
Discovery of Sustainable Refrigerants through Physics-Informed RL Fine-Tuning of Sequence Models

Adrien Goldszal^{1,2,3†*} Diego Calanzone^{2,3*} Vincent Taboga^{2,3} Pierre-Luc Bacon^{2,3}

¹École Polytechnique ²Mila Quebec AI Institute ³Université de Montréal

Abstract

Most refrigerants currently used in air-conditioning systems, such as hydrofluorocarbons, are potent greenhouse gases and are being phased down. Large-scale molecular screening has been applied to the search for alternatives, but in practice only about 300 refrigerants are known, and only a few additional candidates have been suggested without experimental validation. This scarcity of reliable data limits the effectiveness of purely data-driven methods. We present *Refgen*, a generative pipeline that integrates machine learning with physics-grounded inductive biases. Alongside fine-tuning for valid molecular generation, *Refgen* incorporates predictive models for critical properties, equations of state, thermochemical polynomials, and full vapor compression cycle simulations. These models enable reinforcement learning fine-tuning under thermodynamic constraints, enforcing consistency and guiding discovery toward molecules that balance efficiency, safety, and environmental impact. By embedding physics into the learning process, *Refgen* leverages scarce data effectively and enables *de novo* refrigerant discovery beyond the known set of compounds.

1 Introduction

The search for new refrigerants involves balancing often competing objectives: achieving optimal thermodynamic properties, fluid behavior, and heat transfer efficiency while meeting environmental and safety requirements such as low Global Warming Potential (GWP), reduced flammability, and minimal toxicity. The landscape of refrigerant compounds has undergone major shifts in recent decades with the replacement of ozone depleting chlorofluorocarbons (CFC's) with hydrofluorocarbons (HFC's) through the Montréal Protocol [United Nations Environment Programme Ozone Secretariat, 1987]; though, as many HFC's exhibit high GWP, recent efforts such as the Kigali Amendment [kig, 2016] are organizing their phase down. Hydrofluoroolefins (HFO's) are emerging as potential viable candidates, but long term environmental impact studies suggest that most heavily fluorinated compounds such as per- and polyfluoroalkyl substances (PFAS), which include many HFO's, should ideally be eliminated in favor of alternatives [OECD, 2021, Glüge et al., 2024]. While increasing interest is therefore taken in natural refrigerant options such as CO₂ and propane, they present challenges of their own, with propane being highly flammable and CO₂ operating at high pressures.

Traditional *high-throughput screening* has long been a staple in drug and materials discovery, where large molecular databases are filtered down to a small number of promising candidates through property prediction and expert knowledge before experimental validation. In the refrigerant domain, however, such systematic screening efforts have been rare. One of the few large-scale studies of this kind [McLinden et al., 2017] highlighted how limited the chemical space becomes once thermodynamic performance, safety, and environmental constraints are all enforced, identifying

Code is available at <https://github.com/ddidacus/refgen>.

only a handful of plausible candidates. Moreover, accurate evaluation of molecular properties for this filtering process remains difficult due to the scarcity of publicly available ground-truth data. This motivates the recent shift toward learning predictive relationships between physicochemical properties and molecular structure as a way to support refrigerant discovery [Kazakov et al., 2012].

In this work we explore a novel application of large molecule sequence models¹. We present Refgen, a framework for the discovery of new refrigerant candidates adapted to the challenges faced by the industry:

1. We develop state-of-the-art physics-grounded property predictors trained independently from supervised datasets and physics models (Equations of state (EOS), NASA polynomials, group-contribution methods) to compute key thermodynamic and chemical properties of molecules.
2. We build a multi-reward RL pipeline using the property predictors to guide our LLM towards the generation of optimal refrigerant candidates in SMILES format.

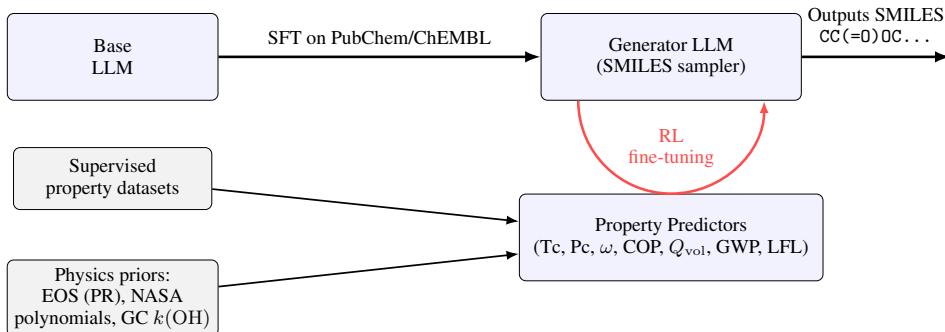


Figure 1: The Refgen framework. The LLM is supervised on molecular corpora for valid SMILES generation. During RL fine-tuning, the grouped predictor outputs form a multi-property reward that hooks into the LLM for policy updates. *Note* : an in-depth pipeline schematic can be found in Appendix A.

2 Methodology

2.1 Computing physics-grounded refrigerant properties

Many complex physical properties are hard to measure in practice. Examples include coefficient of performance in vapor compression cycles and global warming potential of refrigerants. To account for data scarcity, our approach consists of two steps: firstly, we train predictors on sub-properties for which annotations are more readily available; secondly, we use established physical models (e.g. equations of state for thermodynamic properties) to compute the final properties.

Coefficient of Performance and Volumetric Concentration A refrigeration cycle uses the phase changes of a working fluid to transfer heat. The thermodynamic behavior of the molecule in the different steps of the vapor compression cycle is essential as it defines the efficiency of the cycle, called the coefficient of performance (COP). The COP is calculated as the ratio of cooling effect to work input. Another important property which affects the system size is the volumetric cooling capacity (Q_{vol} , in $MJ.m^{-3}$), defined as the refrigeration effect per unit volume of refrigerant vapour entering the compressor. For any working pressure and temperature, COP and Q_{vol} are computed as:

$$COP = \frac{\text{Refrigerating Effect}}{\text{Compressor Work}} = \frac{h_1 - h_4}{h_2 - h_1} \quad Q_{vol} = \frac{h_1 - h_4}{v_1}$$

Where h_1 and v_1 are respectively the specific enthalpy and specific volume of the refrigerant at the compressor inlet (state 1), h_2 the specific enthalpy at compressor outlet (state 2) and h_4 the specific enthalpy of the refrigerant at the evaporator inlet (state 4).

¹A review of related work can be found in appendix B

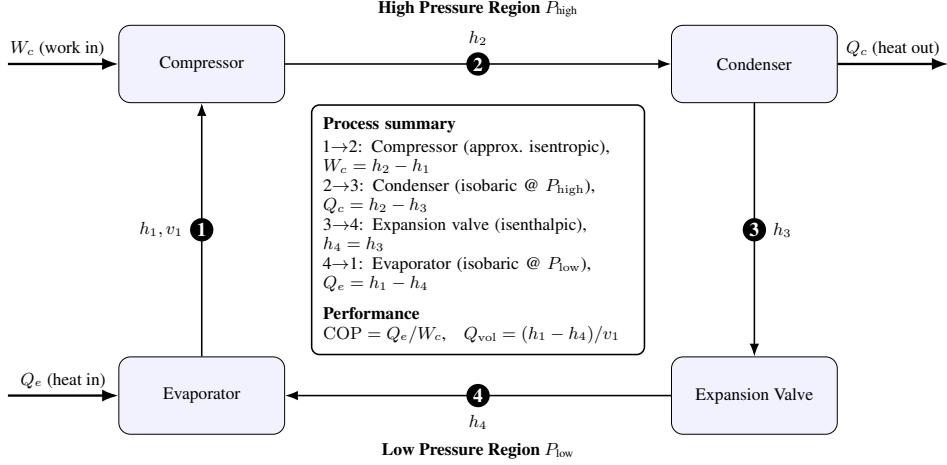


Figure 2: Vapor-compression cycle with four states. Each state is labeled with its specific enthalpy h_i , and state 1 also shows the specific volume v_1 , used in defining the volumetric refrigerating effect Q_{vol} . The coefficient of performance (COP) follows from enthalpy differences between states.

The refrigerant behavior can be modeled with equations of state (EOS) defining the relation between temperature, pressure and enthalpy. We use the modern Peng-Robinson EOS [Peng and Robinson, 1976], which incorporates repulsion and attraction terms, to model the behavior of gases and liquids under different conditions:

$$P = \frac{RT}{V_m - b} - \frac{a\alpha(T)}{V_m(V_m + b) + b(V_m - b)}.$$

where P is pressure, V_m is molar volume, T is temperature, and R is the ideal gas constant. The parameters a , b , and the temperature-dependent factor $\alpha(T)$ are expressed in terms of critical properties and the acentric factor, which needs to be given as input to the model:

$$b = 0.07780 \frac{RT_c}{P_c}, \quad a = 0.45724 \frac{R^2 T_c^2}{P_c}, \quad \alpha(T) = \left[1 + m \left(1 - \sqrt{T/T_c} \right) \right]^2,$$

with

$$m = 0.37464 + 1.54226 \omega - 0.26992 \omega^2.$$

To simulate the thermodynamic cycle of a molecule and COP and Q_{vol} calculations for any working pressures and temperatures, the relationships between real refrigerant enthalpy, entropy and other molecular properties such as pressure and temperature have to be computed. Ideal gas properties are first computed through NASA polynomials, which ensure proper thermodynamic relationships between heat capacity (C_p), enthalpy (H) and entropy (S) before using the EOS to obtain departure functions from ideal to real behaviour. Detailed explanations are in appendix C.

Global Warming Potential The second characteristic to consider is the environmental impact, predominantly measured by the global warming potential (GWP). The GWP₁₀₀ quantifies the cumulative radiative forcing impact of a greenhouse gas (GHG) relative to CO₂, over a 100-year time horizon :

$$\text{GWP}_{100} = \frac{\text{AGWP}_X(100)}{\text{AGWP}_{\text{CO}_2}(100)}$$

where AGWP_X(100) is the Absolute Global Warming Potential of species X, and AGWP_{CO₂}(100) = $9.0 \times 10^{-14} \text{ W m}^{-2} \text{ yr kg}^{-1}$ is the same metric for CO₂, used as a reference.

AGWP_X, as defined in the IPCC reports and detailed in Zieger et al. [2025] combines the atmospheric lifetime of a molecule and its radiative efficiency :

$$\text{AGWP}_X(100) = \text{RE}_X \times \tau_X \times \left(1 - e^{-100/\tau_X} \right)$$

Here, R_{EX} is the **radiative efficiency** of X ($\text{W m}^{-2} \text{kg}^{-1}$), i.e., the instantaneous radiative forcing (net change in radiative flux at the top of the atmosphere caused by the presence or increase of a greenhouse gas) per unit mass increase. τ_{X} is the **atmospheric lifetime** of X in years. The atmospheric lifetime τ is the mean residence time of a molecule before it is removed via chemical reactions or photolysis. As shown in Kazakov et al. [2012], the accuracy in estimating τ_{X} is important mainly for the cases when $\tau_{\text{X}} \ll 100$. For this range of lifetimes, reaction with OH appears to be the dominant loss mechanism. The corresponding lifetime is:

$$\tau_{\text{X}} = \frac{1}{k_{\text{OH}}[\text{OH}]}$$

where k_{OH} is the rate constant ($\text{cm}^3 \text{molecule}^{-1} \text{s}^{-1}$), and $[\text{OH}] \approx 1 \times 10^6 \text{ molecules cm}^{-3}$ [Kazakov et al., 2012].

