles (Alkin et al., 2024; Musaelian et al., 2023), see Fig-181 ure 3. Furthermore, the effective degrees of freedom of a particle system are sometimes orders of 182 magnitude less than the degrees of freedom arising from the discretization, especially in simulations showcasing bulk behavior. Lastly, for particle systems, a neural operator formulation is harder, es-183 pecially with respect to the discretization convergence property (Li et al., 2020a), whereas a neural network is expected to show consistency for increasing input sampling resolutions. 185



Figure 3: Qualitative exploration of scaling limits when modeling particle systems. Starting from 4k input points, we compared training time memory requirements of popular Graph Network-based Simulators (Sanchez-Gonzalez et al., 2020) against UPT++. We compare two GNS versions of 5 and 10 layers with hidden dimensions of 64 and 128, respectively. The tested UPT++ model has 30M parameters and can be trained with up to 2M particles on a single A100 40GB GPU.

206

3 UPT++

Our work builds on Universal Physics Transformers (UPTs) (Alkin et al., 2024) - a recently in-207 troduced framework for scaling neural operators, which follows the encoder-approximator-decoder 208 approach. UPT flexibly encodes different grids, and/or a different number of particles into a uni-209 fied latent space representation and introduces a decoder that queries the latent representation at 210 different locations. We adopt the approach of composite neural operators $\mathcal{G} \approx \hat{\mathcal{G}} \coloneqq \mathcal{D} \circ \mathcal{A} \circ \mathcal{E}$, 211 where $\mathcal{E}: \mathcal{U} \to \mathbb{R}^{n_{\text{latent}} \times h}$ maps a solution state $u^t \in \mathcal{U}$ to a latent space state representation 212 $z^t \coloneqq \mathcal{E}(u^t) \in \mathbb{R}^{n_{\text{latent}} \times h}$, i.e., to n_{latent} tokens of dimension $h, \mathcal{A} : \mathbb{R}^{n_{\text{latent}} \times h} \to \mathbb{R}^{n_{\text{latent}} \times h}$ maps the 213 latent state z^t to a successor latent state $z^{t'} \coloneqq \mathcal{A}(z^t) \in \mathbb{R}^{n_{\text{latent}} \times h}, t' > t$, and $\mathcal{D} : \mathbb{R}^{n_{\text{latent}} \times h} \to \mathcal{U}$ 214 reconstructs a solution state $u^{t'} \in \mathcal{U}$ from the latent space state representation $z^{t'} \in \mathbb{R}^{n_{\text{latent}} \times h}$, i.e. 215 $\boldsymbol{u}^{t'} \coloneqq \mathcal{D}(\boldsymbol{z}^{t'}) \in \mathcal{U}.$

216 Deterministic time-evolution: application to fluid dynamics. We consider fluid dynamics prob-217 lems where the fluid is contained within a domain $\Omega \subset \mathcal{X} \subset \mathbb{R}^3$ but does not fill the entire domain 218 \mathcal{X} . The regions of the domain occupied by the fluid Ω change over time governed by the velocity 219 field of the fluid at each given time. Equivalently, we can consider that we are given two disjoint fluid 220 sets (e.g., air and water) filling the entire domain. Then, we consider the solution state u^t to be a compound function, consisting of a velocity field v^t and an occupancy field o^t , i.e., $u^t = (v^t, o^t)^T$. The velocity field v^t maps a certain coordinate $x \in \mathcal{X}$ to a point-wise velocity at this coordinate, i.e., 222 $v^t: \mathcal{X} \mapsto \mathbb{R}^3$. The occupancy field o^t maps a certain coordinate $x \in \mathcal{X}$ to whether a fluid particle 223 is at this coordinate or not, i.e., $o^t : \mathcal{X} \mapsto \{0, 1\}$. u^t is therefore a function $u^t : \mathcal{X} \mapsto \mathbb{R}^3 \times \{0, 1\}$. 224

225 Stochastic conformation sampling: application to molecular dynamics. We assume that a 226 molecule is spatially located in $\mathcal{X} \subset \mathbb{R}^3$ and the specific conformation of a certain molecule is represented by a vector of continuous density fields, i.e., each component of u^t represents one 227 specific density field associated with the conformation of the molecule at time t. Each density field 228 represents some specific characteristics of the molecular conformation, i.e., it might, for example, 229 be specific to a certain atom type. We construct each density field analogously to Pinheiro et al. 230 (2024) and Dumitrescu et al. (2024). Appendix B.2.1 gives further details on the construction of 231 the density fields. When we make use of d density fields to represent the overall conformation of a 232 molecule, then u^t is a function $u^t : \mathcal{X} \mapsto \mathbb{R}^d$. 233

When applying UPT to such settings, we need to address three main questions, leading to UPT++:

- 1. How to encode complex point-cloud representations into continuous representations, and how to decode point-wise from continuous representations?
- 2. How to formulate deterministic and sampling operators in the latent space?
- 3. How to efficiently train UPT++?

234 235

236

237

238

239 240

241

242

3.1 IMPORTANCE-BASED ENCODING, IMPORTANCE-BASED DECODING

Importance-based encoding. We introduce importance-based encoding, which consists of four 243 conceptual steps, see Figure 4. We start with occupancy-based selection, which accounts for the fact 244 that in many simulations, disjoint sets fill the entire domain. For example, for many fluid dynamics 245 phenomena, the state u^t is a compound state of a materialized fluid field, and either obstacles or a 246 second fluid, potentially in the gaseous phase. Similarly, for molecular dynamics, molecules are rep-247 resented within a box, where not the entire box is filled. Secondly, via *importance-based sampling* 248 we emphasize information within the occupied regions. For example, a sampling strategy for particle 249 methods could be to upsample denser regions to account for the larger concentration of mass there. 250 Similarly, we sample points around atoms according to density fields consisting of 3D Gaussian density distribution centered at each atomic position. After importance-based sampling, the encoder 251 \mathcal{E} first embeds the selected k points into hidden dimension h, adding positional encoding (Vaswani 252 et al., 2017) to the different nodes, i.e., $u_k^t \in \mathbb{R}^{k \times d} \to \mathbb{R}^{k \times h}$. Next, via *local aggregation* we 253 propagate neighboring information to the respective supernodes (radius graph for connectivity to 254 keep discretization convergence), and finally global information aggregation pools the information 255 into a fixed size and uniform latent space via perceiver blocks (Jaegle et al., 2021a;b). The resulting 256 continuous latent space contains n_{latent} latent tokens of dimension h, i.e., $z^t := \mathcal{E}(u^t) \in \mathbb{R}^{n_{\text{latent}} \times h}$. 257

Importance-based decoding. Importance-based decoding reverses the conceptual steps of 258 importance-based encoding. First, via occupancy decoding (Mescheder et al., 2019), we decode 259 an occupancy field to identify where particles or atoms are located in space. Secondly, we point-260 wise decode a field quantity and consider only the occupied points. For example, for fluid dynamics, 261 we consider only the flow velocity in regions where occupancy is predicted. Similarly, for molecular 262 dynamics, we decode whether regions are occupied, and consider our density predictions on those. 263 Analogous to Alkin et al. (2024), the decoder is implemented via a perceiver-like cross-attention 264 layer using a positional embedding of the output positions as query and the latent representation 265 as keys and values. Since there is no interaction between queries, the latent representation can be 266 queried at arbitrarily many positions without large computational overhead. To train the occupancy field and enable positive and negative learning signals, we sample points at all locations of the do-267 main. During inference, we simultaneously predict the occupancy field and the corresponding field 268 quantities and only keep the field at occupied points. For the reconstruction of molecular graph 269 structures, we adopt a method similar to that of Pinheiro et al. (2024); Dumitrescu et al. (2024),



Figure 4: Importance-based encoding and decoding schemes of UPT++, which allow us to encode complex point-cloud representations into continuous representations, and reversely enable pointwise decoding of continuous representations.

utilizing an efficient peak-finding algorithm to identify atom positions, and OpenBabel (O'Boyle et al., 2011) to reconstruct the corresponding molecular bonds.

3.2 LATENT SPACE APPROXIMATOR

289 In UPT++, $\mathcal{A} : \mathbb{R}^h \to \mathbb{R}^h$ is a latent operator, which maps the latent state z^t to a successor latent 290 state $z^{t'} \coloneqq \mathcal{A}(z^t) \in \mathbb{R}^h$ (in the case of a deterministic *time-evolution* operator) or which samples a 291 successor latent state $z^{t'} \in \mathbb{R}^h$ from a conditional probability distribution $p(.|z^t)$, i.e., $z^{t'} \sim p(.|z^t)$ 292 (in the case of a stochastic *sampling operator*). 293

Time-evolution operator. The time-evolution operator, implemented as a transformer (Alkin et al., 2024), $\mathcal{A} : \boldsymbol{z}^t \in \mathbb{R}^{n_{\text{latent}} \times h} \to \boldsymbol{z}^{t'} \in \mathbb{R}^{n_{\text{latent}} \times h}$, propagates the compressed representation forward 295 in time. As n_{latent} is small, forward propagation in time is fast. Notably, the approximator can be 296 applied multiple times, propagating the signal forward in time by Δt corresponding to each call of 297 the approximator. After inferring one timestep, it is not necessary to decode the latent state and 298 encode it again to compute the result of a further timestep. Instead, in inference mode, we can 299 keep the latent representation and apply the forward operator again. We call this process *latent* 300 rollout. Especially when working with many particles, the benefits of latent space rollouts, i.e., fast 301 inference, pay off. However, to enable latent rollouts, the responsibilities of encoder \mathcal{E} , approximator 302 \mathcal{A} , and decoder \mathcal{D} need to be decoupled, which is further discussed in Section 3.3. It should be 303 mentioned that our integration timestep Δt is a multiple of the simulation steps used for dataset 304 generation (sometimes up to a factor thousand larger). 305

Sampling operator. The sampling operator that is implemented via flow matching (Lipman 306 et al., 2022), samples $z^{t'}$ from a conditional distribution $p(.|z^t)$, i.e., sampling is conditioned 307 on the latent state z^t . The learning objective for the approximator is to construct a parameter-308 ized (θ) distribution $p_{\theta}(.|\mathbf{z}^t) \approx p(.|\mathbf{z}^t)$. At inference, we draw samples from $p_{\theta}(.|\mathbf{z}^t)$. Thereby 309 we assume an atomistic timestep Δt , which is the minimum timestep for which $p_{\theta}(|z^t)$ was 310 trained to generate meaningful predictions. Δt is usually equal to or a multiple of simulation 311 timesteps. The actual time difference t' - t between prediction time and condition time is usu-312 ally again a multiple of Δt . We implement rectified flow match (Liu et al., 2022) with adaptions according to Li et al. (2024) to include the conditioning on z^t using classifier-free guid-313 ance (Ho & Salimans, 2022). Appendix B.3 details the training and sampling procedures of the 314 UPT++ sampling operator. In contrast to the time-evolution operator, we extend the sampling op-315 erator to also explicitly depend on the number of atomistic timesteps, i.e., we aim to directly learn 316 $p_{\theta}(\boldsymbol{z}^{t'}|\boldsymbol{z}^{t},N) = p_{\theta}(\boldsymbol{z}^{N \Delta t + t}|\boldsymbol{z}^{t},N) \approx p(\boldsymbol{z}^{t'}|\boldsymbol{z}^{t})$ to draw $\boldsymbol{z}^{t'}$ after $N \Delta t$ steps instead of drawing a chain of N consecutive samples $\boldsymbol{z}^{t+\Delta t} \sim p(.|\boldsymbol{z}^{t}), \ldots, \boldsymbol{z}^{t+N \Delta t} \sim p(.|\boldsymbol{z}^{t+(N-1)\Delta t}).$ 317 318

