LEARNING ESCORTED PROTOCOLS FOR MULTISTATE FREE-ENERGY ESTIMATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Estimating relative free energy differences between multiple thermodynamic states lies at the core of numerous problems in computational biochemistry. Traditional estimators, such as Free Energy Perturbation and its non-equilibrium counterpart based on the Jarzynski equality, rely on defining a switching protocol between thermodynamic states and computing the free energy difference from the work performed during this process. In this work, we present a method for learning such switching protocols within the class of escorted protocols, which combine deterministic and stochastic steps. For this purpose, we use Conditional Flow Matching and introduce Conditional Density Matching (CDM) to estimate changes in free energy. We further reduce the variance in the multi-state setting by coupling multiple flows between thermodynamic states into a flow graph of escorted protocols, enforcing estimator consistency across different transition paths.

1 INTRODUCTION

In recent years there has been significant interest in the application of machine learning to problems in computational biochemistry, such as protein folding (Jumper et al., 2021; Abramson et al., 2024; Bose et al., 2024) and molecular conformer generation (Gómez-Bombarelli et al., 2018) for drug discovery (Wan et al., 2022) and materials design (Merchant et al., 2023). While initial approaches focused on generating single static examples of systems of interest, recent efforts have shifted towards generating the full dynamic ensemble of the system (Noé et al., 2019; Holdijk et al., 2023; Tan et al., 2025; Akhound-Sadegh et al., 2024). This shift towards the entire dynamical ensemble has opened up the possibility of using these methods in computational biochemistry beyond what generating static samples allows. Notable among these is the problem of estimating free-energy differences between different thermodynamic states, a crucial aspect of many tasks in this field, such as binding affinity prediction (Mobley and Gilson, 2017) and other components of the drug discovery pipeline (Cournia and Chipot, 2024). Traditional estimators such as multiwindow Free-Energy Perturbation (FEP) (Zwanzig, 1955; Wang et al., 2015) and its non-equilibrium variants, defined through the Jarzynski equality (Jarzynski, 1997), are among the most commonly used for this purpose.

In this work, we specifically consider the class of Escorted Non-EQuilibrium (E-NEQ) estimators (Vaikuntanathan and Jarzynski, 2008; 2011) based on the aforementioned Jarzynski equality. These estimators combine dynamics that preserve a given stationary distribution $p(x, t)$, such as Langevin dynamics, with a deterministic escorting vector $b(x, t)$ field to transition samples between thermodynamic states. To minimize variance in the finite-sample limit, the preserved stationary distribution and the escorting vector field must together satisfy the continuity equation $\frac{\partial p(x,t)}{\partial t} + \nabla \cdot (p(x, t)b(x, t)) = 0$ (Zhong and DeWeese, 2024). Traditionally, finding such a pair is non-trivial.

To overcome this barrier, we propose learning the stationary distribution and escorting vector field jointly by extending the Conditional Flow Matching (CFM) framework (Lipman et al., 2023) with an additional Conditional Density Matching (CDM) objective (Sec. 3). We hypothesize that learning the stationary distribution and driving force in a unified framework will allow them to compensate for each other's small errors, leading to a more accurate estimator than using only vector fields, such as done in methods such as Targeted Free-Energy Perturbation (TFEP) (Jarzynski, 2002) as illustrated in Fig. 1. To further improve the applicability of the framework, we propose two important practical considerations in the form of Lie-Trotter splitting to reduce the computational cost of the work calculation (Sec. 4.1) and a flow graph construction to reduce the combinatorial complexity of learning protocols for multi-state free-energy estimation (Sec. 4.2).

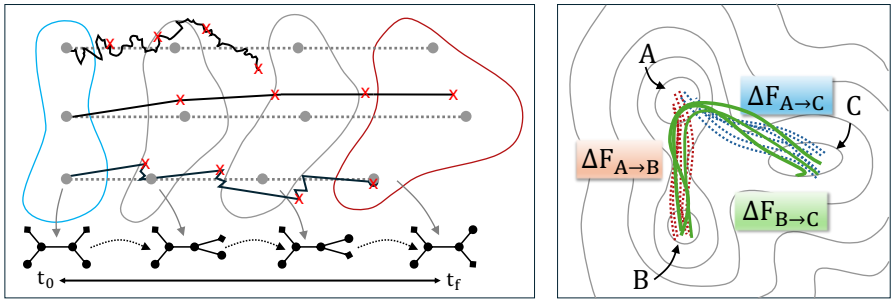

Figure 1: *(a)* Illustration of processes transforming between two thermodynamic states (blue and red) of the same molecular system. Using our proposed Conditional Density Matching framework we learn approximate intermediate densities (grey) along which the learned flow (grey dotted lines) transports samples. The top stochastic line shows how stochastic processes purely targeting the intermediate distributions will not fully reach the target state, while the middle process shows that, due to integration error, deterministic processes may deviate from the true path. Combining stochastic and deterministic processes allows the two approaches to compensate, resulting in more accurate sampling. *(b)* Illustration of combining two free-energy estimates with a central connecting node (A) to obtain the free-energy $\Delta F_{B\to C}$ difference between the nodes not connected.

The remainder of our paper is organized as follows to introduce the core contributions of our work:

- Sec. 2 introduces the theory of non-equilibrium free-energy estimation, focusing on topics such as the Jarzynski equality, escorted switching protocols, and bi-directional estimators.
- Sec. 3 then relates escorted switching protocols with the Flow Matching framework and introduces one of our core contributions in the form of conditional Density Matching. Collectively these form the learning objectives for our proposed Escorted Non-EQuilibrium (E-NEQ) estimator.
- To improve the efficiency of multistate free-energy estimation, Sec. 4.1 proposes a Lie–Trotter splitting scheme and Sec. 4.2 proposes the concept of *Escorted Protocol Flow Graphs* to combine multiple trained escorted switching protocols to reduce training requirements.
- In Sec. 5 the E-NEQ estimator is experimentally validated using the well-known Alanine Dipeptide (ADP) system. ADP is an important benchmark due to its multistate free-energy surface.

2 Background

As stated, in this work we are interested in estimating the relative free-energy differences $\Delta F_{A\to B}$ between pairs of thermodynamic states A and B. Each thermodynamic state A is associated with a potential energy function $U_A : \mathbb{R}^{3N} \to \mathbb{R}$, which defines the distribution of possible microscopic states $x \in \mathbb{R}^{3N}$ through the Boltzmann distribution

$$p_A(x) = \frac{1}{Z_A} \exp\left(-\beta U_A(x)\right), \quad Z_A = \int_{\mathbb{R}^{3N}} \exp\left(-\beta U_A(x)\right) \mathrm{d}x. \tag{1}$$

Generally, $\beta = \frac{1}{k_B T}$, with k_B representing the Boltzmann constant and T the temperature of the system in Kelvin. Using the partition function Z_A, the equilibrium free-energy of state A can be defined as $F_A = -\beta^{-1} \ln Z_A$ and their relative free-energy difference as

$$\Delta F_{A\to B} = F_B - F_A = -\beta^{-1} \ln \frac{Z_B}{Z_A}. \tag{2}$$

As exemplified by the definition of the partition function, estimating free energies by directly integrating over the entire space $\mathbb{R}^{3N}$ is prohibitively expensive. As such, various fields have developed methods to estimate this quantity more efficiently.

2.1 Non-Equilibrium Free-Energy Estimation

We will specifically focus on the class of *alchemical Non-EQuilibrium* (NEQ) free-energy estimation methods, which are a subset of the class of path-based free-energy estimation methods. In path-based free-energy estimation, a central object of in is the switching protocol $U_{A\to B}$.

Definition 2.1 (Switching Protocol)**.** *Given two thermodynamic states A and B with potential energy functions U_A and U_B, a* ***switching protocol*** *$U_{A\to B} : \mathbb{R}^{3N} \times [0, t_f] \to \mathbb{R}$ is a time-dependent potential energy function with boundary conditions $U_{A\to B}(x, 0) = U_A(x)$ and $U_{A\to B}(x, t_f) = U_B(x)$.*

Using such a switching protocol $U_{A\to B}$, a system x initially equilibrated with respect to the potential energy function U_A, can be driven from state A to state B according to the time-dependent dynamics

$$\partial_t \rho(x,t) = \mathcal{L}_t \rho(x,t), \qquad \rho(x,0) = p_A(x), \tag{3}$$

where $\mathcal{L}_t$ is a time-dependent forward (Fokker–Planck/Liouville) generator, specified so that for all t the target distribution $p(x,t) \propto \exp\left(-\beta U_{A\to B}(x,t)\right)$ is stationary for the frozen-time dynamics, i.e. $\mathcal{L}_t p(\cdot,t) = 0$. Notably, it does not have to be the case that the system is in equilibrium at all times, i.e. $\rho(x,t) = p(x,t)$ does not need to hold for all t.

