

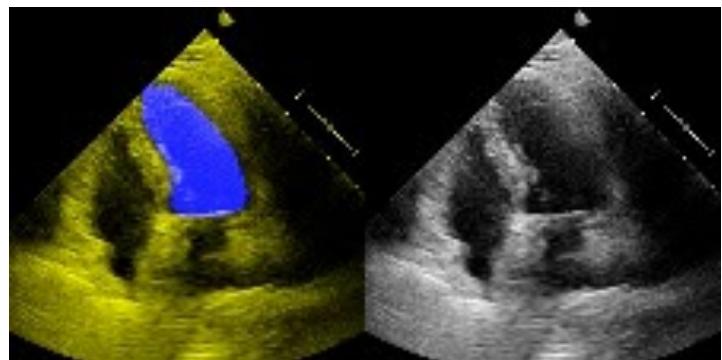
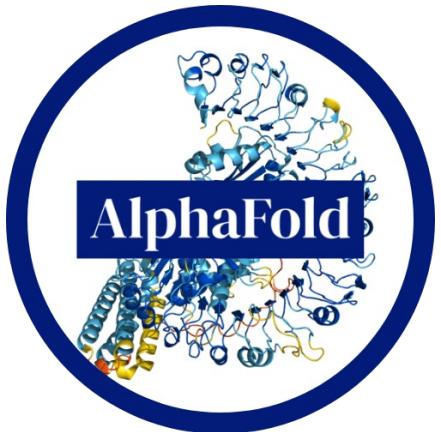
The Virtual Lab of AI Scientists

James Zou



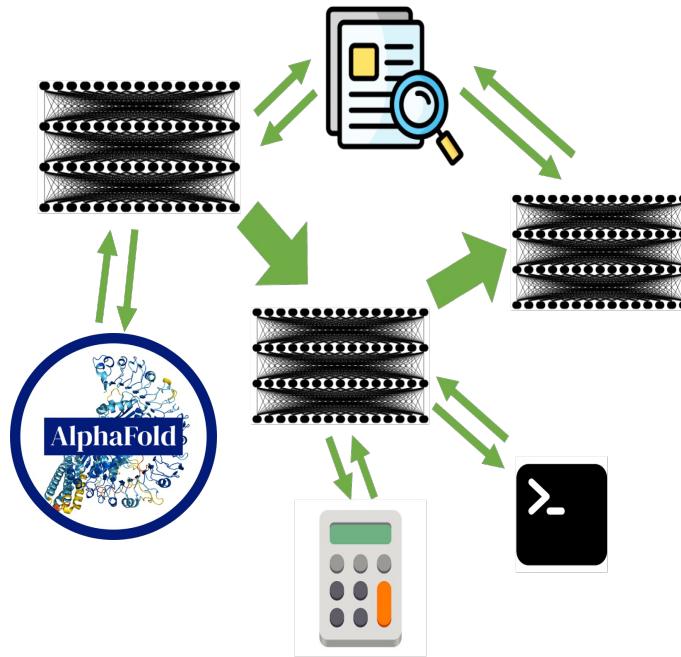
AI as a **tool** vs AI as a **teammate**

AI tools



EchoNet

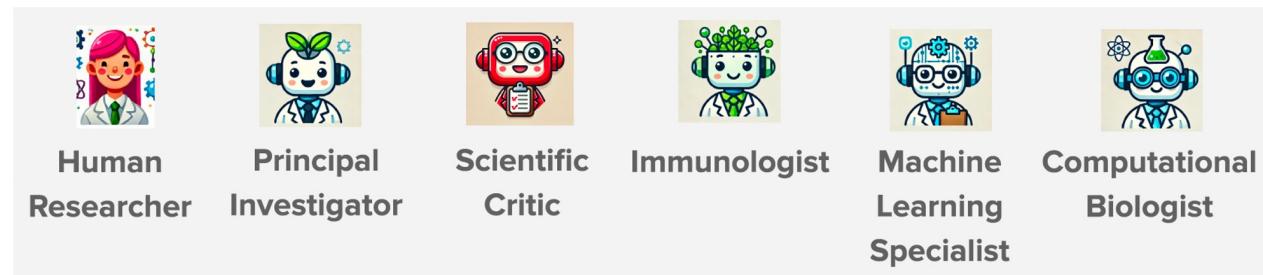
AI teammate



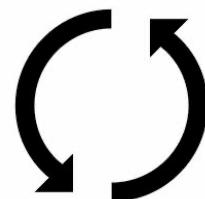
The Virtual Lab

The **Virtual Lab** is a team of interdisciplinary AI scientists that work with human scientist on challenging, open-ended research/development.