319

280

281

282 283 284

285

286 287

288

320 3.3 UPT++ TRAINING PROCEDURE

321

We make use of the decomposition $\mathcal{D} \circ \mathcal{A} \circ \mathcal{E}$ and split up training into 2 stages, Training. 322 323

keeping in mind the motivation to enable latent rollouts, for which the responsibilities of encoder \mathcal{E} , approximator \mathcal{A} , and decoder \mathcal{D} need to be decoupled. Therefore, at the *first training stage*, we



Figure 5: Sketch of the two types of latent operators we use: PDE forward operator for modeling the time-evolution of the Navier-Stokes equations and the sampling operator applicable to molecular dynamics simulations.

train \mathcal{E} and \mathcal{D} by sampling k input points, and – not necessarily related – k' output points. At this stage, we don't apply the forward operator (therefore t' = t), and the encoder-decoder training can be considered as the training of an autoencoder for the compound state u^t . In the second training stage, we freeze the encoder and the decoder weights, and make use of different timesteps t and t', t' > t, to only train the approximator \mathcal{A} given the fixed latent space input z^t and target output $z^{t'}$. Additionally, for training the latent sampling operator, we regularize the latent space via KL-divergence (Zhang et al., 2023; Rombach et al., 2022).

345 3.4 RELATED WORK

347 In recent years, deep learning has started to make a significant impact in the field of computational 348 fluid dynamics (Guo et al., 2016; Li et al., 2020a; Thuerey et al., 2021; Kochkov et al., 2021; Vinuesa & Brunton, 2022; Gupta & Brandstetter, 2022; Lam et al., 2022; Bi et al., 2022; Brandstetter et al., 349 2022b;a; Andrychowicz et al., 2023; Bodnar et al., 2024; Herde et al., 2024). Several of those works 350 have applied Transformers to physical systems. Galerkin Transformer (Cao, 2021) uses Galerkin-351 type attention to address attention complexity, GNOT (Hao et al., 2023) employs linear attention, 352 OFormer (Li et al., 2023a) uses recurrent MLPs to propagate solutions over time and FactFormer 353 (Li et al., 2023b) uses multidimensional factorized attention. However, all these methods apply 354 attention directly to the input points, which is not scalable in our setting. Transolver (Wu et al., 355 2024) reduces the number of tokens by learning a mapping to physics-aware tokens, a concept 356 similar to our use of supernodes. However, their mapping is recomputed in each Transformer layer, 357 whereas we operate within a fixed latent space. OFormer (Li et al., 2023a) also employs a positional 358 embedding combined with a perceiver to query at arbitrary points and perform decoding. However, their approach applies this process directly to the input, followed by a push forward operation, 359 360 resulting in fixed queries that cannot be altered, thus not suitable for our problem setting. CViT (Wang et al., 2024) is the most similar to our decoding method, but it replaces positional embeddings 361 with learned grid features and uses interpolation to generate the queries. 362

Recent advancements in deep learning have led to an increased interest in its application to molecules. *Boltzmann generators* (Noé et al., 2019; Köhler et al., 2021) employ flows to draw asymptotically unbiased samples from the Boltzmann distribution, but lack the ability to generalize across multiple molecules. Unlike UPT++, current approaches apply sampling either in torsion (Jing et al., 2022) or Euclidean space (Klein et al., 2024a;b; Midgley et al., 2024). Recent breakthroughs in computer vision using compact latent spaces have achieved high sampling quality (Rombach et al., 2022), suggesting the potential of analogous approaches in the molecular domain.

370 371

333

334

335 336 337

338

339

340

341

342

343

344

4 EXPERIMENTS

372 373 374

4.1 LAGRANGIAN FLUID SIMULATION

We conducted experiments on two material point method (MPM) (Sulsky et al., 1995) datasets for our 2D experiments, namely WaterDrop and WaterDrop-XL from Sanchez-Gonzalez et al. (2020), which consist of a maximum of 1.1k and 7k particles, respectively, see Appendix A.4. Additionally, we introduce a new dataset of 3D dam break simulations called DamBreak3D generated using the Riemann SPH method (Zhang et al., 2017) and the SPHinXsys library (Zhang et al., 2021). This
dataset has between 145k-215k fluid particles and consists of 800/100/100 trajectories of length 250
steps, obtained after temporal subsampling at every 100th SPH step, more details in Appendix A.7.

387 **GNNs for large particle systems**. As our main baseline, we choose the established particle-based 388 fluid mechanics surrogate GNS introduced by Sanchez-Gonzalez et al. (2020). However, as seen 389 in Figure 3, such GNN-based approaches do not scale well. To the best of our knowledge, there 390 are three main directions for scaling GNNs to larger particle systems: A) evaluating subgraphs and 391 combining the solutions (Bonnet et al., 2022), B) limiting the receptive field of the neural network 392 and applying domain decomposition (Musaelian et al., 2023; Kozinsky et al., 2023), and C) using a 393 hierarchy of coarser graphs (Qi et al., 2017; Fortunato et al., 2022; Lino et al., 2022). Regarding A, to cover a mesh with 150k nodes, AirFRANS (Bonnet et al., 2022) evaluates 100 randomly sampled 394 395 subgraphs of 32k nodes covering the whole domain and averages the outputs on the nodes that have been evaluated multiple times. Regarding B, Allegro (Musaelian et al., 2023) proposes a novel 396 paradigm which, in contrast to message passing, operates on strictly local neighborhoods to allow for 397 straightforward domain decomposition. Although Allegro allows for simulating systems of arbitrary 398 size, the compute requirement scales linearly with the system size – in a scaling example, Allegro 399 distributes 100M atoms over 5k GPUs or roughly 20k atoms per GPU (Kozinsky et al., 2023). Thus, 400 both A and B approaches have a linear or worse scaling of compute with respect to system size. 401 As we aim to develop a framework that scales to at least 200k particles (in our experiments), the 402 hierarchical approach C seems most suitable. Thus, to have a competitive baseline, we develop a 403 multi-scale version of GNS, called MS-GNS, which couples the finer (original) particles with coarser 404 particles consisting of randomly subsampled 12.5% of the finer particles. Our approach is inspired 405 by MS-MGN (Fortunato et al., 2022) and connects the two point clouds by a k-nearest neighbors graph with k = 4 from the fine to the coarse nodes, see Appendix A.5 for more details. 406

407 Model details. In our experiments, UPT++ encodes the first two velocities of the trajectory and 408 the time-evolution operator acts only in the latent space. In contrast, GNS encodes the first five 409 velocities and then autoregressively predicts and integrates the accelerations. Based on the GNS 410 ablation studies in Sanchez-Gonzalez et al. (2020) and Toshev et al. (2023b), we choose to train 411 a model with 10 message-passing (MP) steps and a latent size of 128, denoted by GNS-10-128 in 412 Table 1. With MS-GNS-15, we denote an MS-GNS model with a processor consisting of: 1 MP layer on the fine particles with a latent size 64, 1 downsampling layer, 11 MP layers on the coarse 413 graph with a latent size 128, 1 upsampling layer, and 1 MP layer as the first one. 414

415 **Results**. Our main results are summarized in Table 1, showing that UPT++ performs comparably to 416 GNNs in terms of both IoU and MSE, but can offer more than 50x greater speedups. Notably, we 417 work with the smaller GNS-5-64 model on DamBreak3D as this is the biggest GNS model we could train with one sample per 40GB GPU, compare Figure 3. Qualitatively, the lower *IoU* of UPT++ 418 on the 2D datasets is related to the lack of mass conversation (equivalently volume conservation, 419 as we work with incompressible fluids) in the latent state representation, which manifests itself 420 in having too little or too much fluid along a trajectory. We note that volume conservation is a 421 problem that GNNs also have, as recently discussed in Neural SPH (Toshev et al., 2024), but in 422 contrast to UPT++, GNNs, by construction, preserve the mass, i.e. the number of particles. On the 423 other hand, UPT++ learns a better representation of the velocity field, which we suspect is easier to 424 learn as a continuous function. Figure 1 shows one exemplary trajectory rollout of DamBreak3D, 425 demonstrating that UPT++ can adequately model a 3D particle simulation with 200k particles. 426

427 428

4.2 SAMPLING MOLECULAR CONFORMATIONS

429 430

431 We apply UPT++ to the task of sampling molecular conformations. Molecular conformations are, e.g., important when studying interactions between different molecules or when deriving certain

Table 1: MSE denotes the mean-squared error of the velocity prediction. IoU and MSE are averaged
over all timesteps and all trajectories in the test set of each dataset. UPT++ can model the complex
dynamics while providing a significant speedup.

Dataset	Method	Hardware	Rollout time	Speedup	$IoU (\uparrow)$	$MSE (\downarrow)$
	MPM	6 CPUs	50s	1x	-	-
WaterDrop	GNS-10-128	A40	4.0s	13x	0.91	0.047
-	UPT++	A40	0.53s	94x	0.87	0.036
WaterDrop-XL	MPM	6 CPUs	170s	1x	-	-
	GNS-10-128	A40	44s	4x	0.83	0.16
	UPT++	A40	0.65s	262x	0.81	0.16
	SPH	32 CPUs	1200s	1x	-	-
DamBreak3D	GNS-5-64	A40	100s	12x	0.83	0.54
	MS-GNS-15	A6000	150s	8x	0.91	0.18
	UPT++	A40	2.7s	444x	0.93	0.14

properties for molecules. This might be especially relevant for biomedical chemistry and drug discovery.

We benchmark UPT++ on two small peptide datasets from Klein et al. (2024a), namely the alanine dipeptide dataset (AD), containing a single molecule of 22 atoms, and, further on a small peptide dataset (2AA) containing 400 different peptides, with varying number of atoms (20-50) and different atom types per molecules. We stick to the suggested train/test split provided with 2AA and either evaluate on the whole test set or on an exemplary molecule (AN) from the test set.

