


Reinforcement Learning Guided
de novo generation of
Macrocycles by Repurposing the
Linkinvent Prior

Implementation




Desired Input



ClC1=CC=C(O[*])C(NC(NC2=NC(O[*])=CN=C2)=O)=C1MAC



Linkinvent Format




NC1=NC(O[*])=CN=C1|ClC1=CC=C(C(NC(Cl)=O)=C1)O[*]

- For *de novo* macrocycle generation, we want to link two parts of a molecule designated by dummy atoms “*”
- Linkinvent takes input as two fragments joined by a pipe “|” with attachment points designated by dummy atoms “*”
- This implementation gets around this by automating the fragmentation of macrocycle parent scaffold
 - Generates input in reinvent format
- Handling of macrocycles designated by the addition of a MAC flag at the end of the input SMILES


Implementation



ClC1=CC=C(O[*])C(NC(NC2=NC(O[*])=CN=C2)=O)=C1MAC




- SMILES input as scaffold with dummy atoms
- Get as RDKit mol object



- Detect attachment points
- Calculate shortest path between attachment points
- Limit fragmentation candidates to single bonds not in rings

OC1CC=NC(NC(=O)N[S+])=N1[]Oc1ccc(Cl)cc1[S+]

([#17:1]-[#6:2]1:[#6:3]:[#6:4]:[#6:5](-[#8:6]):[#6:7]:[#6:8]:1)-[#16+:9].[#7:10](-[#6:11](-[#7:12]-[#6:13])=[#8:14])-[#16+:15])>>[#17:1]-[#6:2]1:[#6:3]:[#6:4]:[#6:5](-[#8:6]):[#6:7]:[#6:8]:1)[#7:10]-[#6:11](-[#7:12]-[#6:13])=[#8:14]




- Iterate through fragmentation candidates
- Break bond and add S⁺ as a handle for reaction SMIRKS (least likely Linkinvent token to cause issues)
- Get atomic environment around S⁺ as SMARTS with radius 4
- If SMARTS patterns are unique within the molecule – break
- Reaction stored as “reaction” – logic detects if “reaction” is not None to enable reconnection as postprocessing

Implementation



*OC1CC=NC(NC(=O)N[s+])=N1|*Oc1ccc(Cl)cc1[s+]



- Generated Linkinvent input is processed as normal Linkinvent run

- Agent generates linked fragments as SMILES
- Fragments are recombined to form macrocycle **prior** to scoring



- Scores are generated based on the properties of the entire macrocycle
- Enables RL based on macrocycle properties

Implementation - Limitations

- Will not work for fully symmetrical molecules where the only available single molecule produces two identical fragments
- Will not work for molecules where all bonds are involved in ring systems
- Can only take one line input
 - Only stores one reaction
 - Could save reactions in a dictionary with scaffold strings as keys


Test 1 – Literature Macrocycle Generation

- Identical to tests from Macformer paper
 - Transformer based macrocycle model
 - RL not implemented
- AIM – can we train the Linkinvent agent to generate molecules similar to 5a
 - Literature checkpoint 1 kinase inhibitor for cancer indications
 - Input scaffold with and without the ether attachment points
 - Compare to untrained Linkinvent prior
- RL scoring function to target:
 - Number of rings = 3
 - Linker graph length = 7/5 (without/with ether specified, respectively)
 - Linker num HBA = 2/0 (without/with ether specified)
 - TPSA < 120
- Sampling method:
 - Sample from trained agent 10 times, remove duplicates and unlinked/unmacrocyclised molecules (*/[S+] present in SMILES)