Under such dynamics defined by the protocol $U_{A\to B}$, the second law of thermodynamics tells us that, with thermodynamic work defined as

$$W_{A\to B}(\mathbf{x}) = \int_0^{t_f} \frac{\partial U_{A\to B}(x_t,t)}{\partial t}\, \mathrm{d}t, \tag{4}$$

the average work performed on the system along trajectories $\mathbf{x}$ generated by the dynamics in Eq. (3) is an upper bound on the free-energy difference between the two states

$$\Delta F_{A\to B} \leq \langle W_{A\to B}(\mathbf{x}) \rangle_{U_{A\to B}}. \tag{5}$$

where the shorthand $\langle\cdot\rangle_{U_{A\to B}}$ denotes the ensemble average over trajectories $\mathbf{x}$ generated by the dynamics in Eq. (3). Crucially, equality only holds here for quasi-static (infinitely slow, $t_f = \infty$) processes where the system remains in equilibrium at all times, i.e. $\rho(x,t) = p(x,t)$ for all t.

2.1.1 Jarzynski Equality

For any such switching protocols, Jarzynski (Jarzynski, 1997) showed the remarkable result that instead of restricting to quasi-static processes to obtain an equality for the thermodynamic work we can consider the ensemble average of the exponential of the work.

Theorem 2.1 (Jarzynski Equality (JE) (Jarzynski, 1997))**.** *Given switching protocol $U_{A\to B}$ and time-dependent dynamics with frozen time stationary distribution $p(x,t) \propto \exp\left(-\beta U_{A\to B}(x,t)\right)$ for all $t \in [0,t_f]$ as in Eq. (3), we have that along trajectories $\mathbf{x}$ generated by the dynamics*

$$\left\langle e^{-\beta W_{A\to B}(\mathbf{x})} \right\rangle_{U_{A\to B}} = e^{-\beta \Delta F_{A\to B}}. \tag{6}$$

Important to note here is that the class of dynamical processes considered in the JE only require stationarity. The JE does not depend on ergodicity and is therefore valid for a wide range of dynamics including both time-dependent stochastic dynamics such as underdamped and overdamped Langevin as well as deterministic dynamics, e.g. Hamiltonian. Extra care is needed when momenta are considered (e.g., underdamped Langevin dynamics). In this case, a state-dependent Hamiltonian, rather than the potential energy, should be used in the definition of work in Eq. (4).

Using standard Monte Carlo integration with trajectories $\mathbf{x}_n$ generated according to Eq. (3), the JE provides a consistent but generally biased estimator of the $\Delta\mathcal{F}_{A\to B}$ in the finite-sample setting:

$$\Delta F_{A\to B} \approx \Delta\mathcal{F}_{A\to B} = -\beta^{-1} \ln\left(\frac{1}{N}\sum_{n=1}^{N} e^{-\beta W_{A\to B}(\mathbf{x}_n)}\right), \tag{7}$$

JE Estimator Variance Both the bias and the variance of this estimator grow with the excess work $W^{\text{ex}} = \langle W\rangle_{U_{A\to B}} - \Delta F_{A\to B}$ (Geiger and Dellago, 2010; Gore et al., 2003), which for every time t is lower bounded by the Kullback–Leibler divergence between the two distributions (Vaikuntanathan and Jarzynski, 2009):

$$W_t^{\text{ex}} \geq \beta^{-1} D_{\text{KL}}(\rho(x,t)||p(x,t)), \quad \forall t \in [0,t_f]. \tag{8}$$

As such, intuitively, while the JE is valid for any switching protocol the variance and the bias of the estimator are determined by how much the instantaneous distribution $\rho(x,t)$ *lags* behind the stationary distribution $p(x,t)$ at every time t defined by the switching protocol $U_{A\to B}$.

2.1.2 Escorted Jarzynski Equality

To reduce the amount of lag between the instantaneous distribution $\rho(x,t)$ and the target distribution $p(x,t)$ Vaikuntanathan and Jarzynski (2008) introduced the concept of escorted switching protocols.

Definition 2.2 (Escorted switching protocol)**.** *Given two thermodynamic states A and B with potential energy functions U_A and U_B, an **escorted switching protocol** $(U_{A\to B}, b)$ consists of a time-dependent potential $U_{A\to B} : \mathbb{R}^{3N}\times[0,t_f] \to \mathbb{R}$ and a time-dependent vector field $b : \mathbb{R}^{3N}\times[0,t_f] \to \mathbb{R}^{3N}$ with boundary conditions $U_{A\to B}(x,0) = U_A(x)$ and $U_{A\to B}(x,t_f) = U_B(x)$ and $b(x,0) = b(x,t_f) = 0$.*

Given an escorted switching protocol, the escorted *dynamics* driving a system from state A to state B adds deterministic advection based on the time-dependent vector field b

$$\partial_t \hat{\rho}(x,t) = \widehat{\mathcal{L}}_t \hat{\rho}(x,t) = \mathcal{L}_t \hat{\rho}(x,t) - \nabla\cdot\big(b(x,t)\hat{\rho}(x,t)\big), \qquad \hat{\rho}(x,0) = p_A(x). \tag{9}$$

Here $\mathcal{L}_t$ denotes the forward operator that satisfies frozen time stationarity $\mathcal{L}_t p(\cdot,t) = 0$ for $p(x,t) \propto \exp(-\beta U_{A\to B}(x,t))$, as defined in Eq. (3). Under escorted dynamics, Vaikuntanathan and Jarzynski (2008) showed that with an alternative definition of the work, the JE still holds.

Theorem 2.2 (Escorted Jarzynski Equality (E-JE) Vaikuntanathan and Jarzynski (2008))**.** *Given an escorted switching protocol $(U_{A\to B}, b)$ and time-dependent dynamics as in Eq. (9), we have that along trajectories $\mathbf{x}$ generated by the dynamics*

$$\left\langle e^{-\beta\,\hat{W}_{A\to B}(\mathbf{x})} \right\rangle_{(U_{A\to B},b)} = e^{-\beta\Delta F_{A\to B}}. \tag{10}$$

where the escorted work $\hat{W}_{A\to B}(\mathbf{x})$ along trajectory $\mathbf{x}$ is defined as

$$\hat{W}_{A\to B}(\mathbf{x}) = \int_0^{t_f} \Big(\partial_t U_{A\to B}(x_t,t) + b(x_t,t)\cdot\nabla U_{A\to B}(x_t,t) - \beta^{-1}\nabla\cdot b(x_t,t)\Big)\,\mathrm{d}t. \tag{11}$$

Notably, because b is not involved in enforcing that the dynamics have stationary distribution $p(x,t) \propto \exp\left(-\beta U_{A\to B}(x,t)\right)$ it can be freely chosen to minimize the amount of lag and with that the variance of the escorted estimator $\Delta\widehat{\mathcal{F}}_{A\to B} = -\beta^{-1}\ln(\frac{1}{N}\sum_{n=1}^{N} e^{-\beta\hat{W}_{A\to B}(\mathbf{x}_n)})$. Notably among the choices of b are the case where $b = 0$, which is equivalent to the non-escorted case, as well as the following optimal choice of b as shown by Zhong et al. (2023):

Theorem 2.3 (Optimal Escorted Switching Protocol (Zhong et al., 2023))**.** *If $b(x,t)$ and $p(x,t) \propto \exp\left(-\beta U_{A\to B}(x,t)\right)$ collectively solve the continuity equation $\partial_t p + \nabla\cdot(p\,b) = 0$, then, for every trajectory $\mathbf{x}$ generated by the dynamics Eq. (9), we have $\hat{W}_{A\to B}(\mathbf{x}) = \Delta F_{A\to B}$.*

Crucially, what this shows is that if we learn b and $U_{A\to B}$ collectively to solve the continuity equation, while maintaining the boundary conditions $U_{A\to B}(x,0) = U_A(x)$ and $U_{A\to B}(x,t_f) = U_B(x)$, then a single trajectory $\mathbf{x}$ generated by the dynamics in Eq. (9) suffices to estimate the free-energy difference $\Delta F_{A\to B}$. In the remainder of this work we will explore how we can learn these components using Conditional Flow Matching and our proposed Conditional Density Matching.