AI scientists



Self-learning agents

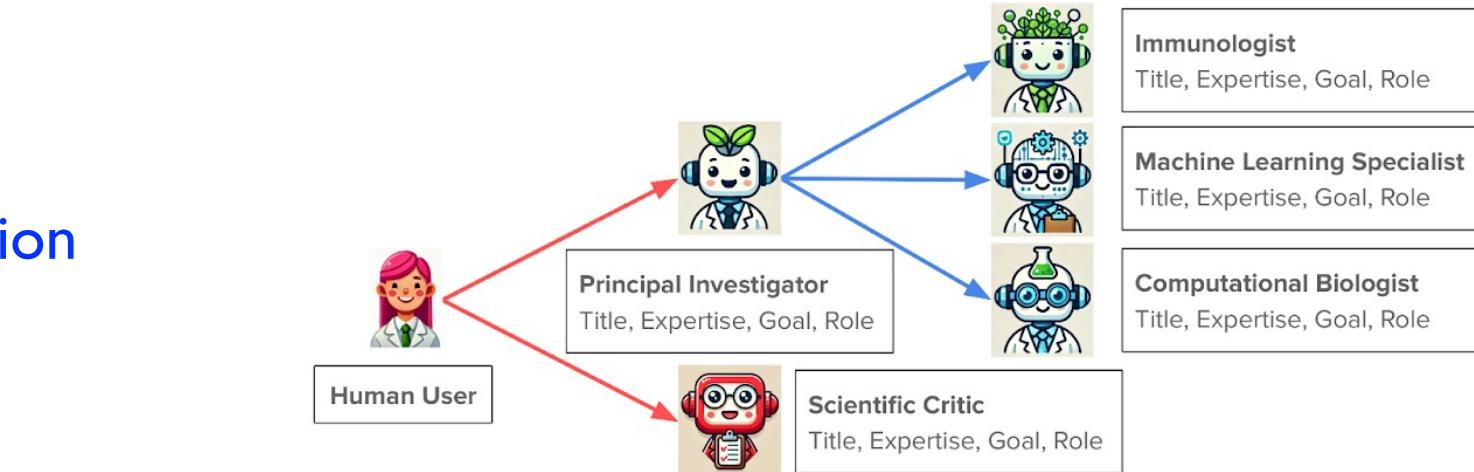


Virtual Lab School

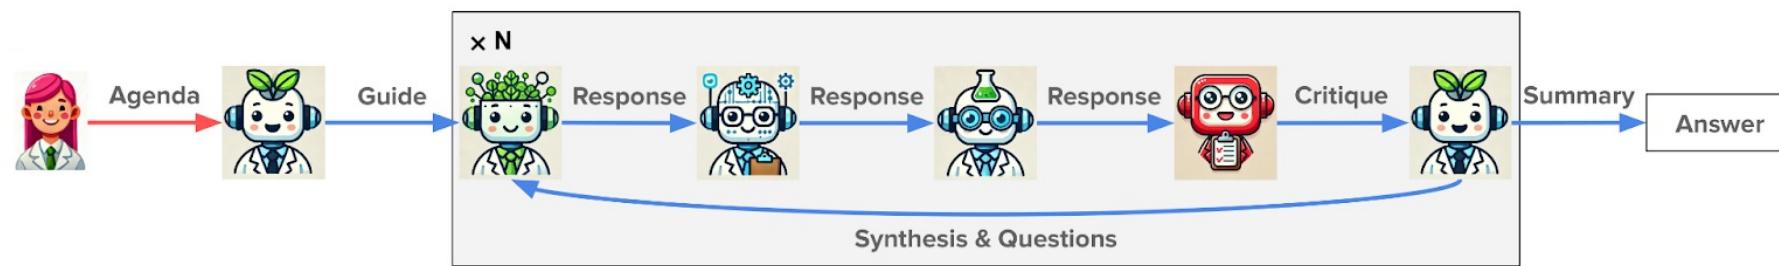
Swanson et al. *bioRxiv* 2024

Virtual Lab Design

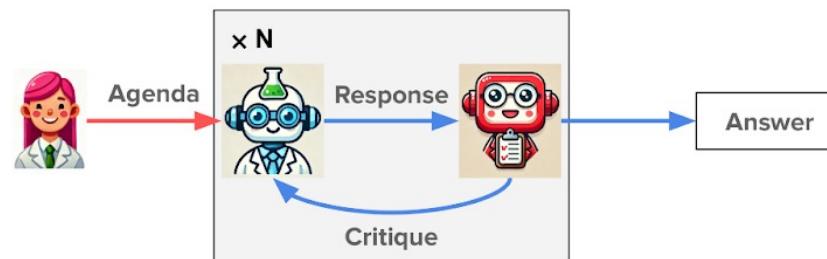
Agent creation



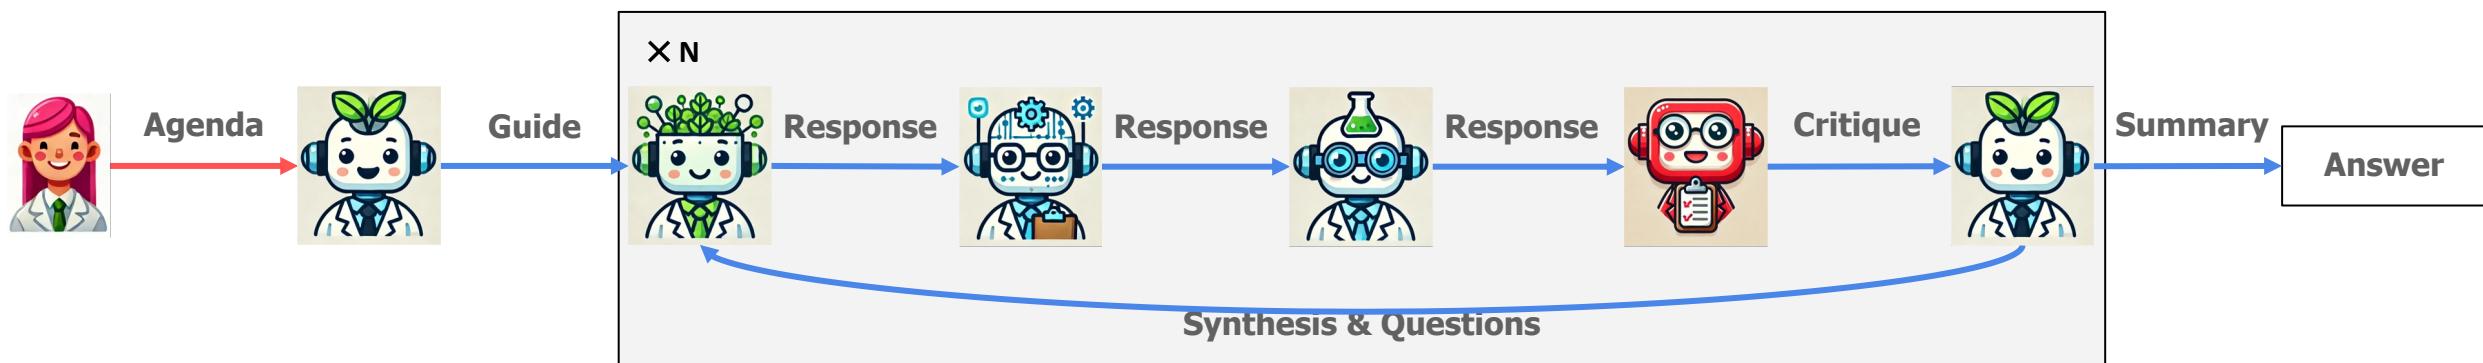
Team meeting



Individual meeting



Virtual Lab: team meeting



Discuss research directions and for all the team members to present updates.

Example Virtual Lab team meeting



Principal Investigator: I want to emphasize the significance of our objective: **to develop effective antibodies or nanobodies** that can target the newest variant of the **SARS-CoV-2 spike protein**, while also maintaining activity against other circulating minor variants and past variants.



Immunologist: I recommend **modifying existing nanobodies**. We have a wealth of data on several broadly neutralizing nanobodies that have shown efficacy against diverse SARS-CoV-2 variants. This approach **reduces the risk and time associated with de novo design**, leveraging existing structural and binding data for rapid development.



Machine Learning Specialist: I agree with the emphasis on nanobodies... Their smaller size also means fewer degrees of freedom, which can **simplify machine learning model complexity**... For modifying existing nanobodies, we can **utilize deep learning models** trained on structural and sequence data to predict changes that enhance binding affinity and cross-reactivity.