Flammability Although flammability requirements are being reconsidered to help find new refrigerants, many safety codes like ASHRAE [2022] still require nonflammable refrigerants. The ASHRAE documentation uses the Lower Flammability Limit (LFL) defined as the minimum concentration capable of propagating a flame through a homogeneous mixture of refrigerant and air at standard conditions (23°C and 101.3 kPa), typically expressed in vol% or kg.m^{-3} .

2.2 Predicting refrigerant properties from SMILES

Building robust property predictors is crucial to guide the search of molecules during post-training RL towards desired tradeoffs. As annotated datasets of thermodynamic property-measurements are scarce, we leverage laws introduced in section 2.1 for more accurate scoring. Our property predictors, similarly to our generative backbone, are sequence models that encode *SMILES* (Simplified Molecular Input Line Entry System, Weininger [1988]): textual representations of molecules in ASCII characters, representing stereochemistry, cycles and branches. A SMILES string is computed by traversing the molecular graph, and is therefore not unique, enabling augmentation. We base our predictors on the *SMiles Transformer Encoder Decoder (SMI-TED)* [Soares et al., 2025], in order to leverage its pre-trained dense embeddings and to avoid handcrafted features.

2.3 RefGen generative model training

Our model is based on Llama 3.2 1B [Grattafiori et al., 2024], an open and highly capable LLM for instruction-conditioned tasks. We adopt the same tokenization scheme, Byte-Pair Encoding, which we found to work well with SMILES. We fine-tune (SFT) the pre-trained Llama model on unconditioned SMILES generation; re-using this LLM allows to inherit instruction-tuning and language understanding capabilities for downstream chemical reasoning tasks. In a subsequent post-training phase (RLFT), the model learns to generate molecular structures that more likely satisfy the acceptance ranges for thermodynamic properties.

Supervised Fine-Tuning (SFT) We fine-tune the model on standard causal language modeling on sequences $\mathbf{x} \in \mathcal{D}_{\text{SFT}}$, $\mathbf{x} = (x_1, x_2, \dots, x_T)$ where $x_t \in \mathcal{V}$ denotes tokens from the vocabulary \mathcal{V} and T is the sequence length. Here, \mathcal{D}_{SFT} is a combination of Pubchem [Kim et al., 2021], ChEMBL [Gaulton et al., 2017] and SureChEMBL [Papadatos et al., 2016] processed with our filtering pipeline, detailed in Appendix E, counting a total of $\sim 37\text{M}$ sequences ($\sim 1.5\text{B}$ tokens). A format for structured SMILES generation is adopted: sequences are delimited with XML-style tokens `<s>` or `<smiles>` and `</s>` or `</smiles>`, e.g. `<s>CC1(F)C(F)C1(F)F</s>`.

RL Finetuning (RLFT) We consider Group-Relative Policy optimization (GRPO) [Shao et al., 2024], which introduces a simplification of the PPO objective based on advantage estimation over multiple rollouts, given (multiple) callable reward functions. We can thus consider the property predictor networks from 2.2 as scoring functions, each depending on the defined acceptance region per thermodynamic property, e.g. number of atoms between 7 and 18, further details in Appendix G.

3 Experimental setup

3.1 Refrigerant property prediction

We compile datasets with annotated properties for SMILES structures: T_c , P_c and w are extracted from Bell et al. [2016–2024]; NASA polynomials for the enthalpy and entropy from Farina. et al. [2021]; radiative efficiencies from Muthiah et al. [2023]; $k(OH)$ reaction constants from McGillen et al. [2020]; LFL values are compiled from Maury et al. [2023] and Bell et al. [2016–2024]. We fine-tune SMI-TED end-to-end on our property-measurement datasets. We apply SMILES augmentation by canonicalizing the structures in the dataset and then by generating multiple graph traversals with RDKIT’s MolToSmiles functionality. Copies of the base SMI-TED model are fine-tuned for the various properties with train/valid/test splits. Additionally, we test our COP predictor on a CoolProp [Bell and contributors, 2025] test split, which also implements detailed EOS for many refrigerant compounds. Accuracy of the GWP predictor is estimated on data from U.S. Environmental Protection Agency, Office of Research and Development (ORD) [2023]. For the reaction rate with the OH radical $k(OH)$, we additionally compare performance with a group contribution method based on Kwok and Atkinson [1995] and open source notes from the EPA AOPWIN software [U.S. Environmental Protection Agency, 2012]. Further details in Appendix D and F.

3.2 Reinforcing optimal refrigerant molecules

To sample optimal refrigerant molecules wrt. constraints on thermodynamic properties, we fine-tune our model with GRPO on completions of $\langle s \rangle$ (unconditioned molecule generation). The reward signal consists of a linear combination of five scoring functions with relative weights obtained via hyperparameter search:

$$R_{\text{total}}(\mathbf{x}) = R_{\text{diversity}}(\mathbf{x}) \cdot \sum_{i=1}^k w_i R_i(\mathbf{x})$$

where w_i represents the relative reward weight, $R_i(\mathbf{x})$ is the reward from the i -th property predictor, and k is the number of objectives. We empirically observe the properties COP, Q_{vol} , T_c to be the most impactful for convergence, hence we allocate more weight. A diversity reward is considered as global scaling factor to ideally re-weight good molecules by novelty. Reward functions and training details are provided in Appendix G.

Reward	w_i
COP & Q_{vol}	0.40
T_c	0.40
Molecular length	0.10
GWP	0.05
LFL	0.05

Figure 3: Reward weights

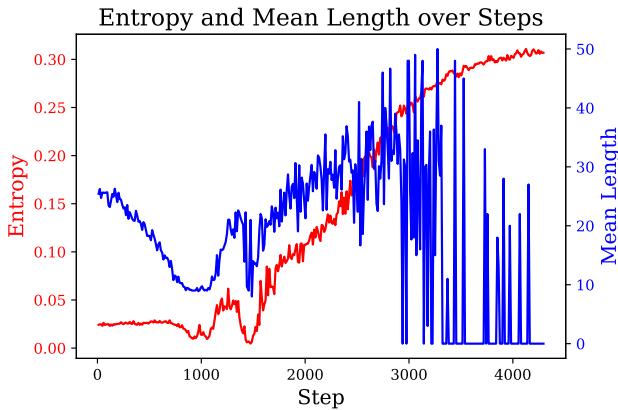


Figure 4: Mode collapse with entropy maximization.

To encourage diversity, we test two approaches: (1) introducing an entropy regularizer, in spirit of Eysenbach and Levine [2022]; (2) using a diversity and repetition reward signal as in Loeffler et al. [2024]. We find the latter to better prevent mode collapse (an instance is displayed in Figure 4) and to yield more stable training dynamics, as entropy maximization clashes with the idea of narrowing down search to a subset of optimal molecules, thus requiring additional techniques such as annealing.

4 Results

We evaluate two main components of the Refgen pipeline: we first assess the property predictor's coverage across the molecular space; and subsequently evaluate our generative model to quantify, within the distribution covered by our predictors, the quality of generated structures.

Property predictor performance We first ground the choice of the SMI-TED backbone for property prediction with a comparison to a gradient-boosting baseline with handcrafted features. Experimental results are reported in Table 4, Appendix D. We find the features learned with SMI-TED to be robust and SMILES augmentation proves beneficial, as redundant representations are mapped, leading to increased test accuracy overall. For the fine-tuned SMI-TED models, we observe low prediction errors in and out of distribution: for COP, we compute the saturation dome reconstruction accuracy by comparing the saturation enthalpy difference RMSE at sampled critical temperatures along the dome between the predicted cycle and CoolProp's very accurate EOS considered as ground-truth. We also compute COP mean absolute error for operating temperatures of 10°C and 40°C (Figure 5 and Table 10). Similarly, we report, for the GWP property, accurate distribution coverage, with the k(OH) group contribution method showcasing the best generalization ability; results for all the predictors are reported in Appendix H.

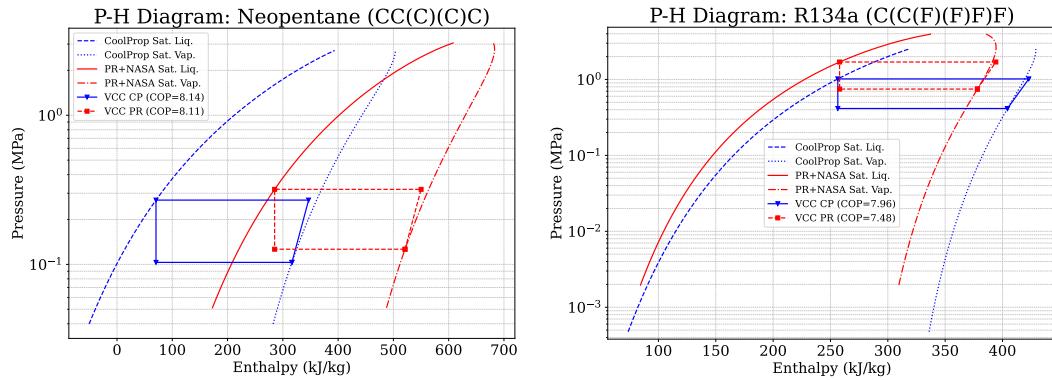


Figure 5: Comparison between ground-truth saturation dome from CoolProp (blue) and predicted (red) for molecules Neopentane (left) and R134a (right). The COP is computed for condenser and evaporator temperatures at 10°C and 40°C respectively. Dome translation shifts are due to varying reference points in the enthalpy and they are irrelevant for the computation of COP or Q_{vol} .

Table 1: Median property scores² for 1000 unique molecules sampled from the base (Refgen-SFT) and fine-tuned (Refgen-RLFT) models .

Setup	T_c [K]	N_{atoms}	GWP	LFL [kg/m ³]	COP	Q_{vol} [MJ/m ³]
Base Model	841	22	0.04	0.07	8.66	0.002
Finetuned Model	403	9	13	0.61	7.99	1.179

De Novo molecular generation We first illustrate the difference between the distribution of generated structures after SFT (base model with no specialization) and after RLFT (fine-tuned model with property optimization) (Table 1, Figure 6). Optimizing for thermodynamic properties significantly changes the target region towards small fluorinated compounds such as HFC's when optimizing for T_c and molecular length, or towards HFO's when adding GWP and flammability constraints. Additionally optimizing for Q_{vol} explores new regions of the chemical space presenting multiple novel SMILES not listed in McLinden et al. [2017].