compd	ring size	X	Chk1 inhibition (IC ₅₀ , nM)
5a	n = 2 15-member	H	10

Ether not specified





Ether specified



Test 1 – Literature Macrocycle Generation

Target Compound



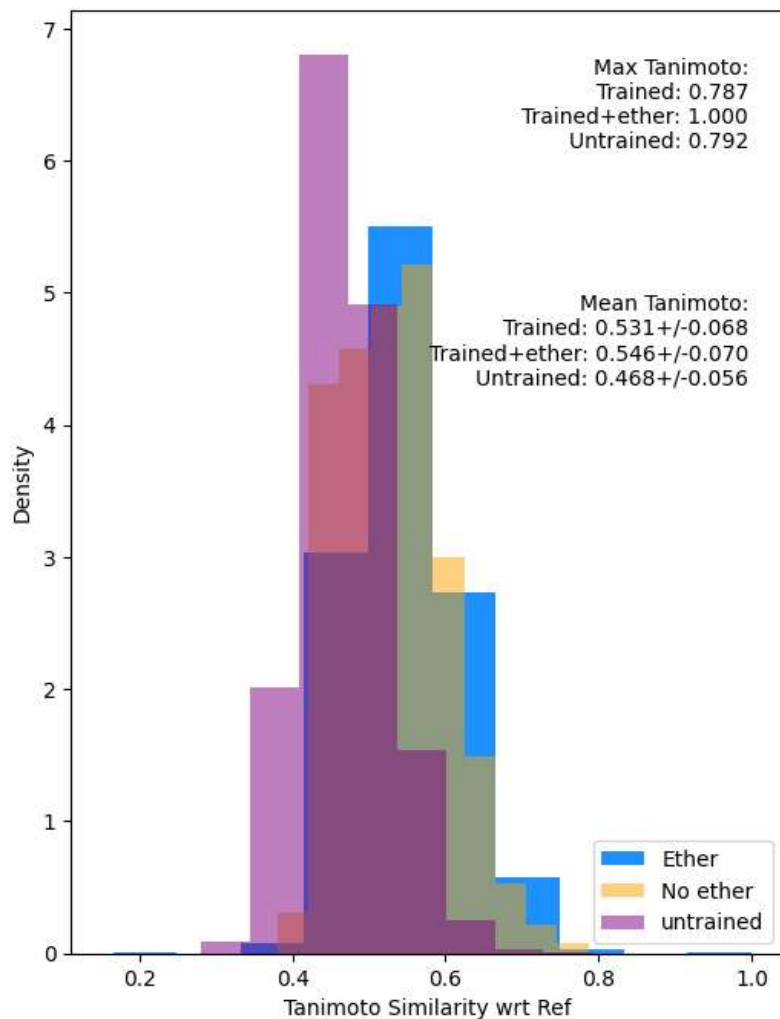

Assessed effects on distributions on Physchem properties and compare to untrained Linkinvent prior

- Properties directly assessed by the scoring function: MolWt (indirectly by linker length and num_rings), number of rings, TPSA, and HBA
- MolWt and nrings distributions when ether is and isn't specified are comparable and both lower average than untrained prior
- Num HBA distributions also shift when ether is specified in RL – but not when ether is not specified. Same trend is seen in TPSA.
- While not directly assessed, distribution of LogP also shifts lower compared to untrained prior
- HBD shifts lower when ether is specified but not when ether isn't specified
 - Suggestive that training agent to produce linkers with zero HBA is easier than training an agent to make a linker with exactly 2
- In all cases, RL with ether specified pushes the mean of property distributions towards the values of the target compound (vertical red dashed lines)

Supportive of the use of RL to tune macrocycle properties

Test 1 – Literature Macrocycle Generation


Target Compound




- **Assessed diversity of generated macrocycles after RL compared to the untrained prior – based on tanimoto similarity of ECFP fingerprints radius 3**
- The target compound was only found when the ether was specified (max tanimoto = 1). Sampling from the reinvent prior afforded a marginally more similar compound than when the ether was not specified (tanimoto = 0.792 vs 0.787)
- By assessing mean tanimoto similarity, we can evaluate if our scoring function is generating macrocycles closer to the target
 - Mean values for ether specified and unspecified are significantly higher than the mean of the untrained compounds tanimoto similarities
 - Mean value higher when ether is specified – but only by 0.007. this is expected when more of the final molecule is specified in the input
 - Indicates we are generating macrocycles closer to the target compound

Test 1 – Literature Macrocycle Generation


Target Compound




Trained Agent – no ether specified



Trained Agent – ether specified



Untrained Linkinvent Prior



- We can visually inspect the most similar compounds to the target
- All methods generate macrocycles using ethers as the attachment point
- All macrocycles generated without the ether specified contain HBA or HBD in the linker – no unsubstituted alkyl chains
- When ether specified the agent generates the target compound and different alkyl chain length analogues
- All closely related compounds generated by the Linkinvent prior have a ring in the linker
 - Suggests this is a key physchem property directed by RL