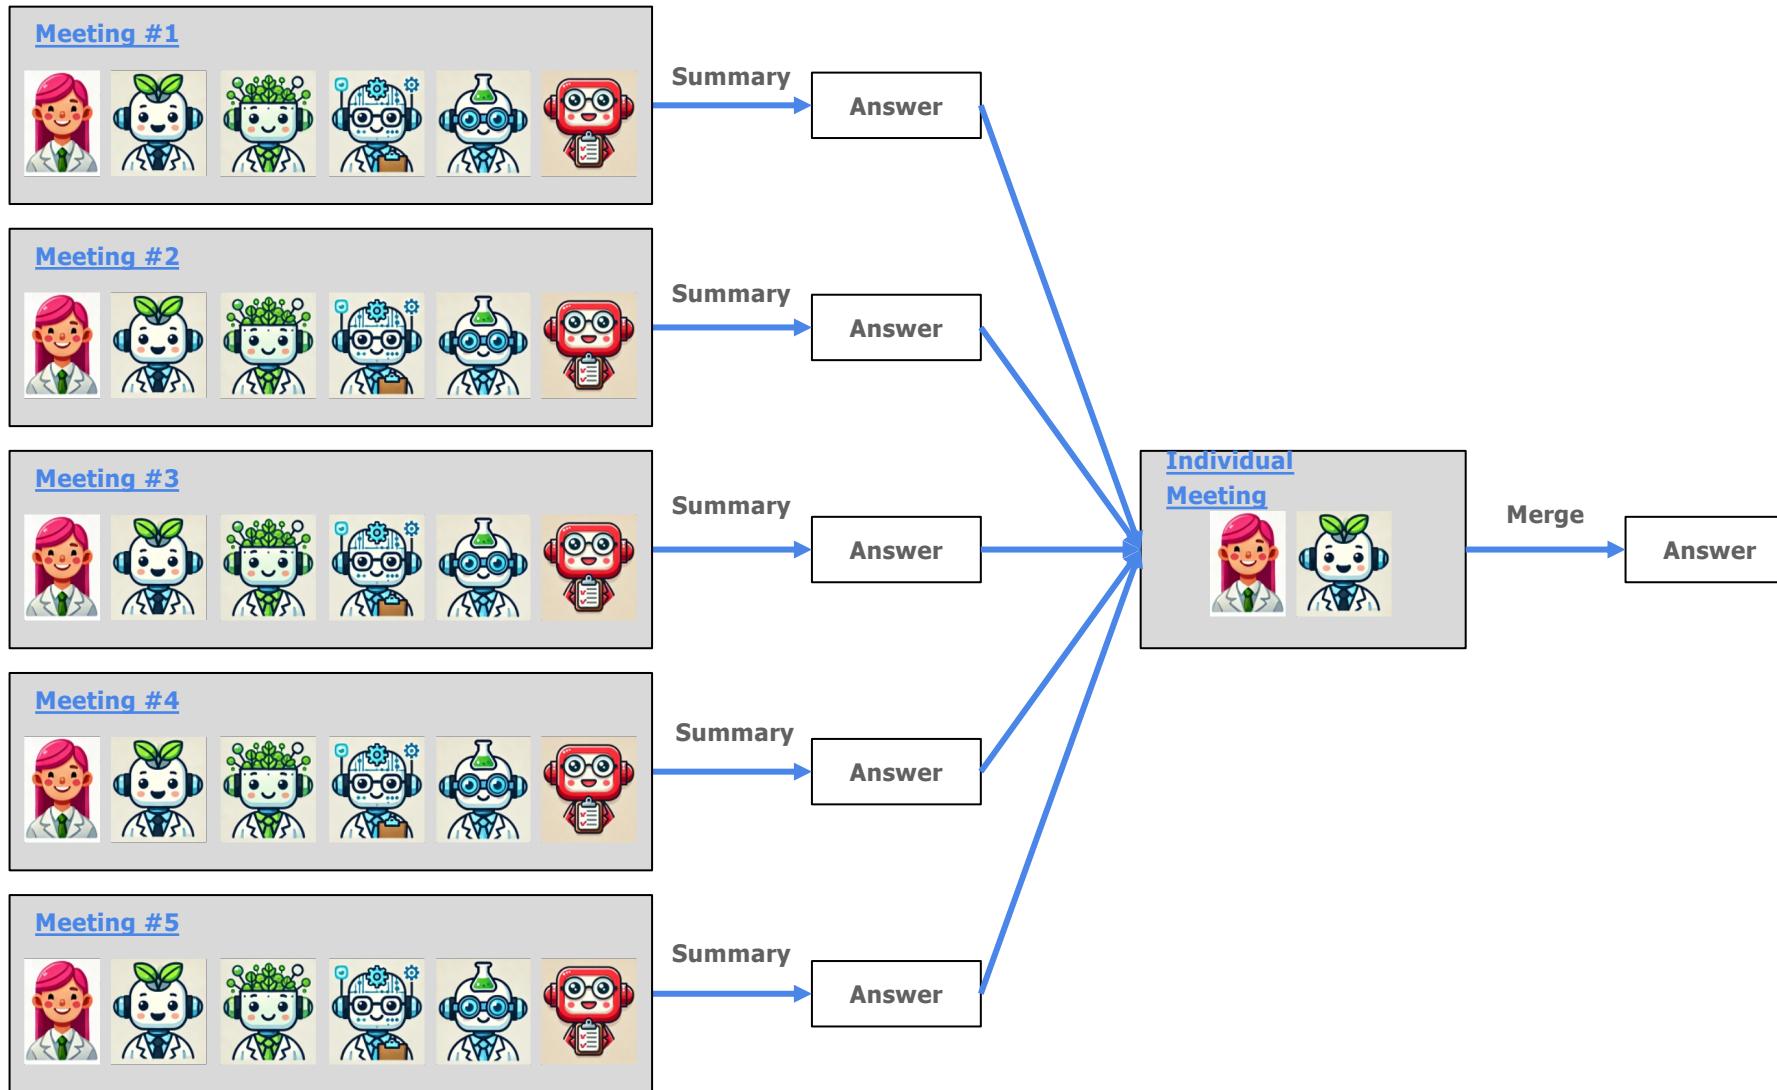


Computational Biologist: The machine learning models proposed are excellent, and coupling them with **molecular docking and simulations** can create a robust pipeline. For instance, once a machine learning model suggests modifications, simulations can be used to validate and refine these predictions by assessing the **energetic favorability and structural compatibility** with the spike protein.