²The COP comparison cannot be accurately made as only 22 molecules had properties in range for which a saturation dome and therefore a COP and Q_{vol} could be computed. Concerning the GWP, comparison is difficult as GWP predictions for these complex molecules is certainly out of distribution

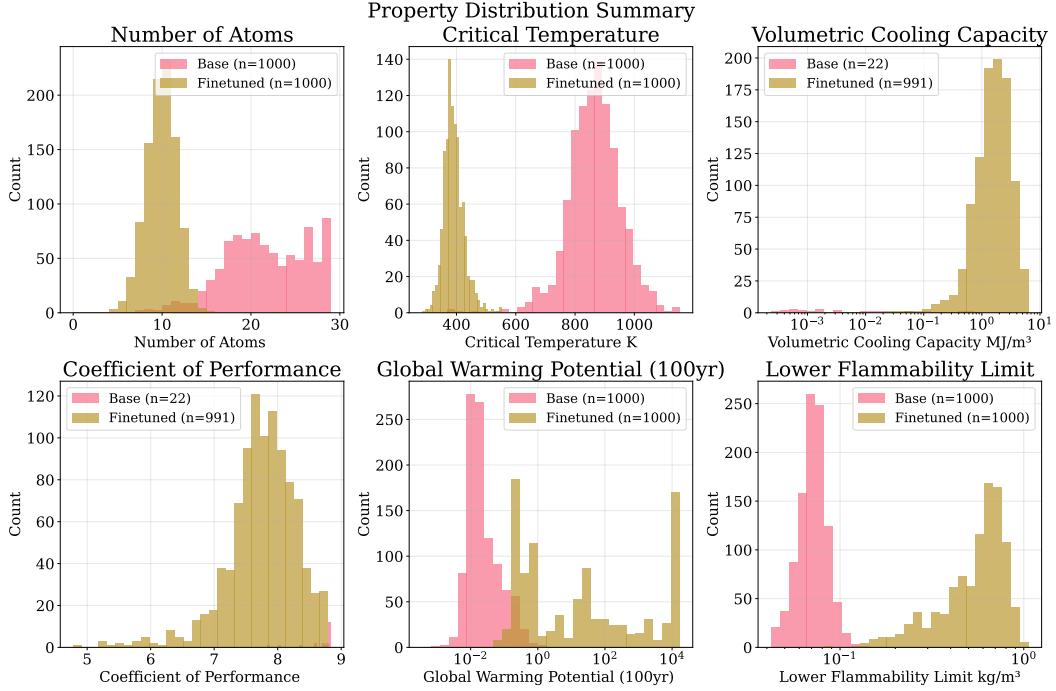


Figure 6: Comparison between distribution plots for our base (Refgen-SFT) and finetuned (Refgen-RLFT) models in terms of length, T_c , Q_{vol} , COP, GWP, LFL.

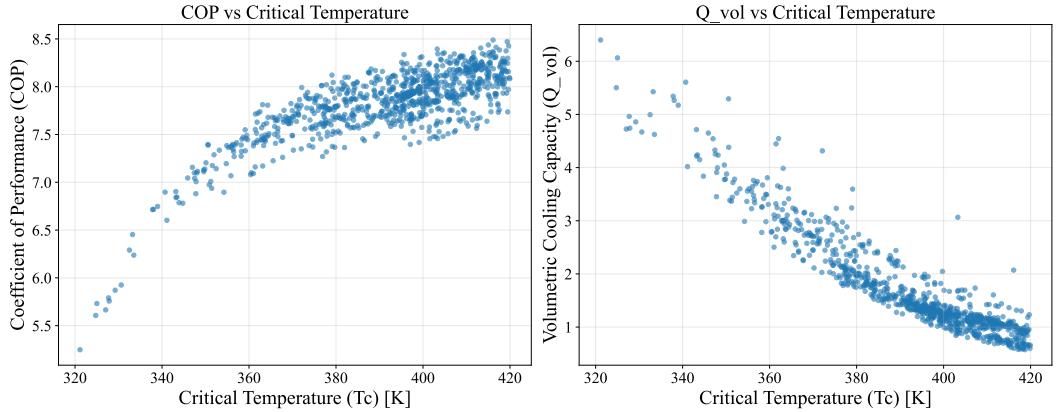


Figure 7: Comparative plots for COP to T_c and Q_{vol} to T_c for compounds generated with Refgen.

Finally, we consider a set of constraints to select optimal candidates based on principles from McLinden et al. [2017]: $COP > 5$; $320K < T_c < 420K$; $LFL > 0.1 kg.m^{-3}$; we also impose $GWP < 10$, which is lower than all existing regulations but still keeping some margin compared to $GWP = 1$ in CO_2 . We filter out $=CF2$ and $-OF$ groups according to stability and toxicity warnings detailed in McLinden et al. [2017]. Generated molecules reveal a tradeoff between COP and Q_{vol} depending on critical temperature, Figure 7: the higher the T_c , the higher the COP, but the lower the Q_{vol} . This behaviour, observed in the reference work of McLinden et al. [2017] confirms the importance of our physics-grounded approach, which ensures thermodynamically sound predictions. In Figure 8, we compare the COP to Q_{vol} tradeoff of the best molecule candidates with respect to R-410A ($COP = 7.39$, $Q_{vol} = 6.61$) as in the reference. R-410A is a heavily used refrigerant blend that is considered to have some of the best thermodynamic properties in the vapour compression cycle; we observe similar or competitive tradeoffs for our compounds, which additionally optimize for diversity and other properties introduced in our new pipeline.

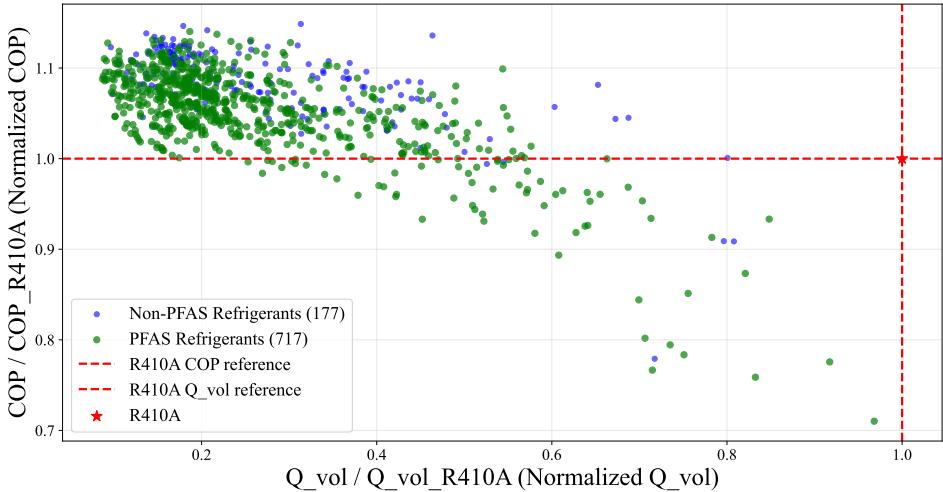


Figure 8: Generated molecules visualized by COP to Q_{vol} ratio, compared to R-410A reference.

Finally, we screen out per- and polyfluorinated substances (PFAS) according to the OECD definition [OECD, 2021] for potential environmental risk of toxicity. We observe $\sim 20\%$ of generated molecules to be non-PFAS, despite no constraint having been added.

SMILES	COP	T_c (K)	GWP100	LFL (kg/m ³)	Q_{vol}
C(F)N(F)N(F)F	6.71	337.82	3.95	0.52	5.34
N(F)C(F)N(F)F	6.72	338.05	0.25	0.54	5.26

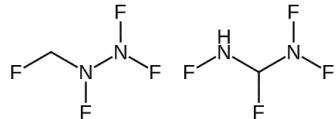


Figure 9: Example refrigerant candidates with properties and 2D structures, full list in Appendix I.

5 Discussion and Conclusion

We introduced `Refgen`, a novel generative pipeline for refrigerant molecules which successfully handles thermodynamic and environmental tradeoffs. Our model re-discovers HFCs and HFOs when appropriate constraints are imposed, while also proposing novel classes of candidates with the provided tradeoffs. Further work is currently focused on scaling up constraints, increasing generation diversity, and most importantly, verifying the validity of sampled candidates in a real lab.

To improve the overall quality of molecules, additional properties could be taken into account such as toxicity or molecular stability. Toxicity prediction highly differs in refrigerant discovery with respect to drug design. In the first case, long term exposures defined by workplace security panels usually define toxicity levels on an individual molecular basis with no clearly defined formula, making estimation very difficult. Molecular stability is equally relevant, as some of the generated candidates contain highly strained 3-atom cycles, which are known to be unstable.

As we sample more sequences from our generative model, we start to observe a saturation in the number of candidates satisfying our constraints, Table 2. This suggests potential limits in exploration and diversity. Further work includes investigating novel generative approaches and techniques to emphasize variety.

Table 2: Selected unique molecules by constraints satisfaction, Section 4.

$N_{generated}$	$N_{filtered}$
65 536	786
1 024 000	894

We aim at integrating our *in-silico* pipeline in the real development process, where candidates sampled from our generative model can be further validated in lab. On this line of research, we seek for experimentation in collaboration with chemists and HVAC experts to assess refrigerant synthesizability, stability, environmental impact and practical efficacy in vapor compression cycles.

References

- Amendment to the montreal protocol on substances that deplete the ozone layer (kigali amendment). Adopted at the Twenty-Eighth Meeting of the Parties to the Montreal Protocol, Kigali, 10–15 October 2016, October 2016. URL https://treaties.un.org/doc/Treaties/2016/10/20161015%2006-28%20PM/Ch_XXVII-2-f.pdf. United Nations Treaty Series, C.N.872.2016.TREATIES-XXVII.2.f; entry into force 1 January 2019 (with certain provisions entering into force 1 January 2033).
- ASHRAE. Designation and Safety Classification of Refrigerants, 2022. URL <https://www.ashrae.org/technical-resources/bookstore/standards-15-and-34>. Packaged with Standard 15-2022.
- Caleb Bell, Yoel Rene Cortes-Pena, and Contributors. Chemicals: Chemical properties component of chemical engineering design library (chedl). <https://github.com/CalebBell/chemicals>, 2016–2024. Accessed: November 18, 2025.
- Ian Bell and contributors. Coolprop: Open-source thermodynamic and transport properties database. <https://pypi.org/project/coolprop/>, 2025. Version 6.8.0, released April 21, 2025. Available on PyPI. Maintained by Ian Bell and contributors.
- Andres M Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, and Philippe Schwaller. Chemcrow: Augmenting large-language models with chemistry tools, 2023. URL <https://arxiv.org/abs/2304.05376>.
- Julian Cremer, Ross Irwin, Alessandro Tibo, Jon Paul Janet, Simon Olsson, and Djork-Arné Clevert. Flowr: Flow matching for structure-aware de novo, interaction- and fragment-based ligand generation, 2025. URL <https://arxiv.org/abs/2504.10564>.
- Ian Dunn and David R. Koes. Exploring discrete flow matching for 3d de novo molecule generation, 2024. URL <https://arxiv.org/abs/2411.16644>.
- Benjamin Eysenbach and Sergey Levine. Maximum entropy rl (provably) solves some robust rl problems, 2022. URL <https://arxiv.org/abs/2103.06257>.
- David S. Farina., Sai Krishna Sirumalla, Emily J. Mazeau, and Richard H. West. Extensive high-accuracy thermochemistry and group additivity values for halocarbon combustion modeling. *Industrial & Engineering Chemistry Research*, 60(41):15492–15501, oct 2021. doi: 10.1021/acs.iecr.1c03076. URL <https://doi.org/10.1021/acs.iecr.1c03076>.
- Anna Gaulton, Anne Hersey, Michal Nowotka, A. Patrícia Bento, Jon Chambers, David Mendez, Prudence Mutowo, Fraser Atkinson, Louisa J. Bellis, Élodie Cibrián-Uhalte, Mark Davies, Neil Dedman, Anders Karlsson, Marta P. Magariños, John P. Overington, George Papadatos, Ilenia Smit, and Andrew R. Leach. ChEMBL: a large-scale bioactivity database for drug discovery, 2017. URL <https://www.ebi.ac.uk/chembl/>. Accessed: 2025-07-25.
- Juliane Glüge, Katharina Breuer, Armin Hafner, Christian Vering, Dirk Müller, Ian T. Cousins, Rainer Lohmann, Gretta Goldenman, and Martin Scheringer. Finding non-fluorinated alternatives to fluorinated gases used as refrigerants. *Environmental Science: Processes & Impacts*, 26(11):1955–1974, 2024. doi: 10.1039/d4em00444b. Published online 2024 Oct 3.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. URL <https://arxiv.org/abs/2407.21783>. Version 3, released November 23, 2024.
- Kehang Han, Balaji Lakshminarayanan, and Jeremiah Liu. Reliable graph neural networks for drug discovery under distributional shift, 2021. URL <https://arxiv.org/abs/2111.12951>.
- Daniel Hesslow, Niccoló Zanichelli, Pascal Notin, Iacopo Poli, and Debora Marks. Rita: a study on scaling up generative protein sequence models, 2022. URL <https://arxiv.org/abs/2205.05789>.