457 **Metrics**. We evaluate the performance of UPT++, implemented via a latent sampling operator, by 458 investigating associated Ramachandran plots (Ramachandran et al., 1963) for the sampled molecule 459 conformations. These plots show the distribution of peptide dihedral angles ϕ and ψ . We quantify 460 the differences of marginalized distributions of angles between our sampled molecule conformations 461 and those ones from a reference simulation in analogy to Yu et al. (2024) by means of the Jenson-462

Model specific details and training details. We use a guided flow-matching model (as explained above) with the same diffusion transformer backbone (Peebles & Xie, 2023) as for the time-evolution operator. As outlined above, we do not only condition on a previous molecule conformation at time *t*, but also on the number of atomistic timesteps *N*. The idea is partly based on ITO (Schreiner et al., 2023). Further, we apply an MD relaxation step before reconditioning (see Appendix B.2.3). Details on importance-based sampling for molecules and on data augmentation we use can be found in Appendices B.2.2 and B.2.4.

Results. Figure 6 shows Ramachandran plots and free energy projection plots for AD conforma-470 tions for 13.6k UPT++ sampled conformations and 800k MD reference simulation conformations, 471 respectively. The Ramachandran plots indicate that modes of reference conformation angles are 472 faithfully restored by sampled UPT++ conformations. The free energy projections show that our 473 model also captures energy minima very well, and that less likely regions in conformation space 474 are explored. For exemplary molecules (AN) from the test set of 2AA Ramachandran plots and free 475 energy projection plots are shown in Figures B2 and B3 in Appendix B.5. All AN plots rely on 476 10k UPT++ sampled conformations and 9.8k MD reference simulation conformations. We investigate the influence of the flow-matching guidance parameter and observe that less guidance seems 477 to capture modes of conformation angles better than high guidance. This is also reflected by the 478 Jensen–Shannon divergences shown in Table 2. 479

Extracting molecule graphs after decoding signals from latent space. To assess if molecule
 graph extraction from the latent space happens uniquely, we evaluate the graph extraction performance of UPT++ in two scenarios: first, as a strict autoencoder, i.e. encoding and decoding without
 the sampling operator, and second, with the sampling operator applied. We consider the extraction
 (for details on the final graph extraction see Appendix B.4) of a molecule as valid when encoded
 and decoded versions have equal InChI (Landrum, 2016) codes, utilizing the RDKit library (Karol, 2018). This ensures chemical identity is preserved. Table 2 shows that the molecule encoder and

Table 2: Summary of results on molecular sampling. We show the reconstruction success rate without applying the sampling operator ("enc-dec"), and with sampling ("sampling+dec").

Dataset	Guidance	Reconstr	. success rate (\uparrow)	JS Ramachandran (↓) (sampling)	
Dutuset	scale	enc-dec	sampling + dec		
AD	3.5	1.0	0.97	0.18	
2AA: AN	1.0	1.0	0.20	0.30	
2AA: AN	3.5	1.0	0.50	0.17	

decoder of UPT++ can accurately extract molecules from the latent space for Alanine dipeptide and all test molecules from the 2AA dataset (success rate 100% when tested with 1000 samples). When applying the sampling operator, the extraction performance decreases slightly (0.97). Results for the AN molecule of the 2AA dataset are worse, which results from the fact that the tested molecules are unseen during training. Further, higher guidance scales increase reconstruction success rates but yield less diverse samples. This behavior is analogous to classifier-free guidance in other domains.



Figure 6: Alanine dipeptide experiments. Left half: Ramachandran plots comparing 14k UPT++ and MD samples. Right half: Free energy surface for the two dihedral angles ψ and ϕ with the same 14k UPT++ samples but 800k MD samples. Our model captures well the energy minima and also explores less probable regions of sample space.

5 CONCLUSION, LIMITATIONS AND FUTURE WORK

We presented UPT++, a generic framework for modeling simulations that are conventionally discretized with particles. UPT++ maps the state of the system to a fixed-sized latent space with novel importance-based encoding and importance-based decoding techniques. The strengths of our approach are its generic architecture, favorable scaling to larger systems, and potential for significant acceleration of numerical simulations. We demonstrated these strengths by training a Lagrangian fluid dynamics surrogate on up to 200k particles, as well as on molecular conformer generation across different peptide molecules.

Amongst others, possible extensions of UPT++ are adoptions to more challenging engineering prob-lems, e.g., complex geometries, solid-liquid interactions, or multi-physics. The fact that the size of latent representations is constant in UPT++ and, therefore, in principle independent of concrete molecule sizes, suggests its application to larger peptides or other larger molecules. One weakness of a field-based approach like UPT++ is that – by construction – UPT++ is not mass conserving. This is in contrast to GNNs, which preserve the mass, i.e., the number of particles. Further, for molecular conformation sampling, we decided to incorporate the conditioning conformations via a classifier-free guidance approach since directly conditioning on previous conformations is problematic, as discussed in literature (Li et al., 2024). We leave the exploration of other conditioning strategies to future work. Lastly, improving the decoding scheme could significantly speed up con-former sampling.

540 ETHICS STATEMENT

Accuracy and Reliability of Simulations. While UPT++ offers a computationally efficient alter native to traditional simulation methods, there is a risk that inaccuracies in the approximations could
 lead to unintended consequences. For example, in a civil engineering context, designs for flood pro tection could rely significantly on the accuracy of our simulations. We strongly emphasize that our
 models should not be blindly relied upon for decision-making in safety-critical areas. Users should
 always corroborate machine learning predictions with established physical models or additional empirical data. Similarly for molecular simulations, the output of our models should be checked with
 experiments or complemented with classical physical simulations before decision-making.

Transparency and Explainability. Given that UPT++ encodes the system's state into a continuous latent representation, it may offer less transparency and explainability compared to methods that directly operate on the physical state. Our latent space approach can make it difficult for users to fully understand how certain predictions or decisions are reached. Lack of interpretability could lead to challenges in trusting and verifying the model's outputs, particularly in critical applications. To mitigate these concerns, we advocate for developing methods that enhance explainability and allow users to inspect and understand the underlying decision processes of UPT++, ensuring its safe and responsible use in real-world scenarios.

Environmental and Social Impact. Our models could have significant societal and environmental impacts. For example, in cases like flood prediction or water resource management, inaccurate predictions may lead to poor planning or resource allocation, disproportionately affecting vulnerable communities. We urge that such models be used responsibly, with attention to fairness, inclusivity, and transparency, especially in areas that affect public health, safety, and well-being.

564 REPRODUCIBILITY STATEMENT

We provide a detailed description of the model in Section A.3 and B.2, along with a comprehensive table for the hyperparameters used in Lagrangian fluid simulations in Table A1 and the hyperparameters used in sampling molecular conformations in Table B1. The data for experiments conducted on WaterDrop and WaterDrop-XL can be found in Sanchez-Gonzalez et al. (2020). The DamBreak3D dataset will be released soon. The data for experiments conducted on the MD experiments can be found in Klein et al. (2024a). Additionally, we provide the anonymized code in the supplementary materials. The finalized code, along with all model checkpoints used in the experiments, will be made publicly available on GitHub.

594 REFERENCES 595

610

616

623

633

- Stefan Adami, Xiangyu Hu, and Nikolaus A Adams. A generalized wall boundary condition for 596 smoothed particle hydrodynamics. Journal of Computational Physics, 231(21):7057–7075, 2012. 597
- 598 Stefan Adami, XY Hu, and Nikolaus A Adams. A transport-velocity formulation for smoothed particle hydrodynamics. Journal of Computational Physics, 241:292–307, 2013. 600
- 601 Benedikt Alkin, Andreas Fürst, Simon Schmid, Lukas Gruber, Markus Holzleitner, and Johannes Brandstetter. Universal physics transformers: A framework for efficiently scaling neural opera-602 tors. arXiv preprint arXiv:2402.12365, 2024. 603
- 604 Marcin Andrychowicz, Lasse Espeholt, Di Li, Samier Merchant, Alex Merose, Fred Zyda, Shreya 605 Agrawal, and Nal Kalchbrenner. Deep learning for day forecasts from sparse observations. arXiv 606 preprint arXiv:2306.06079, 2023. 607
- 608 Carla Antoci, Mario Gallati, and Stefano Sibilla. Numerical simulation of fluid-structure interaction by sph. Computers & structures, 85(11-14):879-890, 2007. 609
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. arXiv preprint 611 arXiv:1607.06450, 2016. 612
- 613 Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Pangu-weather: 614 A 3d high-resolution model for fast and accurate global weather forecast. arXiv preprint 615 arXiv:2211.02556, 2022.
- Cristian Bodnar, Wessel P Bruinsma, Ana Lucic, Megan Stanley, Johannes Brandstetter, Patrick 617 Garvan, Maik Riechert, Jonathan Weyn, Haiyu Dong, Anna Vaughan, et al. Aurora: A foundation 618 model of the atmosphere. arXiv preprint arXiv:2405.13063, 2024. 619
- 620 Florent Bonnet, Jocelyn Mazari, Paola Cinnella, and Patrick Gallinari. Airfrans: High fidelity com-621 putational fluid dynamics dataset for approximating reynolds-averaged navier-stokes solutions. 622 Advances in Neural Information Processing Systems, 35:23463–23478, 2022.
- M. Born and R. Oppenheimer. Zur quantentheorie der molekeln. Annalen der Physik, 389(20): 624 457-484, 1927. ISSN 0003-3804. doi: 10.1002/andp.19273892002. 625
- 626 Jeremiah U Brackbill and Hans M Ruppel. Flip: A method for adaptively zoned, particle-in-cell 627 calculations of fluid flows in two dimensions. Journal of Computational physics, 65(2):314-343, 628 1986. 629
- Johannes Brandstetter, Rianne van den Berg, Max Welling, and Jayesh K Gupta. Clifford neural 630 layers for pde modeling. arXiv preprint arXiv:2209.04934, 2022a. 631
- 632 Johannes Brandstetter, Daniel Worrall, and Max Welling. Message passing neural pde solvers. arXiv preprint arXiv:2202.03376, 2022b. 634
- 635 Shuhao Cao. Choose a transformer: Fourier or galerkin. Advances in neural information processing 636 systems, 34:24924-24940, 2021.
- James Carlson, Arthur Jaffe, and Andrew Wiles. The Millennium Prize Problems. Clay Mathematics 638 Institute and American Mathematical Society, 2006. 639
- 640 D.A. Case, H.M. Aktulga, K. Belfon, I.Y. Ben-Shalom, J.T. Berryman, S.R. Brozell, D.S. Cerutti, 641 T.E. Cheatham, III, G.A. Cisneros, V.W.D. Cruzeiro, T.A. Darden, N. Forouzesh, M. Ghaz-642 imirsaeed, G. Giambaşu, T. Giese, M.K. Gilson, H. Gohlke, A.W. Goetz, J. Harris, Z. Huang, 643 S. Izadi, S.A. Izmailov, K. Kasavajhala, M.C. Kaymak, A. Kovalenko, T. Kurtzman, T.S. Lee, 644 P. Li, Z. Li, C. Lin, J. Liu, T. Luchko, R. Luo, M. Machado, M. Manathunga, K.M. Merz, Y. Miao, O. Mikhailovskii, G. Monard, H. Nguyen, K.A. O'Hearn, A. Onufriev, F. Pan, S. Pan-645 tano, A. Rahnamoun, D.R. Roe, A. Roitberg, C. Sagui, S. Schott-Verdugo, A. Shajan, J. Shen, 646 C.L. Simmerling, N.R. Skrynnikov, J. Smith, J. Swails, R.C. Walker, J. Wang, J. Wang, X. Wu, 647 Y. Wu, Y. Xiong, Y. Xue, D.M. York, C. Zhao, Q. Zhu, and P.A. Kollman. Amber 2024, 2024.