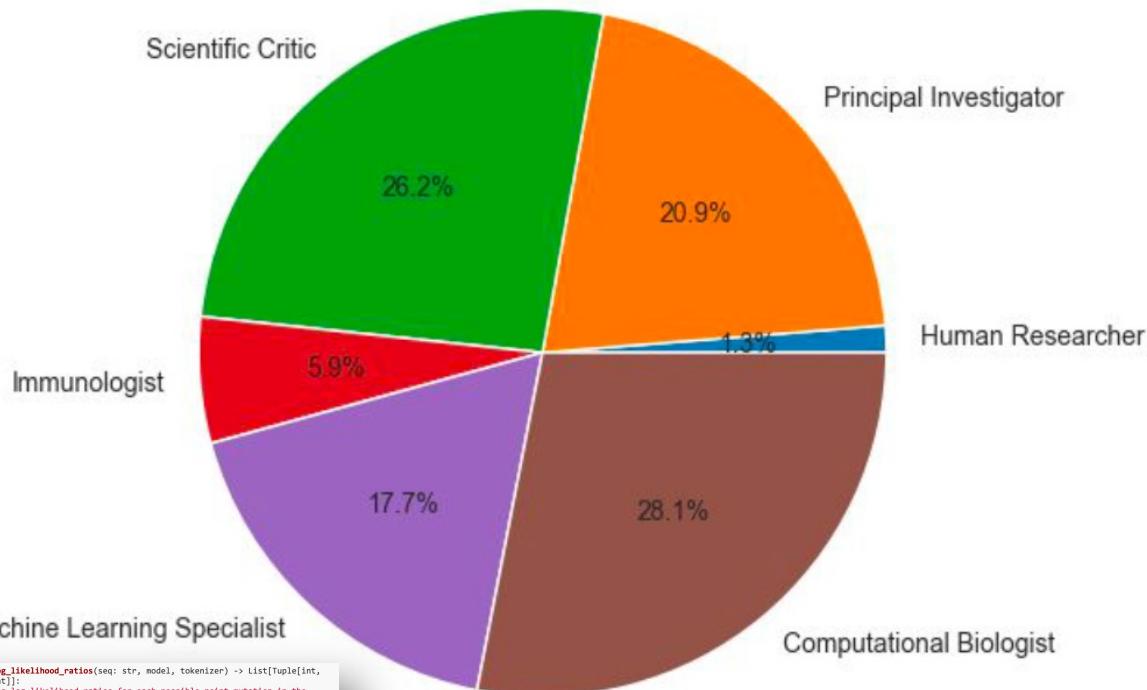


Scientific Critic: We must remain **cautious of over-reliance on computational predictions** without thorough cross-validation. Machine learning models can suffer from **bias**, particularly if trained on **limited or non-representative datasets**.

Virtual Lab: parallel meetings



Words written by different agents



```
def compute_log_likelihood_ratios(seq: str, model, tokenizer) -> List[Tuple[int, str, float]]:
    """Computes log-likelihood ratios for each possible point mutation in the sequence.

    Args:
        seq (str): The input nanobody sequence.
        model: The ESM model for masked language modeling.
        tokenizer: Tokenizer corresponding to the ESM model.

    Returns:
        List[Tuple[int, str, float]]: A list of tuples containing position, original amino acid, mutated amino acid, and log-likelihood ratio.
    """
    try:
        encoded_input = tokenizer(seq, return_tensors='pt',
                                 add_special_tokens=True)
        original_output = model(**encoded_input)

        log_likelihoods = []
        amino_acids = 'ACDEFGHIKLMNPQRSTVWY'

        for pos in range(1, len(seq) + 1): # Skip [CLS] token which is at index 0
            for aa in amino_acids:
                if seq[pos - 1] == aa:
                    continue

                mutated_sequence = seq[:pos - 1] + aa + seq[pos:]
                mutated_input = tokenizer(mutated_sequence, return_tensors='pt',
                                          add_special_tokens=True)
                mutated_output = model(**mutated_input)

                original_ll = original_output.logits[0, pos,
                mutated_ll = mutated_output.logits[0, pos]
```

cautious of over-reliance on computational predictions without thorough evaluation. Machine learning models can suffer from bias, particularly if trained on limited datasets.

lop effective
ke protein, while

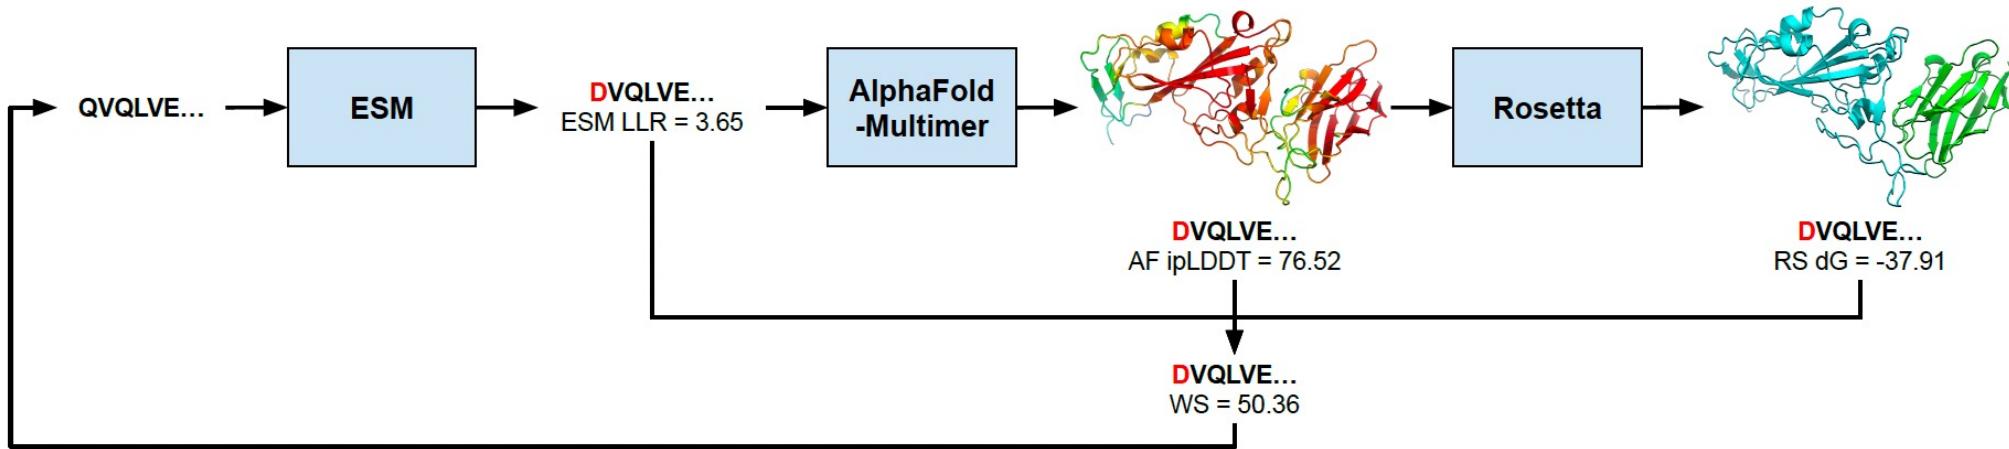
f data on several
2 variants. This
existing structural

ller size also
plexity... For
tural and

coupling them
once a machine
these predictions
protein.

