

293 A Background

294 In this section, we describe Generative Flow Networks in detail (Section A.2), and we perform a
295 survey of related work that tackles the catalyst discovery problem (Section A.3).

296 A.1 Catalysts for hydrogen energy storage

297 Hydrogen Energy Storage (HES) is a scalable, sustainable option for storing excess energy for
298 later use. Although HES has a relatively low efficiency compared to alternatives such as pumped
299 storage hydroelectricity and battery storage [42, 31, 6], it is more scalable and transportable than
300 the alternatives. HES via water electrolysis (water-splitting) can be thought of as two half-cell
301 reactions: the hydrogen evolution reaction and the oxygen evolution reaction. Both HER and OER
302 are electrochemical reactions which occur in the catalyzer at different electrodes. While HER is a one-
303 step reaction, OER consists of four intermediate reactions which, under an applied voltage, produces
304 oxygen gas, protons and electrons under acidic conditions. HER, in this setting, brings electrons and
305 protons together to produce hydrogen gas. As the energy is used for water electrolysis, the hydrogen
306 gas can be collected and stored in tanks, or other storage containers. When the energy is in demand
307 again, fuel cells can extract the energy from the stored hydrogen at about 60% efficiency [14]. An
308 efficient catalyst, used in the fuel cell or in the electrocatalyzer, can significantly speed up the reaction
309 rate and improve the efficiency of the process [42].

310 A.2 Generative Flow Networks

311 A generative flow network (GFlowNets, GFN) is an inference method introduced by [4] that allows
312 one to sequentially construct objects with probability proportional to a non-negative reward function.
313 They have been employed in tasks such as molecule generation, crystal generation, and for causal
314 discovery [22, 1, 3]. A desirable capability of GFlowNets is their ability to generate a diversity of
315 objects that all have high reward (relative to other possible candidates in the space). Within catalyst
316 discovery, this is key, as the reward function may not capture all the relevant predictors of whether a
317 material will be a good catalyst. For example, the reward function may capture a target adsorption
318 energy for the catalyst, but not its stability, which is an important trait (high stability) if the catalyst is
319 to be used in real-world applications. After generating a variety of high-reward catalyst materials, one
320 or more of these materials may also be found to be stable during experimental testing. Furthermore,
321 GFlowNets perform particularly well in cases where the search space is quite large, and outperform
322 methods such as MCMC and reinforcement learning in terms of mode mixing [5].

323 A key component of GFlowNets is the proxy model, which is how the GFlowNet obtains information
324 on the reward landscape, and hence what “desirable”, high-reward samples are. Proxy models in
325 our context are predictive models (possibly machine learning models) that allow one to evaluate the
326 reward of a particular sampled object. In the case of molecule or crystal discovery, it may be that
327 the sampled molecule or crystal does not exist in any database, and as such, relevant properties of
328 the crystal must be computed in real time instead. Here, we note that for certain properties such
329 as the formation energy of a crystal, physics simulations such as density functional theory (DFT)
330 are traditionally performed. These simulations are computationally expensive and not scalable to
331 high batch sizes. Hence, a machine learning model that is able to replace the physics simulation
332 and compute the properties of interest is valuable; for crystals, there are many works that develop
333 machine learning-based approaches to compute properties such as formation energy and adsorption
334 energy [35, 36, 13, 9].

335 Crystal GFlowNet is a GFlowNet designed to discover new materials by sequentially sampling crystal
336 structures [1]. This work uses the GFlowNet paradigm to construct and sample crystal structures that
337 have a high reward. The authors test the method on a reward based upon formation energy, where a
338 high reward corresponds to a low formation energy, and a low reward corresponds to a high formation
339 energy. This allows the framework to be used to sample thermodynamically stable materials with
340 desirable properties. As mentioned earlier, the reward function may use any measurable (or predicted,
341 via a proxy model) characteristic of the material. Sequential crystal construction proceeds in three
342 steps. First, the space group (symmetry operations) of the crystal are selected. Then, the composition
343 (elements and their ratios), and finally, the lattice parameters, which determine the shape of the final
344 lattice. Atom positions are not sampled by the model.

A.3 Machine learning for catalyst discovery

Since traditional catalyst design methods are time-consuming and involve trial and error experiments within a very large search space [18], scientists are turning to machine learning approaches to more efficiently explore the space of good catalysts. Recent work has leveraged both predictive models (which evaluate the performance of catalysts based on their measurable attributes) and generative models (which aim to construct catalyst representations with desirable properties) for catalyst discovery.