661

669

685

- Andrea Colagrossi and Landrini Maurizio. Numerical simulation of interfacial flows by smoothed particle hydrodynamics. *Journal of Computational Physics*, 191, 2003.
- Peter A Cundall and Otto DL Strack. A discrete numerical model for granular assemblies. *geotechnique*, 29(1):47–65, 1979.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Alexandru Dumitrescu, Dani Korpela, Markus Heinonen, Yogesh Verma, Valerii Iakovlev, Vikas
 Garg, and Harri Lähdesmäki. Field-based molecule generation. *arXiv preprint arXiv:2402.15864*, 2024.
- Peter Eastman, Jason Swails, John D. Chodera, Robert T. McGibbon, Yutong Zhao, Kyle A. Beauchamp, Lee-Ping Wang, Andrew C. Simmonett, Matthew P. Harrigan, Chaya D. Stern, Rafal P. Wiewiora, Bernard R. Brooks, and Vijay S. Pande. Openmm 7: Rapid development of high performance algorithms for molecular dynamics. *PLOS Computational Biology*, 13(7): 1–17, 07 2017. doi: 10.1371/journal.pcbi.1005659.
- 667 GW Ford, M Kac, and P Mazur. Statistical mechanics of assemblies of coupled oscillators. *Journal* 668 *of Mathematical Physics*, 6(4):504–515, 1965.
- Meire Fortunato, Tobias Pfaff, Peter Wirnsberger, Alexander Pritzel, and Peter Battaglia. Multiscale
 meshgraphnets. *arXiv preprint arXiv:2210.00612*, 2022.
- Daan Frenkel and Berend Smit. Understanding Molecular Simulation: From Algorithms to Applications. Number 1 in Computational Science Series. Academic Press, San Diego, 2nd ed edition, 2002. ISBN 978-0-12-267351-1.
- Robert A Gingold and Joseph J Monaghan. Smoothed particle hydrodynamics: theory and application to non-spherical stars. *Monthly notices of the royal astronomical society*, 181(3):375–389, 1977.
- Kiaoxiao Guo, Wei Li, and Francesco Iorio. Convolutional neural networks for steady flow approximation. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 481–490, 2016.
- Jayesh K Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized pde
 modeling. *arXiv preprint arXiv:2209.15616*, 2022.
- Zhongkai Hao, Zhengyi Wang, Hang Su, Chengyang Ying, Yinpeng Dong, Songming Liu,
 Ze Cheng, Jian Song, and Jun Zhu. Gnot: A general neural operator transformer for operator
 learning. In *International Conference on Machine Learning*, pp. 12556–12569. PMLR, 2023.
- Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel
 de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes. *arXiv preprint arXiv:2405.19101*, 2024.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. arXiv preprint
 arXiv:2207.12598, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, 2020.
- Håkon Hoel and Anders Szepessy. Classical langevin dynamics derived from quantum mechanics.
 arXiv preprint arXiv:1906.09858, 2019.
- 701 XY Hu and Nikolaus A Adams. An incompressible multi-phase sph method. Journal of computational physics, 227(1):264–278, 2007.

702 Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David 703 Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver io: A 704 general architecture for structured inputs & outputs. In International Conference on Learning 705 Representations, 2021a. 706 Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In International conference on machine 708 learning, pp. 4651-4664. PMLR, 2021b. 710 Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional diffusion for molecular conformer generation. Advances in Neural Information Processing Systems, 711 35:24240-24253, 2022. 712 713 Paul J. Karol. The inchi code. Journal of Chemical Education, 95(6):911–912, Jun 2018. ISSN 714 0021-9584. doi: 10.1021/acs.jchemed.8b00090. 715 716 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), 3rd International Conference on Learning Representations, ICLR 2015, 717 San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. 718 719 Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional net-720 works. In International Conference on Learning Representations, 2017. 721 Leon Klein, Andrew Foong, Tor Fjelde, Bruno Mlodozeniec, Marc Brockschmidt, Sebastian 722 Nowozin, Frank Noé, and Ryota Tomioka. Timewarp: Transferable acceleration of molecular 723 dynamics by learning time-coarsened dynamics. Advances in Neural Information Processing 724 Systems, 36, 2024a. 725 726 Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. Advances in Neural Information Processing Systems, 36, 2024b. 727 728 Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan 729 Hoyer. Machine learning-accelerated computational fluid dynamics. Proceedings of the National 730 Academy of Sciences, 118(21):e2101784118, 2021. 731 Jonas Köhler, Andreas Krämer, and Frank Noé. Smooth normalizing flows. Advances in Neural 732 Information Processing Systems, 34:2796–2809, 2021. 733 734 Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, An-735 drew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces. 736 arXiv preprint arXiv:2108.08481, 2021. 737 Boris Kozinsky, Albert Musaelian, Anders Johansson, and Simon Batzner. Scaling the leading ac-738 curacy of deep equivariant models to biomolecular simulations of realistic size. In Proceedings of 739 the International Conference for High Performance Computing, Networking, Storage and Analy-740 sis, pp. 1–12, 2023. 741 742 Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, 743 Alexander Pritzel, Suman Ravuri, Timo Ewalds, Ferran Alet, Zach Eaton-Rosen, et al. Graphcast: 744 Learning skillful medium-range global weather forecasting. arXiv preprint arXiv:2212.12794, 2022. 745 746 Greg Landrum. Rdkit: Open-source cheminformatics software. 2016. 747 748 Lin Li, Chuan Li, and Emil Alexov. On the modeling of polar component of solvation energy using smooth gaussian-based dielectric function. Journal of Theoretical and Computational Chemistry, 749 13(03):1440002, 2014. 750 751 Shaoning Li, Yusong Wang, Mingyu Li, Jian Zhang, Bin Shao, Nanning Zheng, and Jian Tang. 752 F^{3} low: Frame-to-frame coarse-grained molecular dynamics with se (3) guided flow matching. 753 arXiv preprint arXiv:2405.00751, 2024. 754 Zijie Li, Kazem Meidani, and Amir Barati Farimani. Transformer for partial differential equations' 755 operator learning. Transactions on Machine Learning Research, 2023a. ISSN 2835-8856.

- Zijie Li, Dule Shu, and Amir Barati Farimani. Scalable transformer for pde surrogate modeling. *arXiv preprint arXiv:2305.17560*, 2023b.
 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, An-
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020a.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*, 2020b.
- Mario Lino, Stathi Fotiadis, Anil A Bharath, and Chris D Cantwell. Multi-scale rotation-equivariant graph neural networks for unsteady eulerian fluid dynamics. *Physics of Fluids*, 34(8), 2022.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net, 2017.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019.
 OpenReview.net, 2019.
- Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
- Leon B Lucy. A numerical approach to the testing of the fission hypothesis. *Astronomical Journal*, vol. 82, Dec. 1977, p. 1013-1024., 82:1013–1024, 1977.
- Salvatore Marrone, Matteo Antuono, A Colagrossi, G Colicchio, D Le Touzé, and G Graziani. δ sph model for simulating violent impact flows. *Computer Methods in Applied Mechanics and Engineering*, 200(13-16):1526–1542, 2011.

- Andreas Mayr, Sebastian Lehner, Arno Mayrhofer, Christoph Kloss, Sepp Hochreiter, and Johannes
 Brandstetter. Boundary graph neural networks for 3d simulations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 9099–9107, 2023.
- Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Oc cupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4460–4470, 2019.
- Laurence Midgley, Vincent Stimper, Javier Antorán, Emile Mathieu, Bernhard Schölkopf, and
 José Miguel Hernández-Lobato. Se (3) equivariant augmented coupling flows. *Advances in Neu- ral Information Processing Systems*, 36, 2024.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- Joe J Monaghan. Smoothed particle hydrodynamics. *Reports on progress in physics*, 68(8):1703, 2005.
- Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J Owen, Mordechai Ko rnbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications*, 14(1):579, 2023.

810	Frank Noé Simon Olsson, Jonas Köhler, and Hao Wu, Boltzmann generators: Sampling equilibrium
811	Frank too, one how stone to have a how of the botterial generations. Sumpling equilibrium
0.1.0	states of many-body systems with deep learning. Science, 365(6457):eaaw1147, 2019.
812	
813	Noel M O'Boyle, Michael Banck, Craig A James, Chris Morley, Tim Vandermeersch, and Geof-
814	frey R Hutchison. Open babel: An open chemical toolbox. Journal of cheminformatics, 3(1):

1-14, 2011.
Cobriele Orlando, Daniele Paimondi, Pamon Duran Po

819

831

843

846

847

848

849

- Gabriele Orlando, Daniele Raimondi, Ramon Duran-Romaña, Yves Moreau, Joost Schymkowitz, and Frederic Rousseau. Pyuul provides an interface between biological structures and deep learning algorithms. *Nature Communications*, 13, 02 2022. doi: 10.1038/s41467-022-28327-3.
- Johannes Pahlke and Ivo F. Sbalzarini. A unifying mathematical definition of particle methods. *IEEE Open Journal of the Computer Society*, 4:97–108, 2023. doi: 10.1109/OJCS.2023.3254466.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. pp. 4195–4205, 2023.
- Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. Film: Visual reasoning with a general conditioning layer. pp. 3942–3951. AAAI Press, 2018.
- Pedro O Pinheiro, Arian Jamasb, Omar Mahmood, Vishnu Sresht, and Saeed Saremi. Structure based drug design by denoising voxel grids. *arXiv preprint arXiv:2405.03961*, 2024.
- Jay W. Ponder and David A. Case. Force fields for protein simulations. *Advances in protein chemistry*, 66:27–85, 2003. ISSN 0065-3233. doi: 10.1016/S0065-3233(03)66002-X.
- Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical fea ture learning on point sets in a metric space. Advances in neural information processing systems, 30, 2017.
- ⁸³⁵ G N Ramachandran, C Ramakrishnan, and V Sasisekharan. Stereochemistry of polypeptide chain configurations. *Journal of Molecular Biology*, pp. 95–99, 1963.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10674–10685. IEEE, 2022.
- Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter
 Battaglia. Learning to simulate complex physics with graph networks. In *International conference* on machine learning, pp. 8459–8468. PMLR, 2020.
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini.
 The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
 - Mathias Schreiner, Ole Winther, and Simon Olsson. Implicit transfer operator learning: multiple time-resolution surrogates for molecular dynamics. *arXiv preprint arXiv:2305.18046*, 2023.
 - E. Schrödinger. An undulatory theory of the mechanics of atoms and molecules. *Physical Review*, 28(6):1049–1070, 1926. ISSN 0031-899X. doi: 10.1103/PhysRev.28.1049.
- Jacob Seidman, Georgios Kissas, Paris Perdikaris, and George J Pappas. Nomad: Nonlinear man ifold decoders for operator learning. *Advances in Neural Information Processing Systems*, 35:
 5601–5613, 2022.
- Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020.
- Julian Suk, Christoph Brune, and Jelmer M Wolterink. Se (3) symmetry lets graph neural networks
 learn arterial velocity estimation from small datasets. In *International Conference on Functional Imaging and Modeling of the Heart*, pp. 445–454. Springer, 2023.
- Deborah Sulsky, Shi-Jian Zhou, and Howard L Schreyer. Application of a particle-in-cell method to solid mechanics. *Computer physics communications*, 87(1-2):236–252, 1995.
- 863 Roger Temam. *Navier-Stokes equations: theory and numerical analysis*, volume 343. American Mathematical Soc., 2001.