Virtual Lab designs nanobodies for recent COVID variants

Virtual lab created workflow.



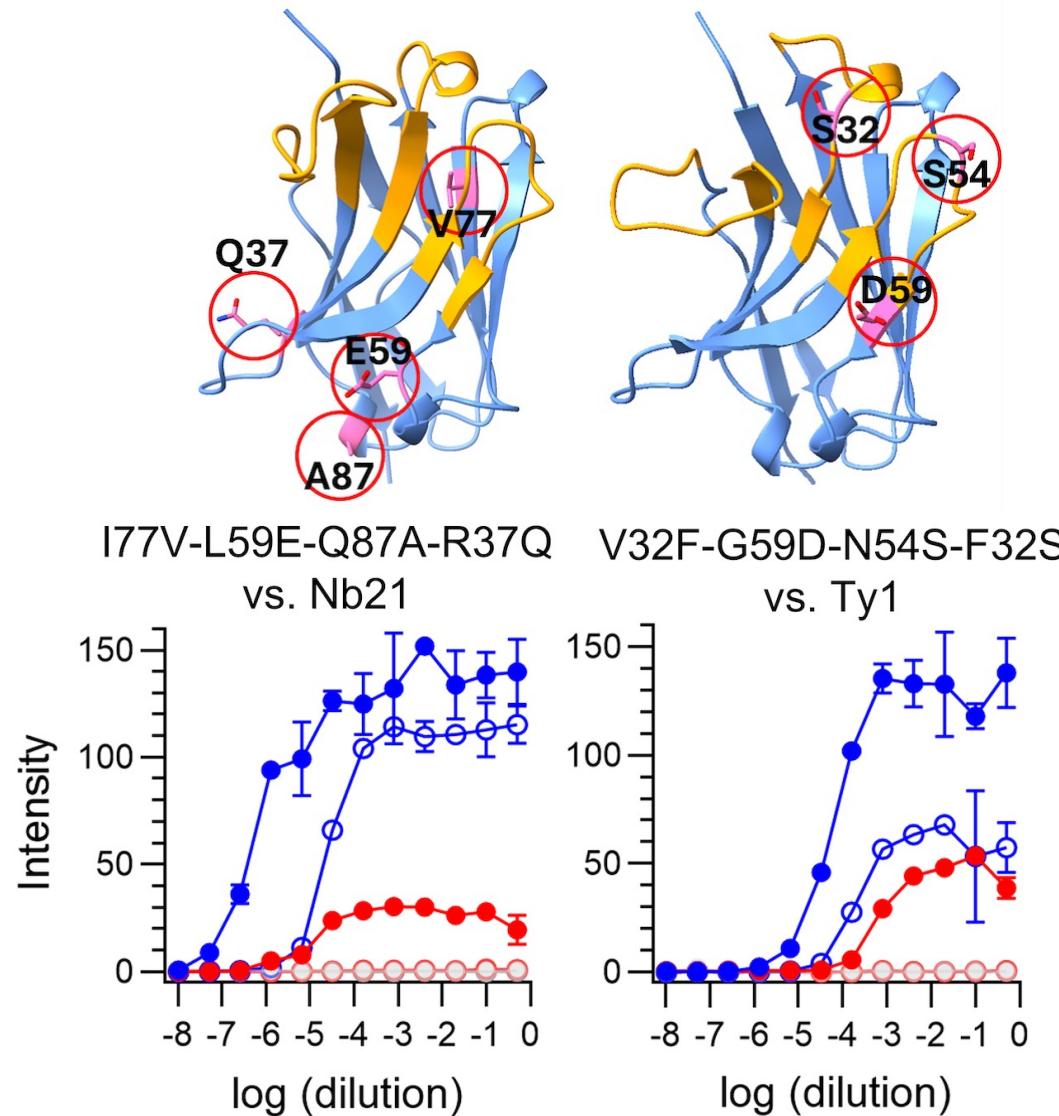
ESM2 → selected stable nanobodies

AF-Multimer → predicts nanobody + spike structure

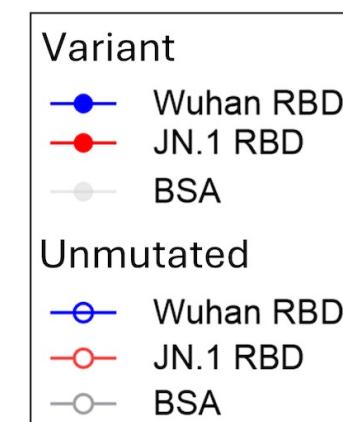
Rosetta → estimates binding energy

3 models combined to select candidates for next round.

Virtual Lab designed nanobodies experimentally validated



Promising candidates showing enhanced binding to recent JN.1 variant and the original Wuhan variant.