2.1.3 Bi-directional Sampling and the Crooks Identity

In addition to introducing the escorting velocity field b, a second approach to reduce the variance of our estimator of the free-energy difference is to use bi-directional estimators. Instead of only evolving samples from the initial thermodynamic state A, it is equally straightforward to start from state B and define a reverse protocol $U_{B\to A}$ with boundary conditions $U_{B\to A}(x,0) = U_B(x)$ and $U_{B\to A}(x,t_f) = U_A(x)$. In the context of non-equilibrium free-energy estimation, this idea underpins the Crooks Fluctuation Theorem (CFT) (Crooks, 1998):

$$\frac{P_{A\to B}(W)}{P_{B\to A}(-W)} = \exp[\beta\,(W - \Delta F_{A\to B})]\,. \tag{12}$$

Here $P_{A\to B}(W)$ is the probability of observing work W for a process starting from $x_0 \sim p_A$ and evolved under the forward dynamics, while $P_{B\to A}(-W)$ is the probability of observing work $-W$ under the time-reversed protocol. This fluctuation theorem holds both for the non-escorted case and the escorted case with $\hat{W}_{A\to B}(\mathbf{x})$ replacing $W(\mathbf{x})$.

For the escorted case, the time-reversed protocol $(U_{B\to A}, \tilde{b})$ is given by $U_{B\to A}(x,t) = U_{A\to B}(x,t_f - t)$ and $\tilde{b}(x,t) = -b(x,t_f - t)$. Consequently the time-reversed escorted dynamics driving the system from state B to state A are given by

$$\partial_t \hat{\rho}_R(x,t) = \tilde{\mathcal{L}}_t \hat{\rho}_R(x,t) - \nabla\cdot\big(\tilde{b}(x,t)\,\hat{\rho}_R(x,t)\big), \qquad \hat{\rho}_R(x,0) = p_B(x). \tag{13}$$

Here $\tilde{\mathcal{L}}_t$ denotes the reverse operator of a non-escorted dynamics that satisfies frozen time stationarity $\tilde{\mathcal{L}}_t \tilde{p}(x,t) = 0$ for $\tilde{p}(x,t) = p(x, t_f - t) \propto \exp(-\beta U_{B\to A}(x,t))$. In the case where dynamics with momenta are considered, the momenta have to be reversed in the reverse generator $\tilde{\mathcal{L}}_t$.

Crucially, the Crooks Fluctuation Theorem expresses $\Delta F_{A\to B}$ as the solution to a single-parameter problem. Rather than directly estimating $\Delta F_{A\to B}$ via the JE, which often suffers from high variance and finite sample bias due to its exponential average, the CFT can be inverted to solve for $\Delta F_{A\to B}$ as the parameter that best fits a given set of observed forward and reverse work for a given protocol.

The remainder of this work will only consider using the CFT instead of the JE. Practically, this is implemented by the Bennett Acceptance Ratio (BAR) estimator (Bennett, 1976).

2.2 Related Work

Traditional Free-Energy Estimation In addition to the traditional alchemical non-equilibrium free-energy estimation approaches introduced in the previous section, a large collection of other approaches have been proposed. We briefly outline the core lines of work here. Closely related to the approaches considered in this work are alchemical *equilibrium* approaches such as (iterative) Free-Energy Perturbation (FEP) (Zwanzig, 1955) and Thermodynamic Integration (TI) (Kirkwood, 1935), which depend on the Zwanzig equation (Zwanzig, 1955). Next are the *non-alchemical* path-based approaches such as transition path sampling (Bolhuis et al., 2002) and nudged elastic band sampling (Henkelman et al., 2000). Non-alchemical approaches rely on frameworks such as Transition State Theory (Truhlar et al., 1996) and the Arrhenius/Eyring equation (Eyring, 1935).

Neural Free-Energy Estimation Methods Within the family of alchemical free-energy estimation approaches, most proposed neural approaches focus on the Targeted Free-Energy Perturbation method (TFEP) (Jarzynski, 2002). The TFEP method can be considered a special case of the escorted switching protocol where only the deterministic vector field b is used to define an invertible mapping. TFEP has been studied within the machine learning context in Wirnsberger et al. (2020); Rizzi et al. (2023; 2021); Erdogan et al. (2024); Zhao and Wang (2023), which all use variants of MLE-trained discrete normalizing flows or Flow Matching. Within the same alchemical family, Máté et al. (2024; 2025) propose a neural version of Thermodynamic Integration. Lastly, most closely related to methods proposed here is the work by He et al. (2025), which also parameterizes an escorted switching protocol. Crucially, their proposed method FEAT uses two different protocols for the forward and backward process between two states and does not consider the multistate setting.

3 Learning Escorted Protocols using Flow and Density Matching

Having discussed the background of (Escorted-) Non-Equilibrium free-energy estimators, we now propose a new method for parameterising and learning the escorted switching protocol $(U^\theta_{A\to B}, b^\phi)$ such that they collectively solve the continuity equation, and adhere to the boundary conditions $U^\theta_{A\to B}(x,0) = U_A(x)$ and $U^\theta_{A\to B}(x,t_f) = U_B(x)$. We propose to parameterise both components by learning b^ϕ_t using standard Conditional Flow Matching (CFM) (Lipman et al., 2023) and U^θ_t using an extension of CFM, which we call Conditional Density Matching.

3.1 Learning b^θ_t using Conditional Flow Matching

Conditional Flow Matching (CFM) is a general framework for learning a vector field v^ϕ_t that drives samples from one arbitrary distribution p_0 to another p_1 along a set of time-dependent intermediate distributions p_t. As it can generally be assumed that there is no access to the ground truth vector field v_t or samples from the intermediate distributions p_t beyond the initial and final distributions, CFM approaches this by considering a *conditional* time-dependent distribution $p_t(x_t \mid z)$ generated by a *conditional* vector field $v_t(x_t \mid z)$ and a coupling distribution $q(z)$. A common choice is to have $q(z)$ defined as an Optimal Transport coupling (Tong et al., 2024) to enforce v_t to follow straight paths.

Using this coupling of conditional and marginal vector fields, Lipman et al. (2023) showed that if we learn the vector field $v^\phi(x_t,t)$ by regressing on the conditional vector field $v_t(x_t, t \mid z)$ using

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{t\sim\text{Uni}(0,1),\, z\sim q(z),\, x_t\sim p_t(x_t|z)} \left[\left\| v^\phi(x_t,t) - v_t(x_t \mid z) \right\|^2 \right] \tag{14}$$

then this is equivalent to regressing on the vector field $v_t(x_t)$ directly.

Crucially, given a paired conditional vector field and conditional distribution, under minor regularity conditions, for any choice of coupling distribution $q(z)$ the marginal vector field $v_t(x_t) = \langle \frac{v_t(x_t|z)p_t(x_t|z)}{p_t(x_t)} \rangle_{q(z)}$ and marginal distribution $p_t(x_t) = \langle p_t(x_t|z) \rangle_{q(z)}$ are shown to jointly solve the continuity equation (Tong et al., 2024). As such, if we set $p_0 = p_A$, $p_1 = p_B$, $t_f^{-1} v^\phi(x, s) = b^\phi(x, t)$ under the rescaling $t = st_f$, and define $p_s(x) = p(x, t) \propto \exp\left(-\beta U^\theta_{A\to B}(x, t)\right)$, then CFM provides a valid approach to finding the escorting vector field b^ϕ for our escorted switching protocol.

3.2 Learning U_t^θ using Conditional Density Matching

This still leaves us with the problem of learning the time-dependent potential $U^\theta(x, t)$. For this purpose, we employ a similar trick to that used by CFM to learn the time-dependent potential $p^\theta(x, t) \propto \exp\left(-\beta U^\theta_{A\to B}(x, t)\right)$ from the conditional time-dependent distribution $p_t(x_t \mid z)$ and consider the following extended version of the Maximum Likelihood Estimation (MLE) objective

$$\mathcal{L}_{\text{DM}} = \mathbb{E}_{t\sim\text{Uni}(0,1),\, x_t\sim p_t(x_t)} \left[-\log p^\theta(x_t, t) \right], \tag{15}$$

which adds an additional time dependence. To highlight the similarity with Flow Matching, we denote this object as the **Density Matching (DM)** objective.