874

875

879

885

887

891

893

900

905

906

864	Nils Thuerey, Philipp Holl, Maximilian Mueller, Patrick Schnell, Felix Trost, and Kiwon Um.
865	Physics-based deep learning. arXiv preprint arXiv:2109.05237, 2021.
866	

- Artur P Toshev, Gianluca Galletti, Johannes Brandstetter, Stefan Adami, and Nikolaus A Adams. 867 Learning lagrangian fluid mechanics with e (3)-equivariant graph neural networks. arXiv preprint 868 arXiv:2305.15603, 2023a. 869
- 870 Artur P. Toshev, Gianluca Galletti, Fabian Fritz, Stefan Adami, and Nikolaus A. Adams. La-871 grangebench: A lagrangian fluid mechanics benchmarking suite. In 37th Conference on Neural 872 Information Processing Systems (NeurIPS 2023) Track on Datasets and Benchmarks, 2023b.
 - Artur P Toshev, Jonas A Erbesdobler, Nikolaus A Adams, and Johannes Brandstetter. Neural sph: Improved neural modeling of lagrangian fluid dynamics. arXiv preprint arXiv:2402.06275, 2024.
- 876 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, 877 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. Advances in neural informa-878 tion processing systems, 30, 2017.
- Ricardo Vinuesa and Steven L Brunton. Enhancing computational fluid dynamics with machine 880 learning. Nature Computational Science, 2(6):358–366, 2022.
- 882 Damien Violeau and Benedict D Rogers. Smoothed particle hydrodynamics (sph) for free-surface 883 flows: past, present and future. Journal of Hydraulic Research, 54(1):1–26, 2016.
 - Sifan Wang, Jacob H Seidman, Shyam Sankaran, Hanwen Wang, George J Pappas, and Paris Perdikaris. Bridging operator learning and conditioned neural fields: A unifying perspective. arXiv preprint arXiv:2405.13998, 2024.
- 888 Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast transformer solver for pdes on general geometries. arXiv preprint arXiv:2402.02366, 2024. 889
- 890 Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual comput-892 ing and beyond. In Computer Graphics Forum, volume 41, pp. 641-676. Wiley Online Library, 2022. 894
- Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geo-895 metric diffusion model for molecular conformation generation. arXiv preprint arXiv:2203.02923, 896 2022. 897
- Ziyang Yu, Wenbing Huang, and Yang Liu. Force-guided bridge matching for full-atom time-899 coarsened dynamics of peptides. arXiv preprint arXiv:2408.15126, 2024.
- Biao Zhang, Jiapeng Tang, Matthias Nießner, and Peter Wonka. 3dshape2vecset: A 3d shape rep-901 resentation for neural fields and generative diffusion models. ACM Transactions on Graphics, 42 902 (4):1–16, July 2023. ISSN 1557-7368. doi: 10.1145/3592442. 903
- 904 Chi Zhang, XY Hu, and Nikolaus A Adams. A weakly compressible sph method based on a lowdissipation riemann solver. Journal of Computational Physics, 335:605–620, 2017.
- Chi Zhang, Massoud Rezavand, and Xiangyu Hu. Dual-criteria time stepping for weakly compressible smoothed particle hydrodynamics. Journal of Computational Physics, 404:109135, 2020. 908
- 909 Chi Zhang, Massoud Rezavand, Yujie Zhu, Yongchuan Yu, Dong Wu, Wenbin Zhang, Jianhang 910 Wang, and Xiangyu Hu. Sphinxsys: An open-source multi-physics and multi-resolution library 911 based on smoothed particle hydrodynamics. Computer Physics Communications, 267:108066, 912 2021.
- 913 Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky TQ Chen. Guided 914 flows for generative modeling and decision making. arXiv preprint arXiv:2311.13443, 2023. 915
- 916 Robert Zwanzig. Nonlinear generalized langevin equations. Journal of Statistical Physics, 9(3): 917 215–220, 1973.

918	CONTENTS
919	CONTENTS

920	А	Lag	rangian fluid simulation	19
921		Δ 1	Smoothed Particle Hydrodynamics (SPH)	10
923		A.1		1)
924		A.2	Scaling limits	19
925		A.3	Implementation details	20
926		A.4	Dataset-specific details	22
927 928		A.5	Baselines	22
929		A.6	Additional results	23
930		A 7	Dem brook 2D dataset	22
931		A./		23
932	р	Mal	anlar Conformation Someling	20
933	Б	NIOI	ecular Conformation Sampling	20
934		B .1	Molecular Dynamics (MD)	26
935 936		B .2	Implementation details	28
937			B.2.1 Density representation	28
938			B 2.2 Importance-based sampling for molecules	29
939				
940			B.2.3 Refinement	29
941			B.2.4 Data augmentation	29
942		D 3	Neural Sampling Operator	30
943		D .5		50
944		B .4	Molecular Graph Reconstruction	32
945		B.5	Additional Results	33
946				
947				

972 А LAGRANGIAN FLUID SIMULATION 973

A.1 **SMOOTHED PARTICLE HYDRODYNAMICS (SPH)**

d

976 In contrast to Eulerian approaches, where discretization of the continuous space is achieved through 977 spatially fixed finite nodes, control volumes, cells, or elements, Lagrangian methods employ finite 978 material points, often termed *particles*, whose movement aligns with the local deformation of the 979 continuum. One of the most prominent Lagrangian discretization schemes is smoothed particle 980 hydrodynamics (SPH), originally proposed by Lucy (1977) and Gingold & Monaghan (1977) for 981 applications in astrophysics. SPH approximates the field properties using radial kernel interpolations 982 over adjacent particles at the location of each particle. The strength of the SPH method is that it does not require connectivity constraints, e.g., meshes, which is particularly useful for simulating systems 983 with large deformations. Since its foundation, SPH has been greatly extended and is the preferred 984 method to simulate problems with (a) free surfaces (Marrone et al., 2011; Violeau & Rogers, 2016), 985 (b) complex boundaries (Adami et al., 2012), (c) multi-phase flows (Hu & Adams, 2007), and (d) 986 fluid-structure interactions (Antoci et al., 2007). SPH approximates the incompressible Navier-987 Stokes equations (NSE) by the so-called weakly compressible NSE, where the weak compressibility 988 assumption typically allows for up to $\sim 1\%$ density deviation Monaghan (2005). This $\sim 1\%$ is 989 enforced for the weakly compressible SPH method while evolving density and momentum: 990

$$\frac{\mathrm{d}}{\mathrm{d}t}(\rho) = -\rho \left(\nabla \cdot \boldsymbol{u}\right),\tag{A.1}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}(\boldsymbol{u}) = \underbrace{-\frac{1}{\rho}\nabla p}_{\text{pressure}} + \underbrace{\frac{1}{Re}\nabla^2 \boldsymbol{u}}_{\text{viscosity}} + \underbrace{\boldsymbol{f}}_{\text{ext. force}}.$$
(A.2)

997 Herein, ρ is the density, **u** the velocity, p the pressure, f an external force, $Re \propto 1/\mu$ the Reynolds 998 number. Solving these equations with standard SPH methods may still produce artifacts, most no-999 tably when particle clumping exceeds the 1% density-fluctuation restriction (Adami et al., 2013; 1000 Toshev et al., 2024).

1001 The Material Point Method (MPM) is another particle-based technique that represents material as an 1002 assembly of material points. The motion of each material point is determined by solving Newton's 1003 laws of motion. MPM adopts a hybrid Eulerian-Lagrangian scheme, which uses moving material 1004 points and a fixed computational grid. MPM is particularly useful in the context of large deforma-1005 tions including fracture and contact scenarios, where traditional mesh-based methods might yield 1006 unrealistic or undesired outcomes due to mesh distortions.

1007 1008

974

975

991

992

993 994

995 996

A.2 SCALING LIMITS 1009

1010 In Figure 3, we compare the memory consumption between UPT++, GNS-10-128, and GNS-5-64 1011 during training. We construct a toy setting based on DamBreak3D, positioning points on a regular 1012 three-dimensional grid with the same particle spacing Δx used in DamBreak3D (see Section A.7 1013 for details). We start with a grid of 16x16x16 points and double the last dimension repeatedly, 1014 increasing the size from 16x16x16 to 16x16x8192. This results in configurations ranging from 4k 1015 points to over 2 million points. For UPT++, we scale the number of input points selected, the number 1016 of supernodes $n_{\rm S}$, and the number of points where we decode the velocity and the occupancy based on the number of particles. We select 25% of the points as input points, use 2.5% as supernodes, 1017 decode the velocity at 10% and decode the occupancy at 20%, which is similar to our setting used 1018 for DamBreak3D, compare Section A1). We fix the number of latent tokens to $n_{latent} = 4096$. We 1019 focus on memory consumption during the first training stage, where the encoder and decoder are 1020 trained, as the second stage requires less memory. 1021

1022 The main memory consumption that can be attributed to UPT++ is the query MLP in the decoder, 1023 which scales linearly with the number of points to decode, both for decoding the occupancy and the velocity. We can further reduce the memory footprint by decoding a smaller fraction of the points, 1024 as can be seen in Figure A1, where we add two scaling plots where we decode only 5% and only 1025 1% of the points.

40 ---- UPT++ UPT++, 5% decode 16 UPT++, 1% decode 8 Memory [GB] 4 2 1 0.5 0.25 52AX ~3¹* 2624 20484 20974 A1944 3²⁴ హి 0× ~64 Number of particles

Figure A1: Memory usage of UPT++ variants with a reduced number of decoding points.