- Moksh Jain, Sharath Chandra Raparthy, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Yoshua Bengio, Santiago Miret, and Emmanuel Bengio. Multi-objective gflownets. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, volume 202 of *Proceedings of Machine Learning Research*, pages 14269–14313. PMLR, 2023. URL <https://proceedings.mlr.press/v202/jain23c.html>.
- Andrei Kazakov, Mark O. McLinden, and Michael Frenkel. Computational design of new refrigerant fluids based on environmental, safety, and thermodynamic characteristics. *Industrial & Engineering Chemistry Research*, 51(38):12534–12542, 2012. doi: 10.1021/ie3016126.
- Sunghwan Kim, Jian Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Sherry He, Qingliang Li, Benjamin A. Shoemaker, Paul A. Thiessen, Bohu Yu, Len Zaslavsky, Jian Zhang, and Evan E. Bolton. Pubchem in 2021: new data content and improved web interfaces. *Nucleic Acids Research*, 49(D1):D1388–D1395, 2021. doi: 10.1093/nar/gkaa971. URL <https://pubchem.ncbi.nlm.nih.gov/>. Accessed: 2025-07-25.
- Eric S. C. Kwok and Roger Atkinson. Estimation of hydroxyl radical reaction rate constants for gas-phase organic compounds using a structure–reactivity relationship: An update. *Atmospheric Environment*, 29(14):1685–1695, 1995.
- Hannes H. Loeffler, Jiazhen He, Alessandro Tibo, Jon Paul Janet, Alexey Voronov, Lewis H. Mervin, and Ola Engkvist. Reinvent 4: Modern ai–driven generative molecule design. *Journal of Cheminformatics*, 16(1):20, 2024. doi: 10.1186/s13321-024-00812-5. URL <https://doi.org/10.1186/s13321-024-00812-5>.
- Mathilda Maury, Didier Mathieu, and Johan Jacquemin. General estimation method for lower flammability limits of organic compounds: The simpler the better. *Fuel*, 345:127753, 2023. doi: 10.1016/j.fuel.2023.127753. URL <https://doi.org/10.1016/j.fuel.2023.127753>.
- Max R. McGillen, William P. L. Carter, Abdelwahid Mellouki, John J. Orlando, Bénédicte Picquet-Varrault, and Timothy J. Wallington. Database for the kinetics of the gas-phase atmospheric reactions of organic compounds. *Earth System Science Data*, 12:1203–1216, 2020. URL <https://doi.org/10.5194/essd-12-1203-2020>.
- Mark O. McLinden, J. Steven Brown, Riccardo Brignoli, Andrei F. Kazakov, and Piotr A. Domanski. Limited options for low-global-warming-potential refrigerants. *Nature Communications*, 8:14476, 2017. doi: 10.1038/ncomms14476. URL <https://www.nature.com/articles/ncomms14476>.
- Balaganesh Muthiah, Shih-Cheng Li, and Yi-Pei Li. Developing machine learning models for accurate prediction of radiative efficiency of greenhouse gases. *Journal of the Taiwan Institute of Chemical Engineers*, 151:105123, 2023.
- Siddharth M. Narayanan, James D. Braza, Ryan-Rhys Griffiths, Albert Bou, Geemi Wellawatte, Mayk Caldas Ramos, Ludovico Mitchener, Samuel G. Rodrigues, and Andrew D. White. Training a scientific reasoning model for chemistry, 2025. URL <https://arxiv.org/abs/2506.17238>.
- OECD. Reconciling terminology of the universe of per- and polyfluoroalkyl substances: Recommendations and practical guidance. Technical Report 61, Organisation for Economic Co-operation and Development (OECD), Paris, 2021.
- George Papadatos, Mark Davies, Nathan Dedman, Jon Chambers, Anna Gaulton, James Siddle, Richard Koks, Sean A. Irvine, Joe Pettersson, Nicko Goncharoff, Anne Hersey, and John P. Overington. SureChEMBL: a large-scale, chemically annotated patent document database. *Nucleic Acids Research*, 44(D1):D1220–D1228, 2016. doi: 10.1093/nar/gkv1253. URL <https://www.surechembl.org/>. Accessed: 2025-07-25.
- Ding-Yu Peng and Donald B. Robinson. A new two-constant equation of state. *Industrial & Engineering Chemistry Fundamentals*, 15(1):59–64, 1976. doi: 10.1021/i160057a011.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. URL <https://arxiv.org/abs/2305.18290>. arXiv preprint arXiv:2305.18290.

- Jerret Ross, Brian Belgodere, Vijil Chenthamarakshan, Inkit Padhi, Youssef Mroueh, and Payel Das. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(12):1256–1264, December 2022. doi: 10.1038/s42256-022-00580-7. URL <https://www.nature.com/articles/s42256-022-00580-7>.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y.K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. arXiv preprint arXiv:2405.20094, 2024. Available at <https://github.com/deepseek-ai/DeepSeek-Math>.
- Eduardo Soares, Emilio Vital Brazil, Victor Yukio Shirasuna, Renato Cerqueira, Dmitry Zubarev, and Kristin Schmidt. Smi-ted: A large-scale foundation model for materials and chemistry. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. URL <https://iclr.cc>. Published as a conference paper at ICLR 2025.
- Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, 2025. URL <https://doi.org/10.48550/arXiv.2412.17780>. Published at ICML 2025, Vancouver, Canada.
- Benjamin I. Tingle, Khanh G. Tang, Mar Castanon, John J. Gutierrez, Munkhzul Khurelbaatar, Chinzorig Dandarchuluun, Yurii S. Moroz, and John J. Irwin. Zinc-22: A free multi-billion-scale database of tangible compounds for ligand discovery. *Journal of Chemical Information and Modeling*, 63(4):1166–1176, February 2023. doi: 10.1021/acs.jcim.2c01253. URL <https://pubs.acs.org/doi/10.1021/acs.jcim.2c01253>.
- United Nations Environment Programme Ozone Secretariat. Montreal Protocol on Substances that Deplete the Ozone Layer. <https://ozone.unep.org/treaties/montreal-protocol>, 1987. Accessed: 2025-08-15.
- U.S. Environmental Protection Agency. Estimation Programs Interface Suite™ for Microsoft® Windows, v 4.11. <https://www.epa.gov/tsca-screening-tools/epi-suitetm-estimation-program-interface>, 2012. United States Environmental Protection Agency, Washington, DC, USA.
- U.S. Environmental Protection Agency, Office of Research and Development (ORD). Ipcc ar4, ar5, and ar6 20-, 100-, and 500-year global warming potentials (gwps). <https://www.data.gov/dataset/ipcc-ar4-ar5-and-ar6-20-100-and-500-year-gwps>, December 2023. Dataset created using LCIA Formatter v1.1.1. Includes data from IPCC AR4, AR5, and AR6 for 20-, 100-, and 500-year time horizons. Names for GHGs follow FEDEFL v1.2. Maintainer: Wesley Ingwersen.
- Clément Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. DiGress: Discrete Denoising Diffusion for Graph Generation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023. URL <https://arxiv.org/abs/2209.14734>. Preprint at arXiv:2209.14734.
- David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988. doi: 10.1021/ci00057a005. URL <https://doi.org/10.1021/ci00057a005>.
- Artem Zholus, Maksim Kuznetsov, Roman Schutski, Rim Shayakhmetov, Daniil Polykovskiy, Sarath Chandar, and Alex Zhavoronkov. Bindgpt: A scalable framework for 3d molecular design via language modeling and reinforcement learning, 2024. URL <https://arxiv.org/abs/2406.03686>.
- Vladimir Zieger, Thibaut Lecompte, Simon Guihéneuf, Yann Guevel, Manuel Bazzana, Thomas Gasser, and Yue He. Climate change metrics: Ipcc ar6 updates, discussions and dynamic assessment applications. *EGUsphere*, 2025. URL <https://doi.org/10.5194/egusphere-2024-3918>. Preprint. Discussion started: 3 February 2025.

A Detailed Refgen pipeline

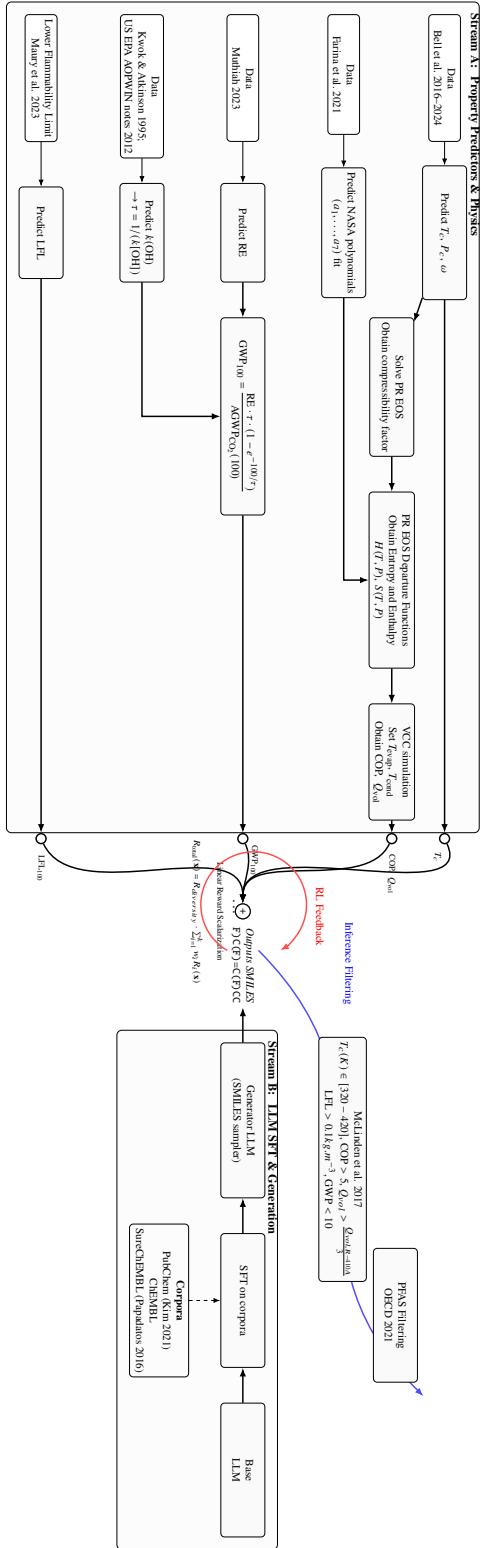


Figure 10: Detailed Refgen pipeline with property prediction details and datasets (Stream A, left) as well as SMILES generation (Stream B, left).