Due to $\mathcal{L}_{\text{DM}}$ being defined as an expectation over the unknown distribution $p_t(x_t)$, learning $U^\theta_{A\to B}(x, t)$ directly using this objective is not possible. However, similar to the CFM objective, we can equivalently express another maximum likelihood objective using the conditional distribution $p_t(x_t \mid z)$ instead. We denote this as the **Conditional Density Matching (CDM)** objective:

$$\mathcal{L}_{\text{CDM}} = \mathbb{E}_{t\sim\text{Uni}(0,1),\, z\sim q(z),\, x_t\sim p_t(x_t|z)} \left[-\log p^\theta(x_t, t) \right]. \tag{16}$$

Similar to the CFM objectives, the DM and CDM objectives have equivalent gradients, $\nabla_\theta \mathcal{L}_{\text{DM}} = \nabla_\theta \mathcal{L}_{\text{CDM}}$, and thus MLE using the conditional distribution will result in the same learned marginal distribution p_t^θ as MLE using the marginal distribution.

In summary, combining the CFM objective reviewed above with the proposed CDM objective, we thus learn a pair of escorting vector field $b^\phi(x, t)$ and potential $U^\theta_{A\to B}(x, t)$ that collective solve the continuity equation. When these two components are collectively used as the escorted switching protocol in an E-NEQ estimator, this results in a low-variance estimate of $\Delta F_{A\to B}$.

4 Efficient Multi-State Free-Energy Estimation

With the previous section proposing a method for learning the escorted protocol $(U^\theta_{A\to B}, b^\phi)$, we now consider some practical restrictions of E-NEQ estimators. Namely, (i) the work calculation becoming prohibitively expensive because a very small global time step $\mathrm{d}t$ being required to deal with the numerical instability of the stationary distribution preserving dynamics, and (ii) in the multi-state setting the number of escorted switching protocols grows exponentially with the number of states. We will address both of these issues in the following two sections. In the appendix we further discuss three more standard, but important, practical considerations.

4.1 Efficient Work Calculation by Lie-Trotter Splitting

Molecular Dynamics simulation is known to quickly become numerically unstable when using large time-steps due to the potential energy including sharply peaked components such as Lennard–Jones potentials. As such, when simulating the dynamics of our escorted protocol in Eq. (9), we are required to use a very small global time-step $\mathrm{d}t$. Unfortunately, as a result of this, in the work calculation we therefore must evaluate the divergence many times, which can be computationally very expensive for any system of considerable size. Ideally, we would therefore like to decouple the divergence calculation from the unstable stationary distribution preserving dynamics.

For this purpose we propose to use Lie–Trotter (Trotter, 1959) splitting to split the combined dynamics into two separate steps; first a step using $b(x, t)$ and then a step using the stationarity-preserving distribution preserving dynamics. Given our escorted dynamics $\widehat{\mathcal{L}}_t = \mathcal{L}_t + \mathcal{E}_t$ where $\mathcal{E}_t = -\nabla \cdot (b(x, t)\rho(x, t))$ is the escort transport, a Lie–Trotter step of size $h = \mathrm{d}t$ frozen at time t

approximates the full evolution by composing the subflows as

$$e^{h\widehat{\mathcal{L}}_t} \approx e^{h\mathcal{L}_t} e^{h\mathcal{E}_t}. \tag{17}$$

Under such split dynamics, a simplified version of the Escorted Jarzynski Equality holds.

Theorem 4.1 (Split Escorted Jarzynski Equality)**.** *Given an escorted switching protocol $(U_{A\to B}, b)$ and split escorted time-dependent dynamics $e^{h\widehat{\mathcal{L}}_t} \approx e^{h\mathcal{L}_t} e^{h\mathcal{E}_t}$, where $\mathcal{L}_t$ is the stationarity-preserving operator (i.e. $\mathcal{L}_t p(x,t) = 0$ with $p(x,t) \propto e^{-\beta U_{A\to B}(x,t)}$) and $\mathcal{E}_t\rho = -\nabla\cdot\big(b(x,t)\rho\big)$ is the escort transport, consider a Lie–Trotter time grid $t_k = kh$ with $N \cdot h = t_f$ and the split update*

$$x'_k = x_k + b(x_k, t_k)h, \qquad x_{k+1} \sim K_{t_k}(x'_k, \cdot), \tag{18}$$

where K_t is any transition kernel that preserves $p(x,t)$ at frozen time t. If $x_0 \sim p_A(x)$, then

$$\left\langle e^{-\beta\,\widehat{W}_{A\to B}(\mathbf{x},h)} \right\rangle_{(U_{A\to B},b,h)} = e^{-\beta\Delta F_{A\to B}} + \mathcal{O}(h), \tag{19}$$

where the split escorted work is defined by

$$\widehat{W}_{A\to B}(\mathbf{x},h) = \sum_{k=0}^{N-1} \left(\frac{\partial U_{A\to B}\big(x'_k,\, t_k\big)}{\partial t} - \beta^{-1}\,\nabla\cdot b(x_k,t_k) \right) h. \tag{20}$$

In particular, (19) holds exactly in the limit $h \to 0$, recovering the Escorted Jarzynski Equality.

Proof. Consider the escort step as a deterministic kernel $L_{t_k}(x'|x) = \delta_{x+b(x,t_k)h}(x')$ and the stationarity-preserving step as a kernel K_{t_k} with stationary distribution $p(x,t_k)$ such that the composed update is the kernel $P_{t_k} = K_{t_k} \circ L_{t_k}$. Then the proof follows from the discrete-time E-JE in Vaikuntanathan and Jarzynski (2011). □

While this still requires us to use the global time discretization $h = \mathrm{d}t$ for each individual substep as well as for the calculation of the work, it allows for smaller internal time discretization within the stationarity-preserving kernel K_t independent of the escort step. This split can be used to significantly reduce the number of divergence evaluations $\nabla_x \cdot b(x,t)$ needed in the work calculation.

For reference, operator splitting is a common topic in MD simulation (Frenkel and Smit). It is, for example, at the core of the BAOAB splitting for underdamped Langevin dynamics and the velocity Verlet integrator for Hamiltonian Monte Carlo (Leimkuhler and Matthews, 2013; Swope et al., 1982).

4.2 Efficient Multi-state FE estimation using Flow-Graph Based MBAR

With the combined CFM and CDM approach to learn the escorted switching protocol $(U^{\theta}_{A\to B}, b^{\phi})$ between the two thermodynamic states A and B, we now consider how to extend this to the case of estimating the relative free-energy difference $\{\Delta F_{i\to j}\}_{i=1,j=1}^{K}$ between a collection of K thermodynamic states. Within the context of the drug-discovery pipeline, lead optimization for example often involves large assays of ligands for which we need to assess their binding affinity (Wang et al., 2015).

Naively, using a single reference state k we can train $K-1$ escorted switching protocols $\{(U^{\theta_i}_{k\to i}, b^{\phi_i}_{k\to i}\}_{i=1}^{K}$ to calculate $\{\Delta\mathcal{F}_{i\to j}\}_{i=1,j=1}^{K}$ using the E-JE. Considering that Free-Energy differences are a state function, we can then obtain all not-directly connected estimates as $\Delta\mathcal{F}_{i\to j} = -\Delta\mathcal{F}_{k\to i} + \Delta\mathcal{F}_{k\to j}$. While this would result in a consistent estimate for all pairs of states, it is generally understood that to obtain accurate multi-state free-energy it is best to obtain all $\{W_{i\to j}\}_{i=1,j=1}^{K}$ individually and use the self-consistent Multistate Bennett Acceptance Ratio (MBAR) (Shirts and Chodera, 2008) to obtain $\{\Delta\mathcal{F}_{i\to j}\}_{i=1,j=1}^{K}$. Following this line of thinking, a straightforward extension of the CFM/CDM framework would therefore be to learn a $U^{\theta_{i,j}}_{i\to j}$ and escorting vector fields $b^{\phi_{i,j}}_{i\to j}$ for each pair of states i and j individually and use them to obtain $\hat{W}_{i\to j}$ according to Eq. (4). However, for a set of K thermodynamic states, this would require $K(K-1)/2$ models, which would quickly become infeasible.

Instead, we propose to construct a Escorted Protocol Flow Graph.