1040 1041

1039

1026

1027

1028

1029

1030

1031

1032

1033 1034

1035

1036

A.3 IMPLEMENTATION DETAILS

The following outlines the implementation details for UPT++. Table A1 gives a detailed overview of the hyperparameters used in training. If there is a change in the dimensions between different blocks, we perform a learnable linear projection. All transformer (Vaswani et al., 2017) and perceiver (Jaegle et al., 2021b) blocks use standard pre-norm architecture as used in ViT Dosovitskiy et al. (2020) and are modulated by the timestep using DiT modulation (Peebles & Xie, 2023). We use layer normalization Ba et al. (2016) in the output of the encoder as well as in the input and output of the forward operator to always keep the latent representation normalized.

All experiments use the AdamW optimizer (Kingma & Ba, 2015; Loshchilov & Hutter, 2019) and follow a learning rate schedule that begins with a linear warmup and transitions to cosine decay (Loshchilov & Hutter, 2017). We perform early stopping and use the best checkpoint in terms of *IoU* evaluated on the validation set of each dataset respectively.

1055 Encoder. First we sample a subset out of all particles and project two consecutive velocities 1056 into a higher dimension using a linear layer. Then we sample $n_{\rm S}$ supernodes from the input point 1057 cloud and perform message passing to points within a radius $r_{\rm S}$ to process the local information. In 1058 message passing, the features of the supernode and the point are concatenated along with a positional 1059 embedding that captures their relative distance. The supernode features are then processed by a transformer to capture global information. To further reduce the number of tokens in the latent space, we apply perceiver pooling with learned queries, followed by layer normalization, producing 1061 a normalized latent representation. 1062

Latent space operator. We apply layer normalization to the latent representation before passing it into the transformer blocks. The output is then added to the original latent representation (after normalization) and passed through another layer normalization step, ensuring the latent representation remains consistently normalized throughout the process. The timestep conditioning of the transformer blocks uses the timestep of the latent representation in the input.

Decoder. The decoder takes the latent representation and processes it with a small Transformer, which is then fed into two perceiver blocks, one for decoding the state and one for decoding the occupancy. The positions where we want to query the latent representation are transformed into a positional embedding and fed through an MLP, resulting in the query used by the perceiver. The keys and values are the outputs of the Transformer, and the result of the perceiver is projected into the input dimension of the physical state or into a two-dimensional output for the occupancy.

Timestep modulation. To incorporate information from the timestep, we encode the current timestep into a positional embedding. We use the transformer positional encoding Vaswani et al. (2017); Gupta & Brandstetter (2022), and, therefore rescale all timesteps to the range [0, 200]. The resulting timestep embedding is then used to perform DiT modulation (Peebles & Xie, 2023) for all transformer and perceiver blocks. DiT modulation involves applying dimension-wise scaling, shifting, and gating operations to both the attention and MLP modules of the transformer and perceiver blocks.

Hyperparameter	WaterDrop	WaterDrop-XL	DamBreak3I
General model parameters			
Number of latent tokens n_{latent}	128	128	512
Timestep embedding dim	192	192	192
DiT conditioning dim	768	768	768
Encoder			
Range of input points selected k	400 - 800	1k-3k	32k-64k
Input features	4	4	6
Node features	96	96	96
Num. supernodes $n_{\rm S}$	128	512	4096
Supernode radius $r_{\rm S}$	0.05	0.05	0.15
Max supernode neighbours	8	8	32
Relative positional embedding dim	96	96	96
Message passing MLP dims	288/96	288/96	288/96
Transformer dim / layers / heads	96/4/2	96/4/2	96/4/2
Perceiver dim / num heads	192/3	192/3	192/3
Forward operator			
Transformer dim / layers / heads	192/12/3	192/12/3	192/12/3
Decoder			
Transformer dim / layers / heads	192/4/3	192/4/3	192/4/3
Query MLP dims	768/768/192	768/768/192	768/768/19
Perceiver dim / num heads	192/3	192/3	192/3
Output features	4	4	6
Number of points to decode (velocity) k'	125	500	16k
Number of points to decode (occupancy) k'_{a}	250	1000	32k
Occupancy radius of a particle	0.01	0.01	0.05
First training stage			
Num. epochs	10	10	10
Learning rate	5e-3	5e-4	5e-4
Weight decay rate	0.05	0.05	0.05
Warmup epochs	2	2	2
Batch size	1024	256	32
Second training stage			
Num. epochs	10	10	10
Learning rate	5e-4	5e-4	5e-4
Weight decay rate	0.05	0.05	0.05
Warmup epochs	2	2	2
······································	-	-	-

Table A1: UPT++	hyperparameters	for th	e application	to Nav	ier-Stokes	equations.

1080

1122

1123

1124

1125

1126 Training. In the first training stage, we train both the encoder and decoder without the time-1127 evolution operator. We do this by sampling states from a trajectory at timestep t, selecting k input 1128 points, decoding at k' output points, and regressing the velocity at these points using an MSE objective. Additionally, we randomly sample k'_{o} points within the domain to obtain the ground truth 1129 occupancy. If a coordinate x lies outside the occupancy radius of all particles, it is labeled as un-1130 occupied; otherwise, it is marked as occupied by the fluid. We train using a CE loss. In the second 1131 training stage, we freeze the encoder and encode two consecutive timesteps, t and t', into latent 1132 space representations z^t and $z^{t'}$. The time-evolution operator uses z^t as input to predict $z^{t'}$, and 1133 we train the time-evolution operator using an MSE objective.

1134 A.4 DATASET-SPECIFIC DETAILS

In the following Table A2, we summarize the size of the used datasets, emphasizing that we work with a median number of particles all the way from 500 through 4,000 to 180,000. Each dataset consists of 1000 trajectories in the training set and 100 trajectories in both the validation and test sets.

1141Table A2: Statistics of particle counts and trajectory length in our Lagrangian fluid dynamics1142datasets.

Dataset	nur	ber of part	Trajectory length	
Duniser	min	median	max	ingereig imgen
WaterDrop	195	548	1,108	1000
WaterDrop-XL	1,948	4,031	7,184	1000
DamBreak3D	144,133	179,312	215,661	250

1149 1150 1151

1153

1140

1152 A.5 BASELINES

Our baselines on the Lagrangian fluid dynamics problems are GNS (Sanchez-Gonzalez et al., 2020) and its multi-scale version MS-GNS. We adopt the Pytorch implementation of GNS from github.com/wu375/simple-physics-simulator-pytorch-geometry to our codebase and reuse most building block in MS-GNS. In the following, we provide more details about MS-GNS and the training hyperparameters.

MS-GNS. To the best of our knowledge, our baseline model MS-GNS is the first multi-scale GNN 1159 for Lagrangian fluid dynamics, and it combines ideas from Fortunato et al. (2022) and our own 1160 importance-based encoder approach. We acknowledge that there are various multi-scale GNNs that 1161 operate on static objects like point sets (Qi et al., 2017; Lino et al., 2022) or meshes (Fortunato et al., 1162 2022; Suk et al., 2023), but because these discretizations are static, one can precompute coarser ver-1163 sions thereof using various different algorithms. However, in Lagrangian numerical methods, we 1164 need to compute a coarse graph at every timestep of the autoregressive rollout evolution, making 1165 advanced algorithms like furthest point sampling (FPS) (Qi et al., 2017) unfeasible - see tested FPS 1166 on DamBreak3D –, taking 10x more time than the model forward evaluation. Thus, for constructing 1167 the coarser graph, we resort to what we do in the sampling-based UPT++ encoding, namely ran-1168 domly picking a subset of the nodes. The only difference to the encoder is that we do not have a fixed number of supernodes, but rather a relative ratio of subsampled nodes, which we set to 0.5^{dim} , 1169 which essentially means that we go to particles with 2x the radius of the finer particles, which also 1170 means that we can just double the cutoff radius for the coarser graph. The obvious disadvantage of 1171 this method is that some fine particles might be far away from the coarser particles, which we rem-1172 edy by constructing the mapping from finer to coarser graph using k-nearest neighbors with k = 41173 -4 basically means that a fine particle sees either its corresponding coarse particle and 3 others, or 1174 just 4 coarse particles. This way every fine node is guaranteed to get access to information from the 1175 coarser nodes, and because we subsample the coarse nodes anew at every timestep, the information 1176 propagation is well distributed. Other than the coarse graph generation and the fine-coarse mapping 1177 graph, our approach is almost equivalent to MS-MGN by Fortunato et al. (2022), with the only dif-1178 ference being that all latent vectors connected with the original finely resolved graph have half the 1179 size of the latent vectors of the coarse graph; this adjustment significantly reduces the memory of the forward pass. Overall, the message-passing steps that operate only on the fine or coarse graphs 1180 are exactly the GNS layers, and the mapping between the resolutions happens along the same k-NN 1181 graph. 1182

The hyperparameter setting for MS-GNS is one of what Fortunato et al. (2022) found to work well, namely a simple V-shaped processor scheme which we call MS-GNS-15 consisting of: 1 MP layer on the fine scale, downsampling layer, 11 MP layers on the coarse graph, upsampling layer, and 1 more MP layer on the fine graph. Regarding the training protocol, we train GNS and MS-GNS with the same optimizer and learning rate scheduler as our other models. A summary of the hyperparameters used for training the GNN baselines is given in Table A3, which complements Appendix A.3.

1189			
1190	Hyperparameter	GNS-10-128/GNS-5-64	MS-GNS-15
1191	Physical input/output features		
1192	Num, input velocities	5	5
1193	Node input features	velocity, boundary dist.	velocity, boundary dist.
1194	Edge input features	displacement	displacement
1195	Node output features	acceleration	acceleration
1196	Include magnitudes	yes	yes
1197	GNN architecture		
1198	MP layers (if appl. fine)	10/5	2
1199	Upsampling layers	_	1
1200	Downsampling layers	-	1
1201	MP layers coarse	-	11
1202	Latent dimension (if appl. fine)	128/64	64
1203	Latent dimension coarse	-	128
1204	Num. MLP layers	2	2
1205	Noise std	6.7e-4	6.7e-4
1206	Training configuration		
1207	Num. epochs	10	10
1208	Learning rate	1e-4	1e-4
1209	Weight decay rate	0.05	0.05
1210	Warmup epochs	2	2
1211	Batch size	{WaterDrop: 2, WaterDrop	p-XL: 10, DamBreak3D: 4}

Table A3: GNN hyperparameters overview.