B Related Work

While traditional screening is limited to known compounds from databases, recent *generative modeling* techniques have shown an ability to explore the chemical space, allowing for the generation of new (*’de novo’*) molecular structures while balancing different property tradeoffs, constraining the molecular space towards a specific region of interest. Methods vary by design choices: using graph representations [Han et al., 2021] or textual e.g. SMILES [?]; using sequence models [Ross et al., 2022, Loeffler et al., 2024], GFlowNets [Jain et al., 2023] or discrete diffusion [Tang et al., 2025, Vignac et al., 2023] and flow matching Dunn and Koes [2024], Cremer et al. [2025]. Recently sequence models, usually in the context of *Large Language Models* (LLMs), gained traction in the molecule discovery community due to their scalability and versatility, allowing to learn from collections of billions of compounds e.g. ZINC Tingle et al. [2023]. Post-training fine-tuning techniques for LLMs, such as as Direct Preference Optimization (DPO) [Rafailov et al., 2023] or Group Relative Policy Optimization (GRPO) [Shao et al., 2024], well transferred from the natural language processing community to chemistry and computational biology, allowing for the development of chemical reasoning models Bran et al. [2023], Narayanan et al. [2025] or models for molecular structures Zholus et al. [2024], Ross et al. [2022] and proteins Hesslow et al. [2022].

C COP calculation and cycle simulation

C.1 The Peng-Robinson Equation of State

To model the behavior of gases and liquids under different conditions, we use an equation of state. While the simplest EOS is usually the *ideal gas law*, more refined models, such as the *Peng–Robinson (PR) equation of state*, incorporate repulsion and attraction between molecules. Its form is:

$$P = \frac{RT}{V_m - b} - \frac{a\alpha(T)}{V_m(V_m + b) + b(V_m - b)}.$$

where P is pressure, V_m is molar volume, T is temperature, and R is the ideal gas constant. The parameters a , b , and the temperature-dependent factor $\alpha(T)$ are expressed in terms of critical properties and the acentric factor:

$$b = 0.07780 \frac{RT_c}{P_c}, \quad a = 0.45724 \frac{R^2 T_c^2}{P_c}, \quad \alpha(T) = \left[1 + m \left(1 - \sqrt{T/T_c} \right) \right]^2,$$

with

$$m = 0.37464 + 1.54226\omega - 0.26992\omega^2.$$

Importantly, these coefficients depend on the critical temperature T_c , critical pressure P_c as well as the acentric factor ω , which need to be inputs to the model.

To quantify how real fluids deviate from ideal behavior, we define the *compressibility factor*:

$$Z = \frac{PV_m}{RT}.$$

For an ideal gas, $Z = 1$. Deviations from unity reflect the presence of intermolecular forces: $Z < 1$ indicates attraction, while $Z > 1$ suggests repulsion or volume exclusion.

The Peng Robinson EOS can be rewritten as a third degree polynomial in the compressibility factor Z .

$$Z^3 - (1 - B)Z^2 + (A - 3B^2 - 2B)Z - (AB - B^2 - B^3) = 0$$

Where

$$A = \frac{a\alpha P}{(RT)^2} \quad \text{and} \quad B = \frac{bP}{RT}$$

The roots of this equation are key to obtaining thermodynamic relationships at the two phase region between liquid and vapor, and obtaining the COP. Solving the equation can lead to two scenarios :

- In the two phase region : 3 real roots

- The smallest root : Z_{liq} of the liquid phase
- The largest root : Z_{vap} of the vapor phase
- The middle root is an artifact with no thermodynamic significance
- Outside the two phase region : 1 real root

C.2 Ideal to Real fluid and gas behaviour with Nasa Polynomials

C.2.1 Enthalpy, Entropy and Heat Capacity

To interpret thermodynamic models, we must define three key quantities: enthalpy, entropy, and heat capacity.

The *enthalpy* H includes internal energy and the energy required to occupy volume:

$$H = U + PV.$$

It reflects the heat required to raise temperature at constant pressure.

The *entropy* S quantifies the degree of disorder or number of accessible microstates. Its temperature derivative is related to the heat capacity:

$$\left(\frac{\partial S}{\partial T} \right)_P = \frac{C_p}{T}.$$

The *heat capacity at constant pressure* is given by

$$C_p = \left(\frac{\partial H}{\partial T} \right)_P.$$

C.2.2 Nasa Polynomial thermodynamic representations

NASA polynomials address the challenge of consistently calculating heat capacity, enthalpy, and entropy while maintaining computational efficiency.

Given the heat capacity $C_p(T)$ as a function of temperature T , the thermodynamic relationships enable us to derive the enthalpy $H(T)$ and entropy $S(T)$ through integration with respect to temperature:

$$\begin{aligned} H(T) - H(T_{\text{ref}}) &= \int_{T_{\text{ref}}}^T C_p(\tau) d\tau \\ S(T) - S(T_{\text{ref}}) &= \int_{T_{\text{ref}}}^T \frac{C_p(\tau)}{\tau} d\tau \end{aligned}$$

where T_{ref} is a reference temperature at which enthalpy and entropy are defined relative to some standard state.

NASA uses a 4th-degree polynomial representation for the dimensionless heat capacity:

$$C_p(T)/R = a_0 + a_1 T + a_2 T^2 + a_3 T^3 + a_4 T^4$$

where R is the ideal gas constant and a_i are the polynomial coefficients. For this form, the indefinite integrals required for enthalpy and entropy can be calculated analytically:

$$\int \frac{C_p(T)}{R} dT = a_0 T + \frac{a_1}{2} T^2 + \frac{a_2}{3} T^3 + \frac{a_3}{4} T^4 + \frac{a_4}{5} T^5 + \text{constant}_H$$

and similarly for the entropy:

$$\int \frac{C_p(T)}{RT} dT = \int \left(\frac{a_0}{T} + a_1 + a_2 T + a_3 T^2 + a_4 T^3 \right) dT = a_0 \ln(T) + a_1 T + \frac{a_2}{2} T^2 + \frac{a_3}{3} T^3 + \frac{a_4}{4} T^4 + \text{constants}$$

A key property that arises from these analytical expressions is the shared coefficients across all three thermodynamic quantities. This is a form of "parameter sharing" ensures the desired consistency. The

resulting representations for heat capacity, enthalpy, and entropy in the NASA 7-coefficient format are:

$$\begin{aligned} C_p(T)/R &= a_0 + a_1 T + a_2 T^2 + a_3 T^3 + a_4 T^4 \\ H(T)/RT &= a_0 + \frac{a_1}{2} T + \frac{a_2}{3} T^2 + \frac{a_3}{4} T^3 + \frac{a_4}{5} T^4 + \frac{a_5}{T} \\ S(T)/R &= a_0 \ln(T) + a_1 T + \frac{a_2}{2} T^2 + \frac{a_3}{3} T^3 + \frac{a_4}{4} T^4 + a_6 \end{aligned}$$

The coefficients

$$a_5 = -H(T_{\text{ref}})/(RT_{\text{ref}})$$

and

$$a_6 = S(T_{\text{ref}})/R - a_0 \ln(T_{\text{ref}}) - a_1 T_{\text{ref}} - \frac{a_2}{2} T_{\text{ref}}^2 - \frac{a_3}{3} T_{\text{ref}}^3 - \frac{a_4}{4} T_{\text{ref}}^4$$

incorporate the integration constants, setting the reference enthalpy and entropy values at a chosen reference temperature T_{ref} (often 298.15 K). By defining $H(T)$ and $S(T)$ through the analytical integrals of the same $C_p(T)$ polynomial, the fundamental thermodynamic relationships between C_p , H and S are inherently satisfied.

C.3 Ideal to Real fluid and gas behaviour : Departure functions

The NASA polynomials only capture ideal gas entropy, enthalpy and thermal capacity. To capture the deviations of these quantities from ideal gas behavior to obtain the real behaviour, we can use the Peng-Robinson EOS, and compute *departure functions*:

$$H^{\text{dep}} = H^{\text{real}} - H^{\text{ideal}}, \quad S^{\text{dep}} = S^{\text{real}} - S^{\text{ideal}}.$$

These corrections are derived analytically from the EOS and depend on Z , a , b , and $\frac{d(a\alpha)}{dT}$. The formulas are as follows :

$$\begin{aligned} H^{\text{dep}} &= RT(Z-1) + \frac{T \frac{da(T)}{dT} - a(T)}{2\sqrt{2}b} \ln \left(\frac{Z + (1 + \sqrt{2})B}{Z + (1 - \sqrt{2})B} \right) \\ S^{\text{dep}} &= R \ln(Z - B) + \frac{\frac{da(T)}{dT}}{2\sqrt{2}b} \ln \left(\frac{Z + (1 + \sqrt{2})B}{Z + (1 - \sqrt{2})B} \right) \end{aligned}$$

where:

- Z is the compressibility factor (solved from the cubic form of PR EOS: $Z^3 - (1 - B)Z^2 + (A - 3B^2 - 2B)Z - (AB - B^2 - B^3) = 0$)
- $A = \frac{a(T)P}{R^2 T^2}$
- $B = \frac{bP}{RT}$
- $\frac{da(T)}{dT}$ is the temperature derivative of $a(T)$, computed as:

$$\frac{da(T)}{dT} = a_c \cdot \frac{-2\kappa}{T_c^{0.5}} \left[1 + \kappa \left(1 - \sqrt{\frac{T}{T_c}} \right) \right] \left(1 - \sqrt{\frac{T}{T_c}} \right)$$

In these expressions, the first term accounts for the volume effect, while the second term corrects for the temperature dependence of the attractive forces.

Real-fluid enthalpy and entropy are then expressed as:

$$H(T, P) = H^{\text{ideal}}(T) + H^{\text{dep}}(T, P, Z, a, b, \frac{d(a\alpha)}{dT}),$$

$$S(T, P) = S^{\text{ideal}}(T, P) + S^{\text{dep}}(T, P, Z, a, b, \frac{d(a\alpha)}{dT}).$$

C.4 Phase equilibrium and saturation pressure

Another important element to reconstruct the vapor construction cycle is the relationship between temperature and pressure at a phase equilibrium. At equilibrium, temperature and pressure are not independent variables, as one uniquely determines the other. It is therefore important to be able to get a certain P_{sat} for a given T at saturation, and vice-versa.

C.4.1 Chemical Potential and Fugacity

The *chemical potential* μ is defined as the change in Gibbs free energy G with respect to the number of moles:

$$\mu = \left(\frac{\partial G}{\partial n} \right)_{T,P}.$$

It governs mass transfer, phase change, and chemical reaction. At phase equilibrium, the chemical potentials in coexisting phases are equal:

$$\mu^{\text{liq}} = \mu^{\text{vap}}.$$

To evaluate chemical potential in real fluids, we introduce the *fugacity* f , related to pressure through the *fugacity coefficient* ϕ :

$$f = \phi P.$$

For an ideal gas, $\phi = 1$, so $f = P$. For real fluids, $\phi \neq 1$, and is computed from the EOS. The chemical potential can then be written as:

$$\mu(T, P) = \mu^{\text{ideal}}(T, P^\circ) + RT \ln \left(\frac{f}{P^\circ} \right),$$

where P° is a reference pressure.

C.4.2 P and T relationship at phase equilibrium

At a given temperature, the vapor and liquid phases are in equilibrium when their fugacities match:

$$f^{\text{liq}}(T, P_{\text{sat}}) = f^{\text{vap}}(T, P_{\text{sat}}),$$

which implies

$$\phi^{\text{liq}}(T, P_{\text{sat}}) = \phi^{\text{vap}}(T, P_{\text{sat}}).$$

Solving this equation yields the saturation pressure P_{sat} at temperature T , a key quantity for modeling phase behavior in refrigerant cycles.