Definition 4.1 (Escorted Protocol Flow Graph (EPFG)). *Given a collection of K thermodynamic states and $K-1$ escorted switching protocols $\{(U^{\theta_i}_{k\to i}, b^{\phi_i}_{k\to i}\}^K_{i=1,i\neq k}$ with central node k an **Escorted Protocol Flow Graph (EPFG)** is constructed by considering the K thermodynamic states to be the nodes and the escorted switching protocols $\{(U_{k\to i}, b_{k\to i}\}^K_{i=1,j=1}$ between all states i and j to represent the edges. Here all edges not connected to the central node k are given by*

$$(U_{i\to j}(x,t), b_{i\to j}(x,t)) = \begin{cases} (U^{\theta_i}_{k\to i}(x, t_f - 2t), -2b^{\phi_i}_{k\to i}(x, t_f - 2t)) & \text{if } 0 \leq t < \frac{t_f}{2} \\ (U^{\theta_j}_{k\to j}(x, 2t - t_f), 2b^{\phi_j}_{k\to j}(x, 2t - t_f)), & \text{if } \frac{t_f}{2} \leq t \leq t_f. \end{cases} \tag{21}$$

Using the EPFG and the escorted protocols defined as its edges, including the constructed concatenated ones, we can now obtain $\{W_{i\to j}\}^K_{i=1,j=1}$ for all combinations of states i and j while only requiring the training of $K-1$ protocols. This enables more accurate free-energy estimates using the MBAR estimator compared to using the pair based approach while minimizing the training costs.

5 Experiments

We evaluate the feasibility of learning the thermodynamic flow network of coupled U^θ and escorting field b^ϕ using the well-known alanine dipeptide (ADP) system. While limited in size, ADP shows similar challenges to larger systems such as high-energy barriers and solvent effects. Furthermore, as ADP in solvent has six distinct metastable states α_L, α_D, β, $C5$ and α' as visualized in Fig. 2. These states are known to exhibit significant differences in conformational free energy it serves as an well-understood benchmark for the multi-state setting.

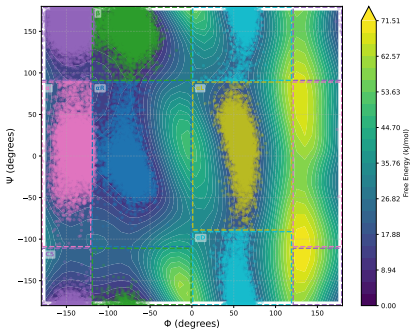

Figure 2: Umbrella sampling FES reconstruction and training samples for ADP.

Data Generation For each metastable state, we generated 10,000 samples using a harmonic flat-bottom constraint on the torsion angles defined by the metastable state boundaries, as specified in the appendix. Training samples for each state are visualized in Fig. 2. We observe that the training samples cover each state but do not represent the correct ratios as observed in full equilibrium sampling.

Baselines To obtain a baseline value to validate our method against, we performed an extensive Umbrella Sampling (US) estimate using a large number of windows and the same setup as used for data generation. Additionally, we compare our method against a neural version of Targeted Free-Energy Perturbation (TFEP). For our comparison we will use the vector field b^ϕ learned using Conditional Flow Matching to implement the TFEP method. As such, the only difference between TFEP and E-NEQ can come from the inclusion of the stationary distribution preserving dynamics.

Model Details We used the same model definitions and training setup for all learned escorted protocols. The escorting vector field b^ϕ is implemented using the SE(3)-equivariant graph neural network (Satorras et al., 2021) with an additional learnable time-embedding component to make it time-dependent (Tan et al., 2025). We use optimal transport coupling q (Tong et al., 2024) to enforce that the escorting vector field and the time-dependent density not only collectively solve the continuity equation but also align with the dynamical optimal transport problem (Benamou and Brenier, 2000). While not discussed in depth in this work, there is a close connection between the dynamical optimal transport problem and the amount of dissipated work in escorted switching protocols (Zhong et al., 2023; Zhong and DeWeese, 2024). We leave exploring this interplay further for future work.

The potential U^θ is parameterised as the negative log probability of a discrete normalizing flow with conditional affine coupling layers and a similar learnable time-embedding component. We chose to use a discrete-time normalising flow here, instead of a more general energy-based model, to minimize the complexity of maximum-likelihood training.

Integration Details For the implementation of the escorted dynamics in Eq. (9), we use the following setup. The escorting vector field is implemented in all experiments using a Runge–Kutta integrator, while we consider three different options for the stationary preserving component of the split dynamics: underdamped Langevin, overdamped Langevin, and Hamiltonian Monte Carlo, all of which use Metropolis-Hastings (MH) correction to preserve the stationary distribution in finite time. The split operators $\mathcal{E}$ and $\mathcal{L}$ are iteratively applied for 100 time steps, during which the work is calculated.

Table 1: Quantitative results of the estimated ΔF between the central α_R state of the escorted flow graph and the directly connected states, as well as the Mean Absolute Error (MAE) over all pairs of states including those not directly connected. The $\mathcal{L}$ column denotes the type of integrator used: Underdamped Langevin (UL), Overdamped Langevin (OL) or Hamiltonian Monte Carlo (HMC). EPFG denotes whether the estimates were obtained using pairwise summation (EPFG=N), or using the escorted protocol flow graph (EPFG=Y).

Method	$\mathcal{L}$	EPFG	α_L	α_D	β	$C5$	α'	MAE
US	-	-	7.42 ± 0.16	12.07 ± 0.40	-1.11 ± 0.03	1.37 ± 0.05	6.55 ± 0.13	-
TFEP	-	N	8.60 ± 0.05	12.39 ± 0.06	0.77 ± 0.04	2.39 ± 0.04	5.78 ± 0.03	1.17
E-NEQ	UL	N	7.93 ± 0.13	12.82 ± 0.16	0.54 ± 0.12	1.47 ± 0.21	5.96 ± 0.06	0.93
(ours)	UL	Y	**7.35 ± 0.12**	12.54 ± 0.15	0.78 ± 0.11	**1.31 ± 0.20**	**6.14 ± 0.06**	0.88
	OL	N	8.22 ± 0.07	12.86 ± 0.08	**0.26 ± 0.12**	1.78 ± 0.10	5.59 ± 0.05	0.96
	OL	Y	8.06 ± 0.05	**12.09 ± 0.06**	**0.24 ± 0.05**	1.69 ± 0.05	7.04 ± 0.04	**0.59**
	HMC	N	7.96 ± 0.12	13.05 ± 0.17	0.62 ± 0.09	**1.57 ± 0.22**	5.97 ± 0.05	0.99
	HMC	Y	7.33 ± 0.11	12.89 ± 0.16	0.50 ± 0.09	1.83 ± 0.20	5.94 ± 0.06	0.95

5.1 Results

In Tab. 1, we report the main quantitative results of learning the escorting vector field b^ϕ and corresponding potential U^θ for use in the E-NEQ estimator, compared to the TFEP approach. We note that the α_R state was chosen as the central node when constructing the flow graph. As such, all reported free-energy differences are relative to α_R.

E-NEQ improves over TFEP for the same amount of divergence calculations Across all three integration approaches, we find that E-NEQ with learned escorted dynamics outperforms TFEP. This holds for both the directly reported ΔF values using the central α_R state and the MAE of the free-energy differences across all pairs of states, with the β state a noticeable exception. In almost all cases the found free-energy differences closely resemble the US baseline with β as notable exception.

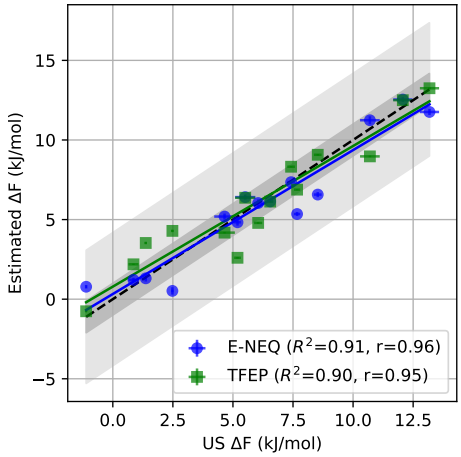

Figure 3: Correlation of estimated free-energy differences from E-NEQ and TFEP against the Umbrella Sampling MBAR baseline.

A further study of the correlation between the estimated free-energy differences and the true (as reported by US) free-energy differences similarly shows that the E-NEQ estimator outperforms TFEP. As highlighted in Fig. 3, both E-NEQ and TFEP show strong correlations, with R^2 and Pearson correlation coefficients of 0.91 and 0.96 for E-NEQ and 0.90 and 0.95 for TFEP. Both methods are firmly within 1 kcal/mol of the true free-energy differences, and in most cases even within 1 kJ/mol. This is well within the generally considered acceptable tolerances for free-energy estimation.