1212 1213 1214

1188

A.6 ADDITIONAL RESULTS

1215 In Figure A2 we present the IoU and the velocity error for full rollouts on the test set. At the 1216 start of the trajectory, GNS has an advantage by numerically integrating the accelerations. However, 1217 during the highly dynamic phase between timesteps 200 and 500, the difference becomes negligible. 1218 Figure A3 and A4 illustrate this by presenting snapshots of a trajectory rollout for both Waterdrop 1219 and Waterdrop-XL. Similarly, as demonstrated for DamBreak3D in Figure 1, UPT++ captures the 1220 overall fluid dynamics in both the smaller Waterdrop and Waterdrop-XL scenarios. In the end of 1221 each trajectory, as the fluid settles at the bottom of the box, GNS inherently preserves mass, or the 1222 number of particles, which enables it to more accurately capture the volume. In contrast, UPT++ latent propagation lacks such constraints, making it unable to accurately capture the fluid's volume. 1223

1224 1225

A.7 DAM BREAK 3D DATASET 1226

We generated the dataset by modifying the 3D dam break test case in the SPHinXsys library (Zhang 1227 et al., 2021)¹. In particular, we modify a) the numerical integration scheme and b) the initial geom-1228 etry of the fluid. 1229

1230 Regarding the integrator, we modify the adaptive-step dual-criteria time stepping scheme (Zhang 1231 et al., 2020) by fixing the step size $\Delta t_{inner} = 0.0008$ of the inner pressure and density relaxation 1232 loop and also by fixing the number of iterations in this inner loop to 5. The reason for this is that we 1233 want equidistantly spaced samples in time to not have to deal with conditioning on the timestep size in the ML problem formulation. The chosen Δt_{inner} is close to the worst-case adaptively estimated 1234 one but still does not significantly change the number of integration steps to reach the end time of 20. 1235 Note that we omit units here as the simulation is of the non-dimensionalized NSE. With the temporal 1236 coarsening level of 100 relative to the inner loop steps, each simulation has $20/(0.0008 \cdot 100) = 250$ 1237 steps (to be more precise, 251 steps, as we record both the very first and last states). 1238

1239 Regarding the parametrization of the geometry, we sample 12 random numbers determining the shape of the top and front of the wave, and after the fluid volume is filled with particles, we add 1240

¹²⁴¹

¹https://github.com/Xiangyu-Hu/SPHinXsys

1268 1269

1270

1272 1273

1274

1276 1277

1278

1283

1293 1294 1295



Figure A2: Mean and standard deviation of the IoU and the velocity error during the rollout of all trajectories in the test set. For WaterDrop and WaterDrop-XL, the first simulation steps are better predicted by GNS, which is expected because GNS numerically integrates accelerations, and also, the last steps, where the fluid is resting at the bottom, are better predicted. During the highly dynamic part, GNS and UPT++ are on par, while UPT++ better predicts the correct velocity. For DamBreak3D, GNS is not able to predict the rollout of the large-scale trajectory, but UPT++ and MS-GNS can handle this task.



Figure A3: Various timesteps along the WaterDrop trajectory are evaluated on a regular grid for comparison purposes. The presence of a point on the grid represents its occupancy, while its color indicates the magnitude of the velocity.

1284 Gaussian noise to the coordinates. The standard deviation of the noise is $\sigma = 0.1 \cdot \Delta x$ with $\Delta x =$ 0.025 being the particle spacing. The computational domain begins at (0,0,0) and spans $L \times H \times I$ 1285 $W = 5.366 \times 2 \times 2$, with the number 5.366 coming from the original dam break experiments 1286 by Colagrossi & Maurizio (2003). The fluid always fills the bottom left part of the domain (at 1287 x = 0, y = 0) spanning the full width, and we modulate the top and front sides by the mentioned 1288 12 numbers defining sinusoidal waves by their amplitude a, period p, and shift s. The top surface of 1289 the fluid is defined by its height $h_{top}(x, z)$ as a function of the length and width (x and z axes), and 1290 the x-coordinate (length) of the front $l_{front}(z, y)$ is defined as a function of the width and height. 1291

$$h_{top}(x,z) = H_{ave} + a_{top,x} \cdot \sin\left(2\pi(p_{top,x} \cdot x/L_{ave} + s_{top,x})\right) \\ + a_{top,z} \cdot \sin\left(2\pi(p_{top,z} \cdot z/W_{ave} + s_{top,z})\right) \\ l_{front}(z,y) = L_{ave} + a_{front,z} \cdot \sin\left(2\pi(p_{front,z} \cdot z/W_{ave} + s_{front,z})\right)$$

 Timestep: 10
 Timestep: 210
 Timestep: 310
 Timestep: 410

 SUB
 Sub

Figure A4: Various timesteps along the WaterDrop-XL trajectory are evaluated on a regular grid for comparison purposes. The presence of a point on the grid represents its occupancy, while its color indicates the magnitude of the velocity.

1314 The values of the average length, height, and width of the fluid are $L_{ave} = 2$, $H_{ave} = 0.7$, 1315 and $W_{ave} = 2$, respectively. The random numbers for a, p, s are sampled uniformly from 1316 $a \sim \mathcal{U}(0, 0.15), p \sim \mathcal{U}(0.25, 2), s \sim \mathcal{U}(0, 1)$. We visualize the first 10 trajectories from the 1317 train split in Figure A5.



Figure A5: Frames 1, 100, and 250 from the first 10 training trajectories of DamBreak3D. The color is the velocity magnitude estimated by subtracting the previous positions from the current ones.

1350 B MOLECULAR CONFORMATION SAMPLING

1352 B.1 MOLECULAR DYNAMICS (MD)

1354 The most fundamental concepts nowadays to describe the dynamics of molecules are given by the 1355 laws of quantum mechanics. The Schrödinger equation is a partial differential equation, that gives the evolution of the complex-valued wave function ψ over time t: $i\hbar \frac{\partial \psi}{\partial t} = \hat{H}(t)\psi$. Here i is 1356 1357 the imaginary unit with $i^2 = -1$, \hbar is reduced Planck constant, and, $\hat{H}(t)$ is the Hamiltonian 1358 operator at time t, which is applied to a function ψ and maps to another function. It determines how 1359 a quantum system evolves with time and its eigenvalues correspond to measurable energy values 1360 of the quantum system. The solution to Schrödinger's equation in the many-body case (particles 1361 1,...,N) is the wave function $\psi(\mathbf{x}_1,...,\mathbf{x}_N,t) : \times_{i=1}^N \mathbb{R}^3 \times \mathbb{R} \to \mathbb{C}$ which we abbreviate as $\psi(\{\mathbf{x}\},t)$. It's the square modulus $|\psi(\{\mathbf{x}\},t)|^2 = \psi^*(\{\mathbf{x}\},t)\psi(\{\mathbf{x}\},t)$ is usually interpreted as 1362 1363 a probability density to measure the positions x_1, \ldots, x_N at time t, whereby the normalization 1364 condition $\int \dots \int |\psi(\{\mathbf{x}\}, t)|^2 d\mathbf{x}_1 \dots d\mathbf{x}_N = 1$ holds for the wave function ψ . 1365

Analytic solutions of ψ for specific operators H(t) are hardly known and are only available for sim-1367 ple systems like free particles or hydrogen atoms. In contrast to that are proteins with many thou-1368 sands of atoms. However, already for much smaller quantum systems approximations are needed. 1369 A famous example is the Born–Oppenheimer approximation, where the wave function of the multi-1370 body system is decomposed into parts for heavier atom nuclei and the light-weight electrons, which 1371 usually move much faster. In this case, one obtains a Schrödinger equation for electron movement and another Schrödinger equation for nuclei movement. A much faster option than solving a sec-1372 ond Schrödinger equation for the motion of the nuclei is to use the laws from classical Newtonian 1373 dynamics. The solution of the first Schrödinger equation defines an energy potential, which can 1374 be utilized to obtain forces \mathbf{F}_i on the nuclei and to update nuclei positions according to Newton's 1375 equation of motion: $\mathbf{F}_i = m_i \, \ddot{\mathbf{q}}_i(t)$ (with m_i being the mass of particle i and $\mathbf{q}_i(t)$ describing the 1376 motion trajectory of particle i over time t). 1377

Additional complexity in studying molecule dynamics is introduced by environmental conditions 1378 surrounding molecules. Maybe the most important is temperature. For bio-molecules it is often 1379 of interest to assume that they are dissolved in water. To model temperature, a usual strategy is 1380 to assume a system of coupled harmonic oscillators to model a heat bath, from which Langevin 1381 dynamics can be derived (Ford et al., 1965; Zwanzig, 1973). The investigation of the relationship 1382 between quantum-mechanical modeling of heat baths and Langevin dynamics still seems to be a current research topic, where there there are different aspects like the coupling of the oscillators or 1384 Markovian properties when stochastic forces are introduced. For instance, Hoel & Szepessy (2019), 1385 studies how canonical quantum observables are approximated by molecular dynamics. This includes 1386 the definition of density operators, which behave according to the quantum Liouville-von Neumann 1387 equation.

The forces in molecules are usually given as the negative derivative of the (potential) energy: $\mathbf{F}_i = -\nabla E$. In the context of molecules, E is usually assumed to be defined by a force field, which is a parameterized sum of intra- and intermolecular interaction terms. An example is the Amber force field (Ponder & Case, 2003; Case et al., 2024):

1394

$$\begin{split} E &= \sum_{\text{bonds } r} k_b (r - r_0)^2 + \sum_{\text{angles } \theta} k_\theta (\theta - \theta_0)^2 + \\ &\sum_{\text{dihedrals } \phi} V_n (1 + \cos(n\phi - \gamma)) + \sum_{i=1}^{N-1} \sum_{j=i+1}^N \left(\frac{A_{ij}}{R_{ij}^{12}} - \frac{B_{ij}}{R_{ij}^6} + \frac{q_i q_j}{\epsilon R_{ij}} \right) \end{split}$$

(B.1)

Here $k_b, r_0, k_\theta, \theta_0, V_n, \gamma, A_{ij}, B_{ij}, \epsilon, q_i, q_j$ serve as force field parameters, which are found either empirically or which might be inspired by theory.

Newton's equations of motions for all particles under consideration form a system of ordinary differential equations (ODEs), to which different numeric integration schemes like Euler, Leapfrog, or, Verlet can be applied to obtain particle position trajectories for given initial positions and initial velocities. In case temperature is included, the resulting Langevin equations form a system of

1404	state the differential equations (CDCs) and I experimentary conclusions for the second Hasheveld be seen
1405	successic differential equations (SDEs), and Langevin integrators can be used. It should be men-
1406	tioned, that it is often necessary to use very small integration timesteps to avoid large approximation
1407	errors. This, nowever, increases the time needed to find new stable molecular configurations.
1/08	
1/100	
1405	
1410	
1411	
1412	
1413	
1414	
1415	
1416	
1417	
1418	
1419	
1420	
1421	
1422	
1423	
1424	
1425	
1426	
1427	
1428	
1429	
1430	
1431	
1432	
1433	
1434	
1435	
1436	
1437	
1438	
1439	
1440	
1441	
1442	
1443	
1444	
1445	
1446	
1447	
1448	
1449	
1450	
1451	
1452	
1453	
1454	
1455	
1456	
1457	

1458 B.2 IMPLEMENTATION DETAILS

We use the same implementation as outlined in Chapter A.3. The differences specific to the MD setting are explained below. Table B1 summarizes the hyperparameters used in the molecular sampling
experiments.