The Peng-Robinson EOS gives a relationship between ϕ and Z , as follows :

$$\ln \phi = Z - 1 - \ln(Z - B) - \frac{A}{2\sqrt{2}B} \ln \left(\frac{Z + (1 + \sqrt{2})B}{Z + (1 - \sqrt{2})B} \right)$$

with the A and B parameters being the same ones presented earlier. We can therefore for example get any saturation pressure P_{sat} from T as follows :

Algorithm 3 Calculation of VCC State Points, COP, and Q_{vol} (Temperature-Driven)

Require: Evaporator Temperature T_{evap} , Condenser Temperature T_{cond} , Fluid Properties params, Pre-calculated PR+NASA Saturation Dome Data `dome_data_pr`

- 1: Extract and prepare valid temperature-saturation property data (\mathbf{T}_{sat} , \mathbf{P}_{sat} , \mathbf{h}_L , \mathbf{h}_V , \mathbf{s}_L , \mathbf{s}_V , \mathbf{v}_V) from `dome_data_pr`.
- 2: Create interpolation functions for saturation properties based on temperature. ▷ Determine Saturation Pressures
- 3: $P_{evap} \leftarrow \text{interpolate}(T_{evap})$ from \mathbf{P}_{sat}
- 4: $P_{cond} \leftarrow \text{interpolate}(T_{cond})$ from \mathbf{P}_{sat} ▷ Define Cycle State Points
- 5: **State 1 (Evaporator Outlet - Saturated Vapor):**

$$\begin{aligned} h_1 &\leftarrow \text{interpolate}(T_{evap}) \text{ from } \mathbf{h}_V \\ s_1 &\leftarrow \text{interpolate}(T_{evap}) \text{ from } \mathbf{s}_V \\ v_1 &\leftarrow \text{interpolate}(T_{evap}) \text{ from } \mathbf{v}_V \\ P_1 &\leftarrow P_{evap} \end{aligned} \quad \text{▷ Get specific volume}$$
- 6: **State 3 (Condenser Outlet - Saturated Liquid):**

$$\begin{aligned} h_3 &\leftarrow \text{interpolate}(T_{cond}) \text{ from } \mathbf{h}_L \\ P_3 &\leftarrow P_{cond} \end{aligned}$$
- 7: **State 4 (Expansion Valve Outlet - Isenthalpic):**

$$\begin{aligned} h_4 &\leftarrow h_3 \\ P_4 &\leftarrow P_{evap} \end{aligned}$$
- 8: **State 2 (Compressor Outlet - Isentropic):**

$$\begin{aligned} s_2 &\leftarrow s_1 \\ P_2 &\leftarrow P_{cond} \end{aligned}$$
- 9: Find T_2 such that the calculated entropy at (T_2, P_2) equals s_2 .
- 10: Calculate h_2 using (T_2, P_2) and the PR+NASA model. ▷ Calculate Performance Metrics
- 11: Compressor Work: $W_c = h_2 - h_1$
- 12: Refrigeration Effect: $Q_e = h_1 - h_4$
- 13: COP $\leftarrow Q_e/W_c$ (if $W_c > 0$ and $Q_e > 0$)
- 14: Volumetric Capacity: $Q_{vol} \leftarrow Q_e/v_1$
- 15: **Return** (h_1, h_2, h_3, h_4) , (P_1, P_2, P_3, P_4) , COP, Q_{vol} .

We can then compute

$$COP = \frac{\text{Refrigerating Effect}}{\text{Compressor Work}} = \frac{h_1 - h_4}{h_2 - h_1} \quad Q_{vol} = \frac{h_1 - h_4}{v_1}$$

D Choice of the property predictor

D.1 The XGBoost baseline

Considering that most datasets of properties obtained in this project are on the smaller side, a natural first baseline to test out were tree based models trained on features extracted from the SMILES representation of molecules.

We train an XGBoost model with the following hyperparameters :

- Estimator number : 1000
- Learning rate : 0.1
- Maximum depth 6
- Subsample 0.8
- Columns (fraction) sampled by tree : 0.8
- Patience 10
- Early stopping metric : Mean Absolute Error

Using cheminformatics libraries such as RDKit, we extract multiple chemical descriptors from SMILES into a large feature vector : Morgan Fingerprints; Molecular Descriptors such as molecular weight, topological polar surface area (TPSA), LogP, and connectivity indices; Fragment Counts: functional group presence is encoded through SMARTS pattern matching, counting occurrences of chemically relevant substructures (e.g., hydroxyl groups, aromatic rings, halogens).

D.2 The SMI-TED Model

We end-to-end finetune the SMI-TED foundation model [Soares et al., 2025] on our property datasets, leveraging the learned encoded latent representations, and adding an MLP head, following the architecture and the code from the authors. SMI-TED hyperparameters are the same as for the final property predictors, and is detailed in Appendix F.2.

D.3 Experiments and results

We compare the two models on our critical temperature dataset, as well as an additional testing set from the CoolProp database. Details on dataset preprocessing are in the property prediction training details in Appendix F.

We compare results with SMILES augmentation, adding 4 additional graph traversals to the initial dataset SMILES, for a total of 5, as it ensured that >95% of all molecules had that many graph traversals and avoided dataset imbalance.

Table 3: Percentage of molecules with 5 unique graph traversals

Target Property	Percentage of Molecules (%)
Tc (Critical Temperature)	97.8%
Pc (Critical Pressure)	97.9%
ω (Acentric Factor)	97.9%
Nasa Polynomials	99.2%
Lower Flammability Limit (LFL)	96.0%
$k(OH)$ Rate Constant	96.4%
Radiative Efficiency (RE)	99.8%

Results are compared and averaged on 5 random seeds.

Table 4: Mean Prediction Performance of XGBoost vs. SMI-TED

Dataset	Augmentation	XGBoost			SMI-TED		
		R ²	MAE	RMSE	R ²	MAE	RMSE
Test Set	No	0.946	17.86	33.00	0.938	15.94	34.64
	Yes	0.942	17.58	33.32	0.946	14.40	31.06
CoolProp	No	0.554	48.92	106.80	0.944	19.72	33.18
	Yes	0.420	56.16	121.68	0.938	17.94	32.68

Not only do the SMILES augmented end-to-end finetuned SMI-TED, avoid the need for handcrafted features, the models showcase better prediction accuracy and especially, better generalization ability on the CoolProp dataset. This shows that despite being trained on very diverse data, the SMI-TED model generalizes better to target refrigerant molecules. This setup is chosen for the final property prediction architecture.

E Supervised fine-tuning stage

E.1 Data curation

We consider the concatenation of Pubchem Compounds Kim et al. [2021], ChEMBL Gaulton et al. [2017] and SureChEMBL Papadatos et al. [2016], amounting to a total of $1.47 \cdot 10^8$ smiles. We apply a filtering pipeline based on the following criteria: (1) excluding molecules with undesired substructures from a predefined list; (2) salts are removed with RDKit’s `SaltRemover`; (3) in case of fragments, keep only the largest one in the molecule with RDKit’s `LargestFragmentChooser`; (4) remove any atom mapping numbers e.g. [CH3:1] → CH3; (5) neutralize the molecule with RDKit’s `Uncharger`; (6) Normalize the molecule with RDKit’s `Normalizer`, handling aromatization, kekulization, normalization of tautomers, charges, and representations. The output subset from this pipeline counts $3.72 \cdot 10^7$ smiles (25.3% of the original dataset, hence 74.7% rejected structures).

E.2 Training parameters

The SFT training employs the following hyperparameters :

- Base Model: Llama-3.2-1B
- Batch Size: 256 (with gradient accumulation: $\frac{256}{\text{max_gpu_batch_size}}$)
- GPU batch size: 16
- Hardware: 4 × NVIDIA A100:80GB
- Learning Rate: $\alpha = 2 \times 10^{-5}$ with linear scheduling
- Optimizer: AdamW with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$
- Sequence Length: maximum 256 tokens
- Training Epochs: 1 epoch over the entire dataset

Let $\mathcal{D}_{SFT} = \{\mathbf{x}\}_{i=1}^N$ be the molecular dataset, θ the model parameters, $\mathcal{T}^{(i)}$ the length of sequence i , and $p_\theta(x_t^{(i)}|x_{<t}^{(i)})$ the probability of generating token $x_t^{(i)}$ given the previous tokens $x_{<t}^{(i)} = (x_1^{(i)}, \dots, x_{t-1}^{(i)})$.

During training, the loss is computed only on the SMILES tokens, not on the property conditioning tokens. This is achieved through attention masking:

$$\mathcal{L}_{SFT} = -\frac{1}{N} \sum_{i=1}^N \sum_{t \in \mathcal{T}_{\text{SMILES}}^{(i)}} \log p_\theta(x_t^{(i)}|x_{<t}^{(i)})$$

where $\mathcal{T}_{\text{SMILES}}^{(i)}$ represents the set of positions corresponding to SMILES tokens in sequence i .

F Training property predictors

F.1 Dataset processing and splits

Smiles canonicalization is not performed to avoid loosing data from multiple data entry points which could have relevant information.

Smiles augmentation is performed, according to results in D showing its benefits to model generalization. For each molecule in the training set, four additional SMILES traversals are generated.

Dataset cleaning : Metals and metalloids are discarded if they exist in the data as they are too far out of distribution from our molecules of interest. This is mostly of use for the thermodynamic properties dataset which is a concatenation of multiple existing datasets extracted from the internet and not already curated. **Single elemental atoms** or ions are also discarded as their properties are often very different from the rest of the molecules and do not represent structures of interest, especially in the case of refrigerant discovery.

Scaffolds Separation of training and testing by molecular scaffolds, also known as core molecular structures, is a common technique to better verify the generalization properties of the model. Molecules with the same scaffolds are grouped together to avoid having too similar structures in different dataset splits. We therefore separate train/validation and test splits by Bemis-Murcko scaffolds, using the RdKit library.

The train / test / validation split is 70/15/15.

Final dataset sizes and splits are as follows :

Table 5: Dataset split sizes for different target properties

Dataset sizes	Train	Augmented Train	Validation	Test
Tc	8079	39615	1358	1612
Pc	8322	40816	1329	1409
ω	4940	24269	1074	968
Nasa Polynomials	11924	59435	2569	2309
RE	55938	279615	11676	14380
k(OH) constant	937	4552	105	131
LFL	1010	4863	180	171

Table 6: External test sets for model evaluation

Property	External Dataset	Test Size
Tc	CoolProp	106
Pc	CoolProp	107
ω	CoolProp	107
GWP	IPCC Reports	220

F.2 SMI-TED hyperparameter details

Hyperparameter details :

- Batch size: 256
- Learning rate: 3×10^{-5}
- Learning rate multiplier: 1.0
- Optimizer: AdamW with $\beta_1 = 0.9$, $\beta_2 = 0.99$
- Loss function: Mean Absolute Error (MAE)
- Early stopping epochs : 10

F.3 Group contribution for kOH constant calculation

Group contribution methods estimate molecular properties by decomposing a molecule into predefined substructures (or “groups”) and summing their individual contributions. Each group has an associated rate or factor derived from experimental data, and the total property (e.g., reaction rate) is the sum of the contributions from all relevant groups in the molecule. This enables rapid predictions without full quantum chemical calculations.

We implemented a group contribution method based Kwok and Atkinson [1995] and the EPA’s AOPWIN software refinements [U.S. Environmental Protection Agency, 2012], which estimates the OH radical reaction rate constants by summing contributions from four main pathways:

- H-abstraction from C–H, O–H, or S–H bonds
- OH addition to unsaturated bonds (C = C, C ≡ C)
- OH addition to aromatic rings

- OH interactions with specific functional groups

The total rate constant is modeled as:

$$k_{OH,tot} = k_{H\text{ abs}} + k_{OH\text{ add, unsat}} + k_{OH\text{ add, ring}} + k_{OH\text{ func. group}}$$

Each term is computed using empirically fitted base rate constants and multiplicative modifiers called **substituent factors**. For example, in H-abstraction from saturated carbon atoms, the rate is determined using formulas like:

$$k(X - CH2 - Y) = k_{\text{sec}} \cdot F(X) \cdot F(Y), \quad k(X - CH < Y, Z) = k_{\text{tert}} \cdot F(X) \cdot F(Y) \cdot F(Z)$$

Here, k_{sec} and k_{tert} are base rate constants for secondary and tertiary hydrogens, and $F(\cdot)$ are substituent factors that encode how reactive neighboring groups are.