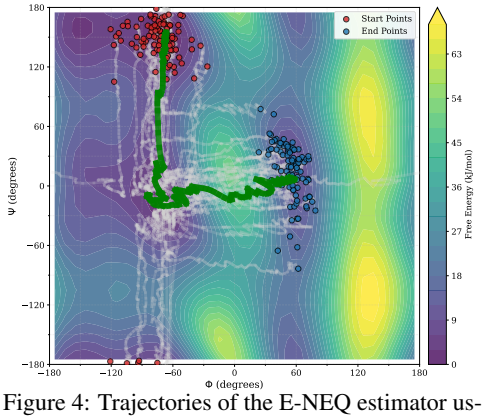

Figure 4: Trajectories of the E-NEQ estimator using a concatenated protocol starting from the β state towards the α_L state using the underdamped Langevin integrator.

Concatenating Protocols improves Multistate Free-Energy Estimation Comparing the results in table 1 between using pairwise consistent estimation between all states using only directly connected states (EPFG=N) and using a EPFG with concatenated protocols (EPFG=Y), we find that the latter is in almost all cases more accurate. Notably, it is not only the MAE that improves, but also the estimated free-energy differences between the states directly connected to the reference α_R state, despite using the same trajectories.

In Fig. 4 we have visualized E-NEQ trajectories starting from the β state towards the α_L state following a concatenated protocol. The trajectories clearly show the concatenation of the characteristically straight paths of OT-CFM starting from the β and visiting the α_R state before finally transitioning to the α_L state.

Table 2: Number of successful trajectories starting in each state for the different integration schemes. The maximum number of trajectories is 50,000, 10,000 to each other state.

f_t^U	dt	acc. rate	α_R	α_L	α_D	β	$C5$	α'
Underdamped Langevin	1e-4	0.94	49459	48871	48383	49183	48328	49419
Overdamped Langevin	1e-8	0.91	49535	49420	49265	49414	49172	49539
Hamiltonian Monte Carlo	1e-4	0.96	44056	46892	38892	40145	32477	47582

Stationary Preserving Dynamics comparison Lastly, we observe that, within the E-NEQ estimator with learned dynamics, underdamped and overdamped Langevin seem to be more stable than Hamiltonian Monte Carlo. Specifically, we found that over longer transitions, as occurs during the concatenation of protocols for multistate free-energy estimation, the HMC approach was more likely to diverge, resulting in discarded samples and ultimately a higher finite sample bias. The discrepancy between the number of successful trajectories starting in each state between the different integrators is reported in Tab. 2. MH correction does not help in this case due to the instability primarily caused by the deterministic vector field putting the samples in high-energy regions of the learned potential. As exemplified by the extremely small step size for overdamped Langevin dynamics, our proposed Lie-Trotter splitting of the work calculation is a necessary step to make the computation feasible within an acceptable compute budget.

6 Discussion

In this work we have proposed a method to learn the escorting switching protocol $(U^{\theta}_{A\to B}, b^{\phi})$ to construct an E-NEQ estimator with minimal variance. For this purpose we considered the framework of Conditional Flow Matching to learn b^{ϕ} and proposed an extension named Conditional Density Matching to learn $U^{\theta}_{A\to B}$. Furthermore, we considered two practical considerations in the form of Lie-Trotter splitting to reduce the computational cost of the work calculation and a flow graph construction to reduce the combinatorial complexity of learning multiple protocols in the multi-state setting. In our experimental evaluation using the ADP system, which has six well-defined metastable states, we have shown our proposed method for learning the escorted protocols to be effective.

Limitations and future work While ADP is a fitting benchmark for our method due to its multiple metastable states, similar challenges to larger systems, and generally a complexity level similar in size to studied in other related work (Rizzi et al., 2021; 2023; Máté et al., 2024; Erdogan et al., 2024; He et al., 2025), it is still a toy problem within the context of practical applications. We therefore emphasize the need to focus on scalability in future work. Preliminary experiments have shown that to achieve this emphasis should be placed on accurately learning the correct time-dependent potential $U^{\theta}_{A\to B}$. For this purpose we believe that recent advances in scaling discrete time normalising flows is an important avenue to explore (Rehman et al., 2025; Zhai et al., 2024).

In addition to this, we would like to note that in this work we choose to align the proposed method as closely as possible to the presented theory in Sec. 2 to provide a well grounded first exploration of neural escorted free-energy estimation. As such, we believe that by lifting some of these strict constraints performance could be improved. For example, removing the Metropolis-Hasting step that is used to ensure that stationary is preserved under discretized dynamics, and using adaptive ODE solvers instead of Runga-Kutta to set the lie-trotter splitting could therefore be beneficial. Both of these are common practice for standard molecular dynamic simulation. Furthermore, extensions of the current framework through hard constraints for enforcing the continuity equation instead of just encoding it in the joint CFM and CDM learning objectives, or incorporating the EPFG already during the training process could be beneficial. Lastly, we note that while we consider Flow Matching as the base for our E-NEQ estimation other generative modelling approachs in the form of Schrodinger Bridge marching and recent advances in Flow-Maps should also be considered in future work.

References

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland,

Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, pages 1–3, May 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07487-w.

Tara Akhound-Sadegh, Jarrid Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, and Alexander Tong. Iterated Denoising Energy Matching for Sampling from Boltzmann Densities, February 2024.

Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, January 2000. ISSN 0945-3245. doi: 10.1007/s002110050002.

Charles H Bennett. Efficient estimation of free energy differences from monte carlo data. *Journal of Computational Physics*, 22(2):245–268, 1976.

Peter G. Bolhuis, David Chandler, Christoph Dellago, and Phillip L. Geissler. Transition Path Sampling: Throwing Ropes Over Rough Mountain Passes, in the Dark. *Annual Review of Physical Chemistry*, 53(1):291–318, October 2002. ISSN 0066-426X, 1545-1593. doi: 10.1146/annurev.physchem.53.082301.113146.

Avishek Joey Bose, Tara Akhound-Sadegh, Guillaume Huguet, Kilian Fatras, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael Bronstein, and Alexander Tong. SE(3)-Stochastic Flow Matching for Protein Backbone Generation, April 2024.

Ricky T. Q. Chen and Yaron Lipman. Flow Matching on General Geometries, February 2024.

Zoe Cournia and Christophe Chipot. Applications of Free-Energy Calculations to Biomolecular Processes. A Collection. *Journal of Chemical Information and Modeling*, 64(7):2129–2131, April 2024. ISSN 1549-9596. doi: 10.1021/acs.jcim.4c00349.

Gavin E. Crooks. Nonequilibrium Measurements of Free Energy Differences for Microscopically Reversible Markovian Systems. *Journal of Statistical Physics*, 90(5):1481–1487, March 1998. ISSN 1572-9613. doi: 10.1023/A:1023208217925.

Ege Erdogan, Radoslav Ralev, Mika Rebensburg, Céline Marquet, Leon Klein, and Hannes Stark. FreeFlow: Latent Flow Matching for Free Energy Difference Estimation. October 2024.

Henry Eyring. The Activated Complex in Chemical Reactions. *Journal of Chemical Physics*, 3:107–115, February 1935. ISSN 0021-9606. doi: 10.1063/1.1749604.

Daan Frenkel and Berend Smit. *Understanding Molecular Simulation: From Algorithms to Applications*.

Philipp Geiger and Christoph Dellago. Optimum protocol for fast-switching free-energy calculations. *Physical Review E*, 81(2):021127, February 2010. ISSN 1539-3755, 1550-2376. doi: 10.1103/PhysRevE.81.021127.

Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science*, 4(2):268–276, February 2018. ISSN 2374-7943. doi: 10.1021/acscentsci.7b00572.

Jeff Gore, Felix Ritort, and Carlos Bustamante. Bias and error in estimates of equilibrium free-energy differences from nonequilibrium measurements. *Proceedings of the National Academy of Sciences*, 100(22):12564–12569, October 2003. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1635159100.

Jiajun He, Yuanqi Du, Francisco Vargas, Yuanqing Wang, Carla P Gomes, José Miguel Hernández-Lobato, and Eric Vanden-Eijnden. Feat: Free energy estimators with adaptive transport. *arXiv preprint arXiv:2504.11516*, 2025.

Graeme Henkelman, Blas P. Uberuaga, and Hannes Jónsson. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *The Journal of Chemical Physics*, 113(22):9901–9904, December 2000. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.1329672.

Lars Holdijk, Yuanqi Du, Ferry Hooft, Priyank Jaini, Berend Ensing, and Max Welling. Stochastic optimal control for collective variable free sampling of molecular transition paths. *Advances in Neural Information Processing Systems*, 36:79540–79556, 2023.