Table B1: UPT++ hyperparameters for the application to molecular sampling.

1465			
1466	Hyperparameter	AD	2AA
1467	General Model Parameters		
1468	Number of latent tokens n_{latent}	32	64
1469	Timestep embedding dim	192	192
1470	DiT conditioning dim	768	768
1471	Encoder		
1472	Range of input points selected	2.7k	6k
1473	Input features	4	4
1474	Node features	96	96
1475	Num. supernodes $n_{\rm S}$	128	512
1476	Supernode radius $r_{\rm S}$	0.05	0.05
1477	Max supernode neighbours	8	8
1478	Relative positional embedding dim	96	96
1479	Message passing MLP dims	288/96	288/96
1480	Transformer dim / layers / heads	96/4/2	96/4/2
1481	Perceiver dim / num heads	192/3	192/3
1482	Forward operator		
1483	Transformer dim / layers / heads	32/22/3	32/22/3
1484	Decoder		
1485	Transformer dim / layers / heads	192/4/3	192/4/3
1486	Query MLP dims	768/768/192	768/768/192
1487	Perceiver dim / num heads	192/3	192/3
1488	Output features	5	5
1489	Number of points to decode (velocity)	125	500
1490	Number of points to decode (occupancy)	250	1000
1491	First stage training		
1492	Num. epochs	1.7k	73
1493	Learning rate	1e-4	1e-4
1494	Schedule	Cosine	-
1495	Batch size	1024	1024
1496	Second stage training		
1497	Num. epochs	4.6k	53
1498	Learning rate	1e-4	1e-4
1499	Batch size	2048	256

1500 1501

1502

1506

1507 1508

1463

1464

B.2.1 DENSITY REPRESENTATION

We represent molecules as density fields, following an approach similar to Pinheiro et al. (2024) and
Dumitrescu et al. (2024). Each atom is represented by a 3D Gaussian-like density (Orlando et al., 2022; Li et al., 2014)

$$D(d,r) = \exp\left(-\frac{d^2}{(0.93 \cdot r)^2}\right),$$
(B.2)

where D is the fraction of occupied volume by an atom with radius r at distance d from its center. While different occupancy radii could be considered for various atom types, we use a uniform radius of r = 0.5 Å for all atom types. The signal of the field for atom type a, with I_a being an index array of all atoms within the molecule corresponding to atom type a, is defined as:

1516 1517

1521 1522

1529 1530

1531

where $m_{I_a[n]}^t$ is the center location of atom $I_a[n]$ at time t and r_a is the radius for atom type a. With that, we obtain one density field u_a per atom type a (with $a \in \{H, C, ...\}$) and the joint signal for all atom types can be summarized by a vector of density fields $u^t(x)$:

 $u_a^t(\boldsymbol{x}) = 1 - \prod_{n=1}^{|I_a|} \left(1 - D\left(\left\| \boldsymbol{x} - \boldsymbol{m}_{I_a[n]}^t \right\|, r_a
ight)
ight),$

$$\boldsymbol{u}^{t}(\boldsymbol{x}) = (u_{\mathrm{H}}^{t}(\boldsymbol{x}), u_{\mathrm{C}}^{t}(\boldsymbol{x}), \ldots). \tag{B.4}$$

(B.3)

1523 B.2.2 IMPORTANCE-BASED SAMPLING FOR MOLECULES

We sample points from the density field vector \boldsymbol{u}^t by first sampling sets of N_{IS} points $(\{\boldsymbol{x}_1,\ldots,\boldsymbol{x}_{N_{\text{IS}}}\})$, where each point \boldsymbol{x}_i is from a normal distribution centered around one of the input molecule atom locations \boldsymbol{m}_k^t at time t (with $k \in \{1,\ldots,N_{\text{atoms}}\}$, where N_{atoms} is the number of atoms for the considered molecule):

$$\boldsymbol{x}_i \sim \mathcal{N}(\boldsymbol{m}_k^t, \boldsymbol{\sigma}^2)$$
 (B.5)

Then we compute the associated signal vectors corresponding to the sampled points x_i , i.e., $u^t(x_i)$. We use $\sigma = 0.5$ Å in all experiments, and add global nodes at the initial atom positions. Additionally, we randomly sample points in the input space and compute their signal.

1536 B.2.3 REFINEMENT

In autoregressive sampling, errors accumulate across inference steps, causing out-of-distribution issues. To mitigate this, we employed the energy minimization procedure described in (Yu et al., 2024), implemented using OpenMM (Eastman et al., 2017).

1542 B.2.4 DATA AUGMENTATION

¹⁵⁴³ During training we apply random rotation, uniform between $[0, 2\pi)$ along the three Euler angles to each training sample.During training we apply random rotation, uniform between $[0, 2\pi)$ along the three Euler angles to each training sample.

1547 1548

1541

1563

1564

1566 B.3 NEURAL SAMPLING OPERATOR

In accordance with literature (Lipman et al., 2022; Liu et al., 2022), we assume a flow Φ to be created via a parameterized (θ) vector field v_{θ} :

$$rac{d\Phi(s, oldsymbol{z})}{ds} = v_{ heta}(s, \Phi(s, oldsymbol{z}), z_{ ext{cond}}, N_{ ext{cond}})$$
 $\Phi(0, oldsymbol{z}) = oldsymbol{z}$

1575

1585

Here *s* serves as the diffusion time. We build upon the idea of classifier-free guidance (Ho & Salimans, 2022) for flow matching (Zheng et al., 2023) to incorporate previous molecule conformations as condition a (z_{cond}). We further build upon ITO (Schreiner et al., 2023) to predict for more than one atomistic time step into the future and therefore also condition on a number of time steps (N_{cond}). Algorithm 1 shows how v_{θ} can be trained given a series of MD trajectories. Algorithm 2 then shows how new samples can be generated using the trained flow matching velocity field v_{θ} and a given previous conformation state as well as the number of atomistic time steps. The flow matching guidance parameter ω and the employed number of ODE steps (N_{ODE}) serve as hyperparameters.

Algorithm 1 Training UPT++ sampling operator

1: Inputs:

• $n_{\mathcal{Z}}$ MD-trajectories $\mathcal{Z} = \{\hat{z}_j^0, \dots, \hat{z}_j^i, \dots, \hat{z}_j^{N_j}\}_{j=0}^{n_{\mathcal{Z}}}$ with samples \hat{z}_j^i taken at times $i \Delta t$ 1587 • max lag N_{max} • p_{cond} probability of conditional training 1590 2: Initialize UPT++ flow matching model v_{θ} 1591 3: $\mathcal{Z}' = \text{Concatenate}\left(\left\{\hat{z}_j^0, \dots, \hat{z}_j^{N_j - N_{\max}}\right\}_{j=0}^{n_{\mathcal{Z}}}\right)$ 1592 4: while not converged do 1593 $\hat{oldsymbol{z}}_{j}^{i} \sim ext{Choice}\left(oldsymbol{\mathcal{Z}}^{'}
ight)$ 5: 1594 $N \sim \text{DiscreteUniform}(1, N_{\text{max}})$ 6: 1596 7: $(\tilde{z}_{\text{cond}}, \tilde{N}_{\text{cond}}) \leftarrow (\hat{z}_i^i, N)$ with probability p_{cond} else \emptyset 1597 $\tilde{s} \sim \text{ContinuousUniform}(0, 1)$ 8: 1598 $\tilde{\boldsymbol{z}}_0 \sim \mathcal{N}(0,1)$ 9: $\tilde{\boldsymbol{z}}_s \leftarrow (1-\tilde{s})\tilde{\boldsymbol{z}}_0 + \tilde{s}\hat{\boldsymbol{z}}_i^{i+N}$ 10: Take gradient step on $\nabla_{\theta} \left\| v_{\theta}(\tilde{s}, \tilde{z}_{s}, \tilde{z}_{\text{cond}}, \tilde{N}_{\text{cond}}) - (\hat{z}_{j}^{i+N} - \tilde{z}_{0}) \right\|^{2}$ 11: 12: end while 1602 13: **Output:** • UPT++ flow matching model v_{θ} 1604 1607 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619

Algo	orithm 2 Sampling UPT++ sampling operator
1:	Inputs:
	• trained UPT++ flow matching model v_{θ}
	• Condition state z_{cond}
	• Forward sampling timesteps N_{cond}
	• Number of ODE steps N_{ODE}
	• Guidance parameter ω
2:	$ ilde{oldsymbol{z}}_0 \sim \mathcal{N}(0,1)$
3:	$h \leftarrow \frac{1}{N_{\text{opt}}}$
4:	$v_{\theta, \text{guided}}(.,.) \leftarrow (1-\omega) v_{\theta}(.,.,\emptyset) + \omega v_{\theta}(.,.,z_{\text{cond}}, N_{\text{cond}})$
5 : 1	for $s=1,\ldots,N_{ODE}$ do
6:	$\tilde{z}_{sh} \leftarrow \text{ODEStep}(v_{\theta, \text{guided}}((s-1)h, \tilde{z}_{(s-1)h}), h)$
7:	end for
8:	Output:
	• Sample \tilde{z}_1

1674 B.4 MOLECULAR GRAPH RECONSTRUCTION

We present a systematic procedure for reconstructing the molecular graph, from the predicted den-sity distribution. Figure B1 provides a visual representation of certain steps.

- 1. Evaluate positions on our occupancy field and remove all values below a threshold of 0.5 (used throughout our experiments). Note that reconstructing the molecule based solely on the density channels yields similar results.
- 2. Peak finding: Identify local maxima in the thresholded occupancy map using a maximum filter.
 - 3. For each density channel, select the top N_{α} values, where N_{α} is the expected number of atoms of element α in the molecule.
 - 4. Reconstruct bonds ² using OpenBabel (O'Boyle et al., 2011).
 - 5. Validate the chemical equivalence of the encoded molecule using InChI codes (Landrum, 2016).



Figure B1: Top row: Densities across all channels in a single plot, distinct colors represent different atomic channels. All values exceeding an occupancy threshold are assigned a uniform value. Top left: Original molecular conformation obtained by applying the encoder/decoder only. Bottom left: Raw density maps for each channel, predicted by our latent sampling operator. The density values increase concentrically towards the atomic positions. Bottom right: Extracted peaks indicating atomic positions.

²https://github.com/guanjq/targetdiff/blob/main/utils/reconstruct.py

1728 B.5 ADDITIONAL RESULTS

1730 In Figures B2, B3, we visualize results corresponding to those in Table 2 analogously to Figure 6.



Figure B2: 2AA: AN (3.5). *Left half:* Ramachandran plots comparing 10k UPT++ and 9.8k MD samples. *Right half:* Free energy surface of the same UPT++ and MD samples.



Figure B3: **2AA: AN (1.0).** *Left half:* Ramachandran plots comparing 10k UPT++ and 9.8k MD samples. *Right half:* Free energy surface of the same UPT++ and MD samples.