G Reinforcement Learning Fine-Tuning stage

G.1 Reward functions

G.1.1 Molecule Validity

As the base model has been finetuned to generate valid SMILES, we need to ensure that our RL finetuning does not make the distribution diverge to the point it stops generating valid molecules. We therefore build a reward for molecule validity alongside property-specific rewards:

$$R_{\text{final}}(\mathbf{x}) = R_{\text{validity}}(\mathbf{x}) + R_{\text{properties}}(\mathbf{x})$$

where:

$$R_{\text{validity}}(\mathbf{x}) = \begin{cases} 1.0 & \text{if SMILES is valid and parseable} \\ 0.0 & \text{otherwise} \end{cases}$$

$$R_{\text{properties}}(\mathbf{x}) = \begin{cases} \text{property-based score} & \text{if SMILES is valid and parseable} \\ 0.0 & \text{otherwise} \end{cases}$$

This additive formulation ensures that invalid molecules receive zero total reward (since both validity and property rewards are zero), while valid molecules receive their property-based score plus a +1.0 validity bonus, strongly encouraging the model to maintain chemical validity while optimizing for desired characteristics.

We verify two elements for validity :

- We ensure the SMILES is a valid molecule through the RdKit library
- We also ensure that each atom in the the molecule has the correct valence (also through RdKit). We consider molecules with radicals to be invalid as they are very unstable.

This second element in the validity reward was added after later analysis that, as the model strayed from the initial distribution, it tended to generate molecules with radicals (i.e with squared brackets in the molecule), which had good properties, but are very unstable in practice. Stability, despite not being its own reward or a property analyzed in detail in this project, is key for refrigerants.

G.1.2 Critical temperature

Critical temperature is an essential property of a refrigerant and is used here, as in Kazakov et al. [2012] and McLinden et al. [2017], as the main thermodynamic criterion for the selection of molecules. They adopt lower bounds of 300K for the minimal operating temperature of the condenser in the cycle, as well as a generous upper bound of 550K, explaining that machines with centrifugal compressors can work with refrigerants with T_c up to 470K+. The ideal range however, has been shown, and is explicitly given in McLinden et al. [2017] as [320K - 420K].

We therefore build a tiered exponentially decreasing reward as such, with a reward of 1 in the ideal range, and attaining 0.5 in the larger 300-550 range.

Let T be the temperature in Kelvin. The function $f(T)$ is defined as follows:

$$f(T) = \begin{cases} e^{-k_{\text{left}}(T_{\text{plateau_min}} - T)} & \text{if } T < T_{\text{plateau_min}} \\ 1.0 & \text{if } T_{\text{plateau_min}} \leq T \leq T_{\text{plateau_max}} \\ e^{-k_{\text{right}}(T - T_{\text{plateau_max}})} & \text{if } T_{\text{plateau_max}} < T \leq T_{\text{target_max_temp}} \\ R_{\text{target}} \cdot e^{-k_{\text{left}}(T - T_{\text{target_max_temp}})} & \text{if } T > T_{\text{target_max_temp}} \end{cases}$$

where $T_{\text{plateau_min}}=320K$, $T_{\text{plateau_max}}=420K$, $T_{\text{target_min_temp}} = 300K$, $T_{\text{target_max_temp}} = 550K$, the target reward at boundary temperature $R_{\text{target}} = 0.5$ and $k_{\text{left}} = \frac{-\ln(R_{\text{target}})}{T_{\text{plateau_min}} - T_{\text{target_min_temp}}}$ and $k_{\text{right}} = \frac{-\ln(R_{\text{target}})}{T_{\text{target_max_temp}} - T_{\text{plateau_max}}}$

G.1.3 Molecular length

The length penalty is essential to obtain tractable molecules that are not completely out of distribution. Moreover, refrigerants are small molecules. Empirical results and thermodynamic results show that refrigerants have rarely over 15-18 atoms. In their work, Kazakov et al. [2012] filter PubChem for under 15 atoms while in their final work [McLinden et al., 2017], a larger range of up to 18 atoms are analyzed.

We use RdKit to count the number of atoms. A custom gaussian-like reward function with a plateau and custom curve steepness is used, to avoid generating individual atoms, as well as molecules over 18 atoms.

Let x be the measured property value (e.g., number of atoms). The function $g(x)$ is defined as follows:

$$g(x) = \begin{cases} e^{-0.5\left(\frac{|x - P_{\text{start}}|}{\sigma}\right)^S} & \text{if } x < P_{\text{start}} \\ 1.0 & \text{if } P_{\text{start}} \leq x \leq P_{\text{end}} \\ e^{-0.5\left(\frac{|x - P_{\text{end}}|}{\sigma}\right)^S} & \text{if } x > P_{\text{end}} \end{cases}$$

where the start of the ideal range $P_{\text{start}} = 7$, the end of the ideal range $P_{\text{end}} = 18$, $\sigma = 3$ and $S = 4$.

G.1.4 COP & Q_{vol}

The Coefficient of Performance of the thermodynamic cycle associated to the specific refrigerant molecule is defined at specific evaporator and condenser temperatures. We use the same operating conditions used for AC systems suggested by McLinden et al. [2017] :

- $T_{\text{evap}} = 10^{\circ}\text{C}$
- $T_{\text{cond}} = 40^{\circ}\text{C}$

Empirically, most generated molecules have a high COP. We therefore build a **increasing logistic function** around the Q_{vol} such that the midpoint is about a third of the value of R-410, as done in McLinden et al. [2017] for their filtering. The function is as follows :

$$R(x) = \frac{L}{1 + e^{-k(x - x_0)}}$$

where : $L = 1$ is the upper asymptote, $k = 2$ is the steepness and $x_0 = 2$ is the midpoint. An additional constraint is adding on COP, with the reward being 0 if $COP < 5$.

G.1.5 GWP

Modern day constraints on refrigerants are still somewhat lenient despite being increasingly restrictive. In their detailed screening paper, Kazakov et al. [2012], use a limit of 200, according to recent

european guidelines of refrigerants in the automotive industry in Europe. An ideal objective however, is the use of refrigerants with GWP < 1, i.e with lower global warming potentials than CO₂. This is the case of HFO's which are the current new viable candidates that are emerging in the industry.

We therefore want to push our generation towards GWP < 1, we use a decreasing exponential with a plateau. The reward function for GWP, is given by:

$$R(\text{GWP}) = \begin{cases} 1 & \text{if } \text{GWP} \leq \text{plateau_end} \\ e^{-\text{decay_rate} \cdot (\text{GWP} - \text{plateau_end})} & \text{if } \text{GWP} > \text{plateau_end} \end{cases}$$

With plateau_end = 0.5 and decay_rate = 0.15 This gives a slow decay approaching 0 at 25-30.

G.1.6 Flammability

Flammability is also an important factor for refrigerants. We follow, as Kazakov et al. [2012] and McLinden et al. [2017], the ASHRAE standard, with molecules with a Lower Flammability Limit of > 0.1 kg/m³ considered in a good range and almost non flammable.

We use the same increasing logistic function as for Q_{vol} pressure with $L = 1.0$, $x_0 = 0.1$ and $k = 70$

G.1.7 The diversity reward

The diversity reward aims at penalizing molecules which have appeared often in the recent batches. This seems crucial after initial testing showed lack of exploration and mode collapse at training.

A memory size limit is chosen depending on how far back one wants to penalize generation. The diversity reward stores canonical SMILES and their counts, timestamps, and fingerprints, which are stored in a FIFO manner. The diversity reward works through two different mechanisms:

- A molecule count penalty based on how many times the SMILES has appeared in the memory dictionary.
- A Tanimoto similarity penalty based on fingerprints.

Molecule Count Penalty (Exponential) The core of the penalty system is an exponential function that increases the penalty as the count of a molecule rises. The penalty for a molecule that has appeared c times is calculated with the formula:

$$P(c) = 1 - e^{-k \cdot (c-1)}$$

Where c is the count of how many times the exact same molecule has been seen and k is a rate constant, which is set to 0.5.

Tanimoto Similarity Penalty When a new molecule is not an exact match but is similar to a molecule already in memory, the penalty is scaled by the Tanimoto similarity score. This ensures that very similar molecules are penalized almost as much as exact matches. The formula for the penalty due to a similar molecule is:

$$P_{\text{similar}} = P(c_{\text{similar}}) \cdot S_{\text{Tanimoto}}$$

Where $P(c_{\text{similar}})$ is the penalty calculated above for the similar molecule that is already in memory, which has been seen c_{similar} times and S_{Tanimoto} is the Tanimoto similarity score (a value between 0.0 and 1.0) between the new molecule and the similar molecule from memory.

Final Diversity Reward The system calculates a penalty for the new molecule being an exact match (P_{exact}) and for its similarity to all molecules in memory that are above the similarity threshold ($P_{\text{similar},1}, P_{\text{similar},2}, \dots$). To determine the final penalty, it takes the single harshest (maximum) penalty from all possibilities:

$$P_{\text{max}} = \max(P_{\text{exact}}, P_{\text{similar},1}, P_{\text{similar},2}, \dots)$$

The final diversity reward ($R_{\text{diversity}}$) is then calculated as 1 minus this maximum penalty. The reward is a value between 0.0 (for a highly unoriginal molecule) and 1.0 (for a completely novel molecule).

$$R_{\text{diversity}} = 1 - P_{\text{max}}$$

G.2 GRPO training parameters

- Learning rate : $5e^{-7}$ with linear scheduler
- KL β : 0
- Batch size : 16
- Number of generations per batch : 4
- Optimizer : AdamW with $\beta_1 = 0.9$, $\beta_2 = 0.999$
- Patience 500 with minimum $\delta = 0.1$
- Temperature = 1
- Top p sampling : 0.9

Reward weights :

Table 7: Reward weights used in training

Metric	Weight
COP & Q_{vol}	4.0
T_c	4.0
Molecular length	1.0
GWP	0.5
LFL	0.5

The choice of the weights was made through simple parameter search and tuning, and reflects the difficulty of achieving the different desired properties. Q_{vol} and T_c are harder to optimize for but essential properties and therefore are given higher weights.

G.3 GRPO Inference details

Molecule generation at inference time is generated through 'Top-p' nucleus sampling with a threshold of 0.9, as well as a temperature of 1.

H Property prediction results

H.1 Coefficient of Performance (COP)

Table 8: Prediction Performance of Thermodynamic Properties (Test Set)

Setup	R ²	MAE	RMSE
Tc	0.975	12.4696	21.835
Pc	0.876	148931	568426
ω	0.8749	0.0398	0.0868
Overall Nasa Polynomial Mean	0.9647	237.2239	856.0378
Target: a1	0.9576	0.1223	0.1746
Target: a2	0.9904	0.00072	0.00099
Target: a3	0.9864	7.42×10^{-7}	1.11×10^{-6}
Target: a4	0.9829	3.37×10^{-10}	5.06×10^{-10}
Target: a5	0.9800	5.22×10^{-14}	7.82×10^{-14}
Target: a6	0.9971	1659.9358	2264.8633
Target: a7	0.8588	0.5101	0.7571

To test nasa polynomials and the Peng-Robinson model, we measure errors on the saturation dome construction and the associated COP prediction for the **CoolProp** [Bell and contributors, 2025] fluids. The COP is calculated for condenser and evaporator temperatures of 10°C and 40°C.