C. Jarzynski. Nonequilibrium Equality for Free Energy Differences. *Physical Review Letters*, 78(14): 2690–2693, April 1997. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.78.2690.

C. Jarzynski. Targeted free energy perturbation. *Physical Review E*, 65(4):046122, April 2002. ISSN 1063-651X, 1095-3787. doi: 10.1103/PhysRevE.65.046122.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, August 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03819-2.

John G. Kirkwood. Statistical Mechanics of Fluid Mixtures. *The Journal of Chemical Physics*, 3(5): 300–313, May 1935. ISSN 0021-9606. doi: 10.1063/1.1749657.

Benedict Leimkuhler and Charles Matthews. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Applied Mathematics Research eXpress*, 2013(1):34–56, January 2013. ISSN 1687-1200. doi: 10.1093/amrx/abs010.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching for Generative Modeling, February 2023.

Bálint Máté, François Fleuret, and Tristan Bereau. Neural Thermodynamic Integration: Free Energies from Energy-based Diffusion Models, June 2024.

Bálint Máté, François Fleuret, and Tristan Bereau. Solvation free energies from neural thermodynamic integration. *The Journal of Chemical Physics*, 162(12), 2025.

Amil Merchant, Simon Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk. Scaling deep learning for materials discovery. *Nature*, 624(7990):80–85, December 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06735-9.

David L Mobley and Michael K Gilson. Predicting Binding Free Energies: Frontiers and Benchmarks. 2017.

Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.

Danyal Rehman, Oscar Davis, Jiarui Lu, Jian Tang, Michael Bronstein, Yoshua Bengio, Alexander Tong, and Avishek Joey Bose. Fort: Forward-only regression training of normalizing flows. *arXiv preprint arXiv:2506.01158*, 2025.

Andrea Rizzi, Paolo Carloni, and Michele Parrinello. Targeted Free Energy Perturbation Revisited: Accurate Free Energies from Mapped Reference Potentials. *The Journal of Physical Chemistry Letters*, 12(39):9449–9454, October 2021. ISSN 1948-7185, 1948-7185. doi: 10.1021/acs.jpclett.1c02135.

Andrea Rizzi, Paolo Carloni, and Michele Parrinello. Multimap targeted free energy estimation. *Proceedings of the National Academy of Sciences*, 120(46):e2304308120, November 2023. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2304308120.

Victor Garcia Satorras, Emiel Hoogeboom, Fabian B. Fuchs, Ingmar Posner, and Max Welling. E(n) Equivariant Normalizing Flows. *arXiv:2105.09016 [physics, stat]*, June 2021.

Michael R. Shirts and John D. Chodera. Statistically optimal analysis of samples from multiple equilibrium states. *The Journal of Chemical Physics*, 129(12):124105, September 2008. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.2978177.

William C. Swope, Hans C. Andersen, Peter H. Berens, and Kent R. Wilson. A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters. *The Journal of Chemical Physics*, 76(1):637–649, January 1982. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.442716.

Charlie B. Tan, Avishek Joey Bose, Chen Lin, Leon Klein, Michael M. Bronstein, and Alexander Tong. Scalable Equilibrium Sampling with Sequential Boltzmann Generators, February 2025.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport, March 2024.

G.M. Torrie and J.P. Valleau. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *Journal of Computational Physics*, 23(2):187–199, February 1977. ISSN 00219991. doi: 10.1016/0021-9991(77)90121-8.

Hale F Trotter. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.

Donald G. Truhlar, Bruce C. Garrett, and Stephen J. Klippenstein. Current Status of Transition-State Theory. *The Journal of Physical Chemistry*, 100(31):12771–12800, January 1996. ISSN 0022-3654. doi: 10.1021/jp953748q.

Suriyanarayanan Vaikuntanathan and Christopher Jarzynski. Escorted Free Energy Simulations: Improving Convergence by Reducing Dissipation. *Physical Review Letters*, 100(19):190601, May 2008. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.100.190601.

Suriyanarayanan Vaikuntanathan and Christopher Jarzynski. Dissipation and lag in irreversible processes. *EPL (Europhysics Letters)*, 87(6):60005, September 2009. ISSN 0295-5075, 1286-4854. doi: 10.1209/0295-5075/87/60005.

Suriyanarayanan Vaikuntanathan and Christopher Jarzynski. Escorted free energy simulations. *The Journal of Chemical Physics*, 134(5):054107, February 2011. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.3544679.

Jiří Vymětal and Jiří Vondrášek. Metadynamics As a Tool for Mapping the Conformational and Free-Energy Space of Peptides — The Alanine Dipeptide Case Study. *The Journal of Physical Chemistry B*, 114(16):5632–5642, April 2010. ISSN 1520-6106. doi: 10.1021/jp100950w.

Fangping Wan, Daphne Kontogiorgos-Heintz, and Cesar de la Fuente-Nunez. Deep generative models for peptide design. *Digital Discovery*, 1(3):195–208, June 2022. ISSN 2635-098X. doi: 10.1039/D1DD00024A.

Lingle Wang, Yujie Wu, Yuqing Deng, Byungchan Kim, Levi Pierce, Goran Krilov, Dmitry Lupyan, Shaughnessy Robinson, Markus K. Dahlgren, Jeremy Greenwood, Donna L. Romero, Craig Masse, Jennifer L. Knight, Thomas Steinbrecher, Thijs Beuming, Wolfgang Damm, Ed Harder, Woody Sherman, Mark Brewer, Ron Wester, Mark Murcko, Leah Frye, Ramy Farid, Teng Lin, David L. Mobley, William L. Jorgensen, Bruce J. Berne, Richard A. Friesner, and Robert Abel. Accurate and Reliable Prediction of Relative Ligand Binding Potency in Prospective Drug Discovery by Way of a Modern Free-Energy Calculation Protocol and Force Field. *Journal of the American Chemical Society*, 137(7):2695–2703, February 2015. ISSN 0002-7863. doi: 10.1021/ja512751q.

Peter Wirnsberger, Andrew J. Ballard, George Papamakarios, Stuart Abercrombie, Sébastien Racanière, Alexander Pritzel, Danilo Jimenez Rezende, and Charles Blundell. Targeted free energy estimation via learned mappings. *The Journal of Chemical Physics*, 153(14):144112, October 2020. ISSN 0021-9606, 1089-7690. doi: 10.1063/5.0018903.

Shuangfei Zhai, Ruixiang Zhang, Preetum Nakkiran, David Berthelot, Jiatao Gu, Huangjie Zheng, Tianrong Chen, Miguel Angel Bautista, Navdeep Jaitly, and Josh Susskind. Normalizing flows are capable generative models. *arXiv preprint arXiv:2412.06329*, 2024.

Lu Zhao and Lei Wang. Bounding free energy difference with flow matching. *Chinese Physics Letters*, 40(12):120201, December 2023. ISSN 0256-307X, 1741-3540. doi: 10.1088/0256-307X/40/12/120201.

Adrianne Zhong and Michael R. DeWeese. Beyond Linear Response: Equivalence between Thermodynamic Geometry and Optimal Transport, April 2024.

Adrianne Zhong, Benjamin Kuznets-Speck, and Michael R. DeWeese. Time-Asymmetric Fluctuation Theorem and Efficient Free Energy Estimation, December 2023.

Robert W Zwanzig. High-temperature equation of state by a perturbation method. ii. polar gases. *The Journal of Chemical Physics*, 23(10):1915–1922, 1955.

A Appendix

A.1 Bennett Acceptance Ratio (BAR) and Multistate Bennett Acceptance Ratio (MBAR)

In this section we briefly introduce the Bennett Acceptance Ratio (BAR) and Multistate Bennett Acceptance Ratio (MBAR) estimators as introduced in the main text. First, the BAR estimator can be interpreted as a maximum likelihood solution for the logistic model that is implied by the Crooks Fluctuation Theorem (CFT). Given a collection of forward $\mathbf{x}_i^F$ and backward $\mathbf{x}_i^B$ trajectories and their corresponding work values $W_{A\to B}(\mathbf{x}_i^F)$ and $W_{B\to A}(\mathbf{x}_i^B)$ respectively, BAR solves:

$$\sum_{i=1}^{N_A} \frac{1}{1+\exp\beta^{-1}(W_{A\to B}(\mathbf{x}_i^F)-\Delta F)} = \sum_{i=1}^{N_B} \frac{1}{1+\exp\beta^{-1}(-W_{B\to A}(\mathbf{x}_i^B)-\Delta F)} \tag{22}$$

which can be solved using standard root-finding algorithms. The approximate solution for ΔF found using this estimator is known to have minimal variance among estimators that use both forward and backward trajectories (Bennett, 1976).