One major thread of predictive machine learning for catalyst design are predictive models to replace density functional theory (DFT) calculations of adsorption energies. The Open Catalyst Dataset 2020 [8] and 2022 [37] are datasets of 1.2 million and 62 thousand relaxations respectively. The datasets span a wide variety of materials, surfaces and adsorbates, with the inclusion of oxides in OC22. Machine learning models such as FAENet, Schnet, PaiNN, and Graphormer [13] are neural networks (e.g. graph neural networks, transformers) that can be trained to perform the OC20 Initial Structure to Relaxed Energy (IS2RE) task, which involves predicting the relaxed energy of an adsorbate-catalyst system from their initial atomic positions. Machine learning models have also been used to predict formation energies, perform relaxations, predict forces [9]. Predictive machine learning models have also been used to predict the thermodynamic stability of catalysts [44]. All of these predictive models have the potential to be included as proxy models within generative approaches.

Recent generative approaches to catalyst discovery include reinforcement learning [23], large language models [24], variational autoencoders [34], and generative adversarial networks (GANs, citation) which aim to generate catalysts *de novo*, by meeting criteria such as target adsorption energies and encoding this in the reward or loss. AdsorbRL [23] employs Deep Q-learning to find compositions of catalysts that have either minimal or maximal adsorption energy. They only model composition, omitting space groups, crystal symmetries, surface selection, and atom positions. [34] addresses the problem by employing a VAE to generate catalysts by sampling from the latent space, but we note that reliance on the latent space makes controllable generation and systematic exploration of possible catalysts difficult. [20] takes the approach of GANs to generate catalysts for the ammonia formation reaction, but does not take catalyst stability into account during generation. LLM-based tools such as [24, 30] have so far been mainly applied to help screen known catalysts and perform a literature search.

376 B Effect of relaxation on lattice parameters

377 After the GFlowNet generates samples, we relax them using M3GNet PES [9], as detailed in
 378 Section 2.1. Figure 3 demonstrates the change in lattice parameters of the cubic platinum lattice
 379 before and after relaxation. We note that the lattice parameters relax to the minimum of the total
 380 energy of the system.




Figure 3: For the HER case study, the platinum samples generated by the GFlowNet relax to the minimum energy every time. The green line represents the total energy of the system.

381 Table 1 details the search space for the Hydrogen evolution reaction case study (Section 2.2). We
 382 note that in general, the Catalyst GFlowNet’s constraints and search space can be set to be much
 383 larger than that used in the HER case study. In future work, we will consider a larger search space up
 384 to 80-100 atoms, more elements and space group options, while enforcing a neutral charge constraint.

Table 1: Hydrogen evolution reaction catalyst search space.

Setting	Value	Units
Possible elements	Pt, Ag, Au, Pd, Ir, Ni, W, Co, Cu, Mo, Rh, Nb	
Min. # of different elements in unit cell	1	
Max. # of different elements in unit cell	1	
Min. atoms in unit cell	2	
Max. atoms in unit cell	4	
Min. atoms per element	2	
Max. atoms per element	4	
Enforce neutral charge?	No	
Possible space groups	225, 229	
Min. lattice parameter length	2	Angstroms
Max. lattice parameter length	6	Angstroms
Min. angle between lattice vectors	60	Degrees
Max. angle between lattice vectors	140	Degrees

385 **C Table of hydrogen evolution reaction samples**

Table 2: Proportions of sampled structures for the case study: hydrogen evolution reaction.

Composition	Space Group	Overpotential (eV)			Samples	
		Experimental ²	DFT	Proxy Model	Count	Percentage
Pt	225	0.016	0.03	0.04	185	43.53
Rh	225	0.071*	0.06	0.03	144	33.88
Pd	229	0.036*	0.09*	0.12	54	12.71
Co	225	0.224*	0.25	0.16	24	5.65
Ir	225	0.015	0.08	0.14	10	2.35
Mo	229	0.394*	0.53	0.21	5	1.18
Cu	225	0.445	0.27	0.39	3	0.71
W	229	0.318*	0.59	0.30	0	0.00
Nb	229	0.488*	0.56	0.35	0	0.00
Ni	229	0.268	0.23	0.32	0	0.00
Au	225	0.337	0.63	0.51	0	0.00
Ag	225	0.403	0.58	0.62	0	0.00
Sum					425	100.00




Figure 4: A linear fit that maps log of exchange current density to experimental overpotential using data from [38] and [11]. These values are used column 3 (from the left) of Table 2.

²When available, the experimental overpotential was added from [11]. For the remaining elements, marked with an asterisk, the overpotentials were inferred using a linear fit between the log of experimental exchange current density (from [38]) and the available experimental overpotentials (from [11]). Figure 4 shows the linear fit and predictions.