Table 9: Thermodynamic Property Prediction on CoolProp

Property	R ²	MAE	RMSE
T _c	0.926	20.69	37.60
P _c	0.769	502134	1451979
ω	0.636	0.057	0.125

Table 10: Saturation Dome RMSE and COP MAE on CoolProp

Metric	N _{samples}	Mean	Median	Std Dev
Dome latent heat RMSE (J)	110	30.9	17.5	41.4
ΔCOP (Ø)	82	0.252	0.12	0.343

H.2 Global Warming Potential (GWP)

H.2.1 Radiative Efficiency (RE)

The SMI-TED finetuned model exhibits good performance on radiative efficiency prediction, as for the other properties.

Table 11: Prediction Performance of radiative efficiency

Setup	Test Set		
	R ²	MAE	RMSE
RE	0.91	0.016	0.025

H.2.2 k(OH) constant

We compare three methods for k(OH) prediction before testing on the database of 220 GWP values :

1. Direct SMI-TED prediction of k(OH) through the McGillen et al. [2020] dataset
2. Using a custom group contribution (GC) method inspired by the Kwok and Atkinson [1995] paper and notes from the U.S. Environmental Protection Agency [2012] AOPWIN software.
3. Training a SMI-TED model on correction factors between the ground truth dataset and the group contribution method

Important remark : the [Kwok and Atkinson, 1995] group contribution method gives a reaction rate of 0 for molecules not reacting according to the 4 main reaction pathways detailed in the method. This usually means the true reaction rate is very slow. We decide to put these values at the arbitrarily small $k_{OH} = 10^{-20}$. For the SMI-TED models, considering the very similar results, training was done on 5 random seeds.

Generalization ability: For the k(OH) constant, more than just good predictions on the McGillen et al. [2020] dataset, we try to consider generalization ability. The group contribution method we implemented is adapted from the AOPWIN software from the EPA (which is not open source) [U.S. Environmental Protection Agency, 2012], which has been tuned on many different molecules and is considered to have good generalization ability. We therefore compare our own group contribution implementation based on open source AOPWIN notes, the SMI-TED k(OH) prediction model the error correction model against the AOPWIN group contribution predicted values (as a form of '*ground truth*') on 72k molecules extracted from PubChem following the methodology in Kazakov et al. [2012]: Molecules with less than 18 atoms and only C, F, H, S, Cl, Br atoms.

We analyze 3 elements : (1) The quality of our group contribution reconstruction (2) the generalization ability of the SMI-TED k(OH) prediction model (3) How much the error correction model ends up deviating from the AOPWIN method.

Results are averaged on 5 random seeds for the SMI-TED based models.

Table 12: Average k(OH) prediction performance comparison with AOPWIN

Setup	Log scale metrics					Original scale metrics		
	R²	<i>R_{pearson}</i>	$\tau_{kendall}$	MALE	10^{MALE}	% <2x	<5x	<10x
Our GC method	0.6623	0.8138	0.6338	0.4094	2.5668	61.28	77.92	87.07
SMI-TED k(OH) Pred.	0.4301	0.6426	0.5004	0.6697	4.7074	30.64	61.55	77.03
Our GC + kOH Correction	0.5849	0.7643	0.5569	0.6052	4.0641	38.67	67.56	79.85

H.2.3 Total GWP prediction

Table 13: GWP prediction performance comparison

Setup	R²	Factor error (10^{RMSE_{log10}})	% <2x	% <5x	% <10x
Our GC method	0.535	11.46	37.73	65.45	75.91
SMI-TED k(OH) Pred.	0.079	12.18	24.16	49.89	66.64
Our GC + kOH Correction	0.434	12.03	34.17	59.84	73.84

Conclusion It seems as if the implemented group contribution method, based on the AOPWIN ('ground truth baseline') software achieves >60% of predictions with an error under a factor 2 of the software, showing that, despite some edge reaction cases not detailed in their work, we managed to rebuild the AOPWIN method pretty well. The group contribution's generalization capabilities seem to enable it have better performance in GWP prediction compared to the direct SMI-TED k(OH) predictors. In fact, the error correction term actually impedes generalization and does not help with prediction.

We therefore deem it best to use the group contribution by itself for k(OH) constant prediction.

H.3 Flammability

Table 14: Prediction Performance Flammability

Setup	Test Set		
	R²	MAE	RMSE
LFL	0.83	0.16	0.61

I Final best molecules

Only non-PFAS molecules are kept. We add the $Q_{vol} > \frac{Q_{vol-R-410A}}{3}$ filter suggested by McLinden et al. [2017] to ensure reasonable system size, not added before because of the limited amount of molecules satisfying this constraint. A list of all best candidates is given.

Table 15: Physicochemical and Environmental Properties of Best Generated Compounds

SMILES	COP	T _c (K)	p _{evap} (kPa)	GWP100	LFL (kg/m ³)	Q _{vol}	Num Atoms	
C(F)N(F)N(F)F	6.71	337.82	1068.09	3.95	0.52	5.34	7	
C1(F)N(F)N1F	7.40	350.56	932.38	0.38	0.42	5.29	6	
N(F)C(F)N(F)F	6.72	338.05	1050.42	0.25	0.54	5.26	7	
C(F)(N(F)F)N(F)F	5.76	327.85	1182.34	0.22	0.57	4.74	8	
N1(F)N(F)N1F	7.72	362.07	753.04	0.42	0.44	4.54	6	
N1(F)C(F)N1F	7.71	361.48	725.39	0.49	0.41	4.45	6	
C1(F)N(F)O1	7.99	372.12	640.08	0.52	0.33	4.31	5	
NC(F)N(F)F	7.81	363.11	579.40	0.27	0.43	3.99	6	
N1(F)C(F)(N(F)F)O1	7.36	355.25	617.32	0.75	0.59	3.60	8	
N(F)OC(F)N(F)F	7.55	359.12	527.72	0.73	0.62	3.50	8	
C(F)(F)C1(F)N(F)O1	7.35	356.30	593.85	0.64	0.57	3.48	8	
C(F)(C1(F)N(F)O1)F	7.44	359.69	553.55	0.64	0.56	3.31	8	
C1C(F)(N(F)F)N1F	7.64	366.07	511.37	0.52	0.44	3.17	8	
N1C(F)(N(F)F)N1F	7.75	368.05	494.60	0.57	0.61	3.16	8	
C1=C(F)O1	8.39	403.27	393.13	0.08	0.14	3.06	4	
C1(F)N(F)N(F)N1F	7.87	373.59	463.07	0.13	0.59	3.04	8	
C1(F)N(F)N(F)O1	8.01	376.75	424.91	0.22	0.50	2.98	7	
C1(F)C(F)(N(F)F)N1F	7.48	364.03	492.50	0.56	0.63	2.98	9	
N1(F)C(F)N(F)N1F	7.88	374.77	445.63	0.16	0.60	2.95	8	
N1(F)N(F)C(F)N1F	7.88	375.36	438.01	0.15	0.60	2.91	8	
N1(F)C(F)N(F)O1	8.02	377.78	406.60	0.23	0.51	2.88	7	
N1(F)C(F)(N(F)F)N1	7.79	370.94	431.26	0.73	0.60	2.83	8	
N(F)C1(F)N(F)N1F	7.91	375.60	417.94	0.53	0.62	2.83	8	
N1C(F)(N(F)F)O1	8.01	376.65	393.64	0.71	0.51	2.80	7	
N1(F)C(F)(N(F)F)N1F	7.65	369.09	437.30	0.58	0.68	2.78	9	
N(F)C(F)N(F)N(F)F	7.69	368.44	409.19	0.23	0.68	2.76	9	
N1(F)N(F)N(F)N1F	7.98	380.43	406.94	0.17	0.59	2.74	8	
C(F)(C1(F)N(F)N1F)F	7.62	368.95	430.36	0.54	0.66	2.72	9	
C(F)(F)C1(F)N(F)N1F	7.62	369.56	428.57	0.54	0.66	2.72	9	
C(F)C1(F)N(F)N1F	7.86	377.53	401.27	0.50	0.47	2.68	8	
N1(F)N(F)C(F)O1	8.10	382.82	347.83	0.25	0.50	2.57	7	
N1(F)N(F)N1C(F)F	7.97	381.52	373.18	0.54	0.61	2.55	8	
N(F)N(F)C(F)N(F)F	7.78	373.30	365.08	0.22	0.67	2.54	9	
N1(F)C(F)N1N(F)F	7.98	381.33	372.24	0.50	0.61	2.54	8	
C(F)N(F)N(F)N(F)F	7.83	375.13	346.73	3.87	0.67	2.46	9	
N1(F)N(F)N(F)O1	8.17	389.04	334.60	0.26	0.51	2.44	7	
NC(F)N(F)N(F)F	8.04	379.35	311.08	0.23	0.62	2.44	8	
C(F)C1(F)N(F)N(F)1	7.88	379.72	357.74	0.57	0.49	2.43	8	
N1(F)N(C(F)F)N1F	8.04	384.80	343.60	0.52	0.61	2.42	8	
N(F)NC(F)N(F)F	8.10	382.07	301.84	0.25	0.63	2.42	8	
C1(F)N(F)N1N(F)F	8.05	384.84	344.78	0.44	0.60	2.39	8	
N1(F)N(F)C1(F)N(F)F	7.79	377.65	357.47	0.52	0.67	2.37	9	
N(F)C1(F)N(F)O1	8.19	388.06	298.65	0.68	0.52	2.28	7	
C(F)(N(F)F)N(F)N(F)F	7.65	374.01	338.49	0.21	0.70	2.27	10	
C1(F)N(F)C1(F)N(F)F	7.79	379.55	335.69	0.53	0.64	2.24	9	
N1(F)C(F)C1(F)N(F)F	7.76	379.54	338.57	0.56	0.64	2.24	9	
C(F)(F)NC(F)N(F)F	7.79	377.70	306.51	0.23	0.67	2.24	9	
N1(F)N(F)N1N(F)F	8.09	389.17	321.74	0.46	0.62	2.24	8	
N1(C(F)F)N(F)N1F	8.11	389.52	310.28	0.54	0.60	2.24	8	
C1(F)N(F)N1C(F)F	8.05	388.19	309.18	0.50	0.44	2.20	8	
C(F)C(F)N(F)F	8.07	389.43	286.64	3.54	0.40	2.16	7	
CC(F)N(F)F	8.16	390.68	274.20	1.88	0.23	2.13	6	
C(F)(F)OC(F)N(F)F	7.71	377.07	295.55	0.69	0.67	2.12	9	
C(F)(N(F)N(F)F)N(F)F	7.73	377.42	300.02	0.21	0.70	2.08	10	
C(F)(N(F)F)C(F)N(F)F	7.59	377.37	305.02	5.07	0.73	2.07	10	
C(F)(F)OC1(F)N(F)O1	7.84	378.21	283.25	0.71	0.66	2.07	9	
C1=C(F)N1F	8.49	416.10	253.52	0.14	0.19	2.07	5	
N1(F)OC(F)N1F	8.31	399.74	261.63	0.20	0.50	2.05	7	
N1(F)OC(F)(N(F)F)O1	7.93	382.48	28	275.23	0.32	0.68	2.02	9
C1(F)OC(F)(N(F)F)N1F	7.80	380.46	282.67	0.22	0.73	2.00	10	
C(F)(C(F)N(F)F)N(F)F	7.66	379.58	286.00	4.87	0.73	1.99	10	