Adjusted to the multistate setting, the MBAR estimate is given by the set of equations (Shirts and Chodera, 2008):

$$F_k = -\beta^{-1}\ln\left(\sum_{j=1}^{K}\sum_{n=1}^{N_j} \frac{\exp\left(-\beta W_{i\to j}(\mathbf{x}_n^{i\to j})\right)}{\sum_{l=1}^{K} N_l \exp\left(F_l - \beta W_{l\to j}(\mathbf{x}_n^{l\to j})\right)}\right) \tag{23}$$

where $W_{i\to j}(\mathbf{x}_n^{i\to j})$ is the work done on the n-th trajectory moving from state i to state j. Note that here we obtain absolute free-energies as F_l with respect to the reference state k. Thus, to obtain the free-energy differences between all pairs of states, we can use the following equation:

$$\Delta F_{i\to j} = F_j - F_i. \tag{24}$$

A.2 Full Algorithmic Overview

Below we briefly provide the full algorithmic overview of the both the two state setting for free-energy estimation using E-NEQ as well as the multistate setting. Code is available at: https://github.com/iclranon862/Anonymous-Repo

A.2.1 Two-State Setting

Given our Conditional Flow and Density Matching based method, we break down the algorithm to obtain the free-energy difference between two thermodynamic states into the following steps:

1. **Generating Training Samples**: Given the two thermodynamic states of interest generate the equilibrium distributions of each state using standard molecular dynamics simulations. This results in two sets of samples $\{\mathbf{x}_n^A\}_{n=1}^{N_A}$ and $\{\mathbf{x}_n^B\}_{n=1}^{N_B}$ from respectively state A and state B.
2. **Learning Escorted Dynamics**: Learn the parameterised vector field b^θ and corresponding time-dependent potential $U_{A\to B}^\theta$ using Conditional Flow Matching and Conditional Density Matching respectively.
3. **Running the Escorted Protocol**: Using the learned b^θ and $U_{A\to B}^\theta$, run fixed length non-equilibrium trajectories using the escorted dynamics in Eq. (9) and determine for each trajectory the work done on the system using Eq. (4). Apply the Jarzynski Equality to estimate the free-energy difference between the two states.
4. **Running the Reversed Protocol** [Optional, but highly recommended]: Run the reversed protocol as defined in Eq. (13) and determine the work done on the system using the same method as above. Use the BAR estimator to obtain the minimum-variance estimate of the free-energy difference between the two states using the collection of work values from both the forward and reversed protocol.

A.2.2 MULTI-STATE SETTING

Extending to the multi-state setting using the proposed escorted protocol flow graph, we can break down the algorithm to estimate the free-energy differences between a set of thermodynamic states into the following steps:

1. **Generating Training Samples** Given a set of K thermodynamic states of interest and their corresponding potential energies $\{U_k\}_{l=1}^{K}$, generate a dataset of N_k equilibrium samples $\{\mathbf{x}_n^k\}_{n=1}^{N_k}$ from each state k using standard molecular dynamics simulations.
2. **Learning $K-1$ Escorted Protocols** Select one of the thermodynamic states as the reference state and learn the $K-1$ escorted protocols using the method proposed in Sec. 3. Subsequently, construct the fully connected flow graph using the concatenated potentials and escorting vector fields as defined in Eq. (21).
3. **Running the Escorted Protocols** Use the $K(K-1)/2$ escorted protocols in combination with one of the following methods to estimate the free-energy differences between all pairs of states:
   - Run $K(K-1)/2$ fixed length non-equilibrium trajectories using the escorted dynamics to obtain an estimate of the free-energy differences between all pairs of states using the Jarzynski Equality.
   - Run $K(K-1)/2$ fixed length non-equilibrium trajectories using the escorted dynamics, and another $K(K-1)/2$ fixed length non-equilibrium trajectories using the reversed escorted dynamics to obtain a pair-wise estimate of the free-energy differences between all pairs of states using the BAR estimator.
   - Similar to above, but now use the MBAR estimator using the pairwise work values from the forward and reversed protocols to obtain a self-consistent estimate of the free-energy differences between all pairs of states.

In our experiments, we used the third option using the MBAR estimator.

A.2.3 PRACTICAL DETAILS

So far we have assumed each thermodynamic state k to have its own potential energy function U_k. In practice, this is only the case for a small number of free-energy estimation problems, such as solvation free-energy. For most other forms of free-energy, such as conformational free-energy which we consider in our experimental evaluation, the states have the same potential energy U but correspond to different regions of the phase-space. In this case, a state-restricted potential energy function U_i has to be defined as:

$$U_i(x) = \begin{cases} U(x), & x \in \Omega_i \\ +\infty, & x \notin \Omega_i \end{cases} \tag{25}$$

where Ω_i is the region of space that is considered to be part of thermodynamic state i. Restricting potential energies to specific regions of space is a common technique in Molecular Dynamics simulations (Torrie and Valleau, 1977).

Equivariances Molecular systems are in general invariant to rigid-body transformations, such as rotations and translations. Including these transformations in architecture design of b^θ and U^θ is thus desirable and can significantly improve the performance of the method. However, in the case of Flow Matching extra care needs to be taken to ensure that the target vector field v_t also takes these transformations into account. This is formalised in the framework of Riemannian Flow Matching (Chen and Lipman, 2024; Bose et al., 2024).

Thermodynamic Irrelevant Degrees of Freedom In addition to the reduction in complexity that can be achieved by considering the symmetry structure of the system, an additional inductive bias can be introduced by considering the degrees of freedom of the system that do not contribute to the free-energy difference between the states of interest. For example, in the case of Alanine Dipeptide (ADP) a number of Carbon atoms are each connected to a single other heavy atom and have their remaining valences satisfied by hydrogen atoms. While these Hydrogen atoms considerably fluctuate during the simulation and can cause significant spikes in the potential energy, they generally have a

uniform contribution to the free-energy across all the states of interest. As such, we do not have to include them in the representation of the system used for U^θ and b^θ.

Note that removing degrees of freedom that are considered thermodynamic irrelevant has to be done with care, as incorrect removal can have a detrimental effect on the performance of the method. For example, while the atoms in the ADP system attached to terminal Carbon atoms are considered thermodynamic irrelevant, the same is not true for the atoms attached to central heavy atoms.

A.3 Experimental Setup

A.3.1 Molecular Dynamics Simulation Details

The simulation was run at 300 K using a standard Langevin middle integrator with a BAOAB splitting scheme and a time step of 2 fs. Samples were saved every 1000 steps. The Amber 14-alndx force-field was used with the GBn2 implicit solvent model. The same force field and implicit solvent model were used as the U_A and U_B potentials in the Umbrella Sampling simulations."

A.3.2 Metastable State Definitions

For the Alanine Dipeptide (ADP) system, we use the following definition of the metastable states based on the ϕ and ψ dihedral angles:

$$\alpha_R = \{-120 \leq \phi \leq 0, -110 \leq \psi \leq 90\} \tag{26}$$

$$\beta = \{-120 \leq \phi \leq 0, 90 \leq \psi \leq 180\} \cup \{-120 \leq \phi \leq 0, -180 \leq \psi \leq -110\} \tag{27}$$

$$\begin{aligned} C5 = {} & \{-180 \leq \phi \leq -120, -180 \leq \psi \leq -110\} \cup \{-180 \leq \phi \leq -120, 90 \leq \psi \leq 180\} \\ & \cup \{120 \leq \phi \leq 180, -180 \leq \psi \leq -110\} \cup \{120 \leq \phi \leq 180, 90 \leq \psi \leq 180\} \end{aligned} \tag{28}$$

$$\alpha' = \{-180 \leq \phi \leq -120, -110 \leq \psi \leq 90\} \cup \{120 \leq \phi \leq 180, -110 \leq \psi \leq 90\} \tag{29}$$

$$\alpha_L = \{0 \leq \phi \leq 120, -90 \leq \psi \leq 90\} \tag{30}$$

$$\alpha_D = \{0 \leq \phi \leq 120, -180 \leq \psi \leq -90\} \cup \{0 \leq \phi \leq 120, 90 \leq \psi \leq 180\} \tag{31}$$

adapted from (Vymětal and Vondrášek, 2010).

We have visualised the 5 metastable states in figure 2.