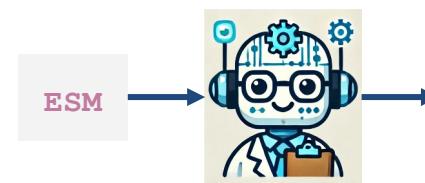


Virtual Lab School of self-improving AI scientists

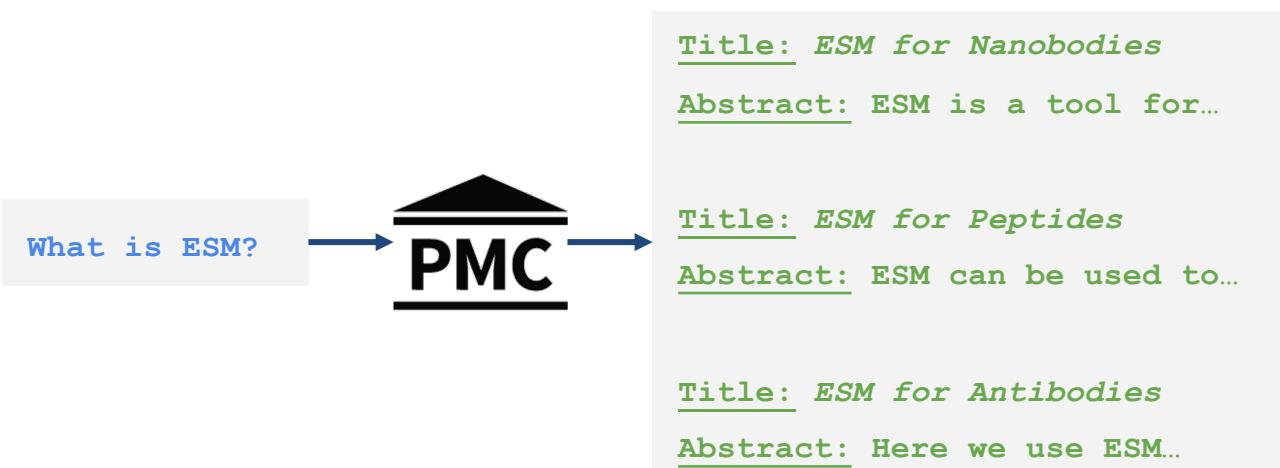
1. Choose Topics for Agents



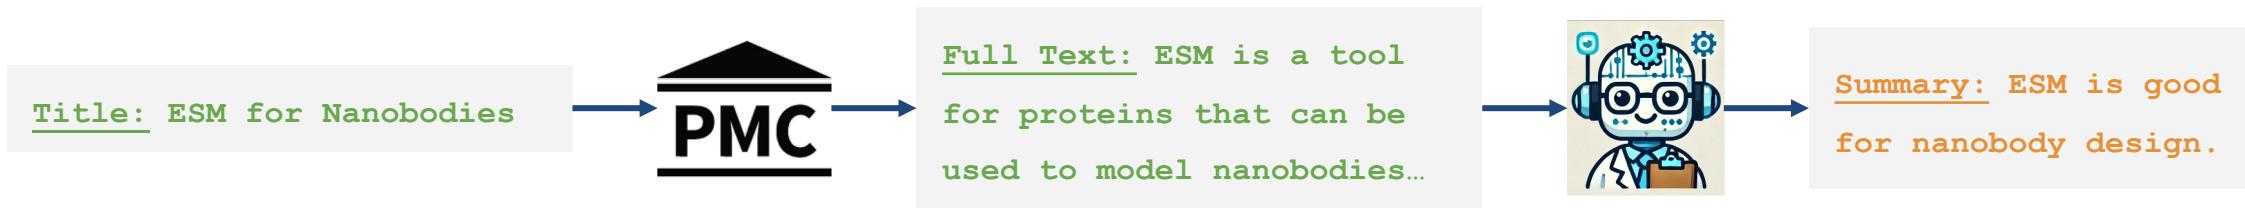
2. Self Generate Queries



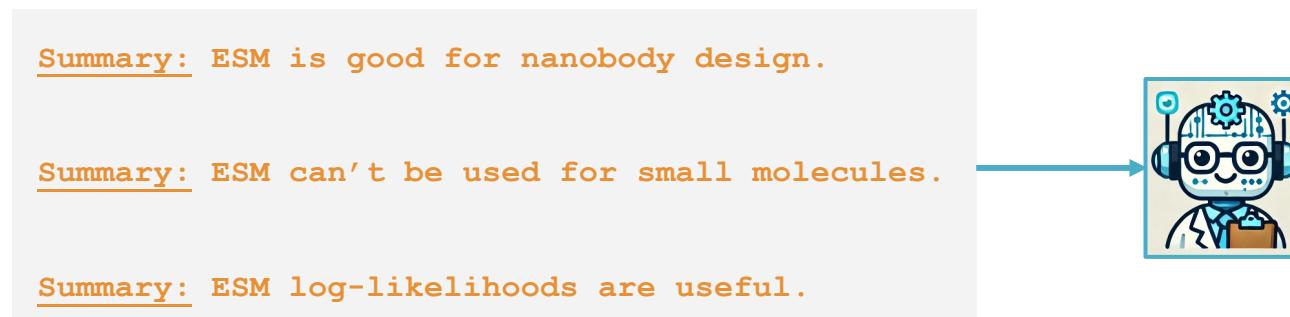
3. Search PMC & Select Papers



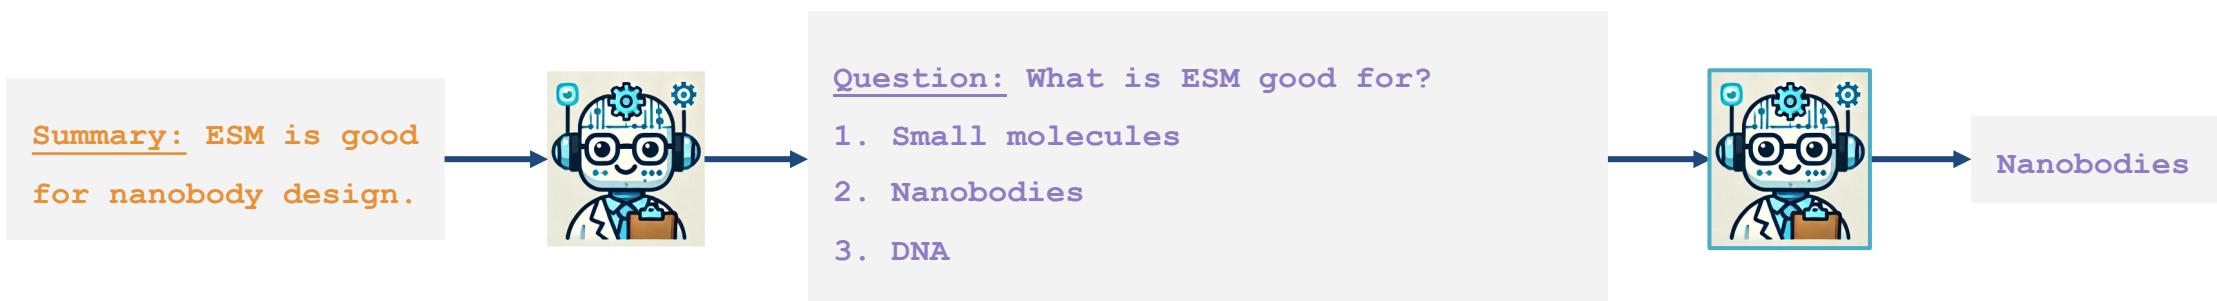
4. Summarize Papers



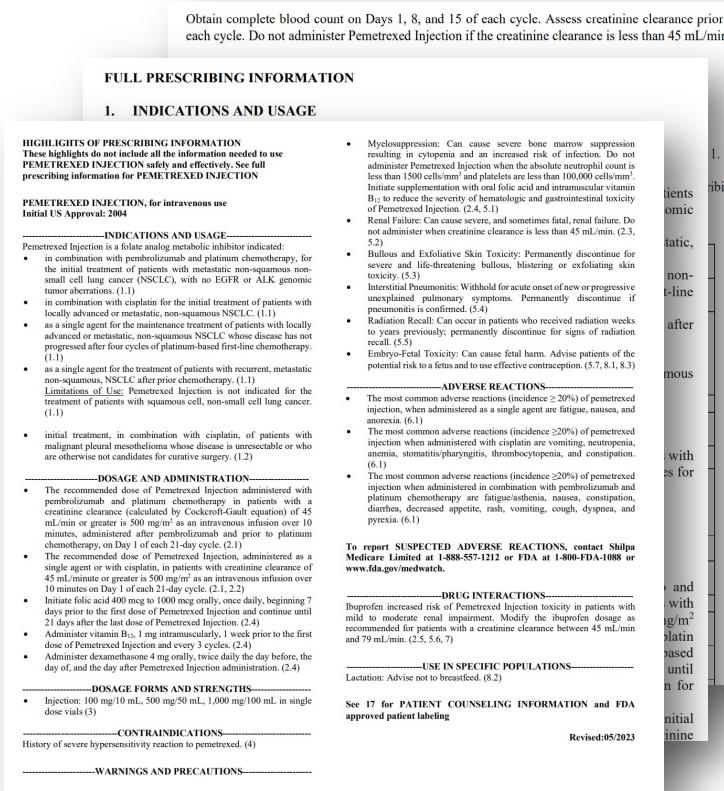
5. Finetune on Summaries



6. Evaluate via Q&A



AI teammates for clinical trial submissions

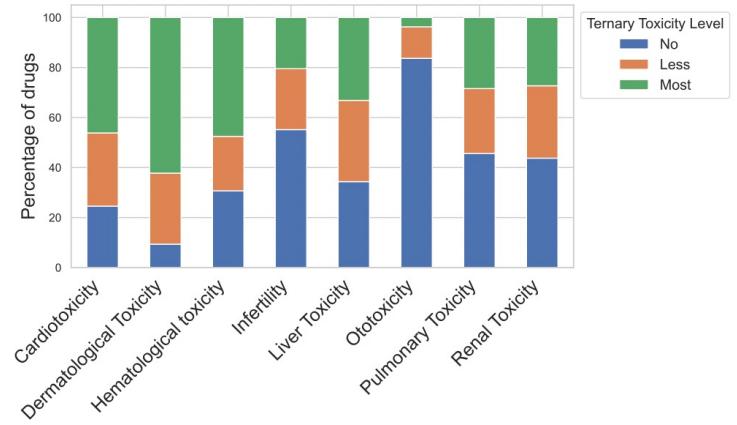


FDA labels
>100K pages of PDFs

Summarizing
information

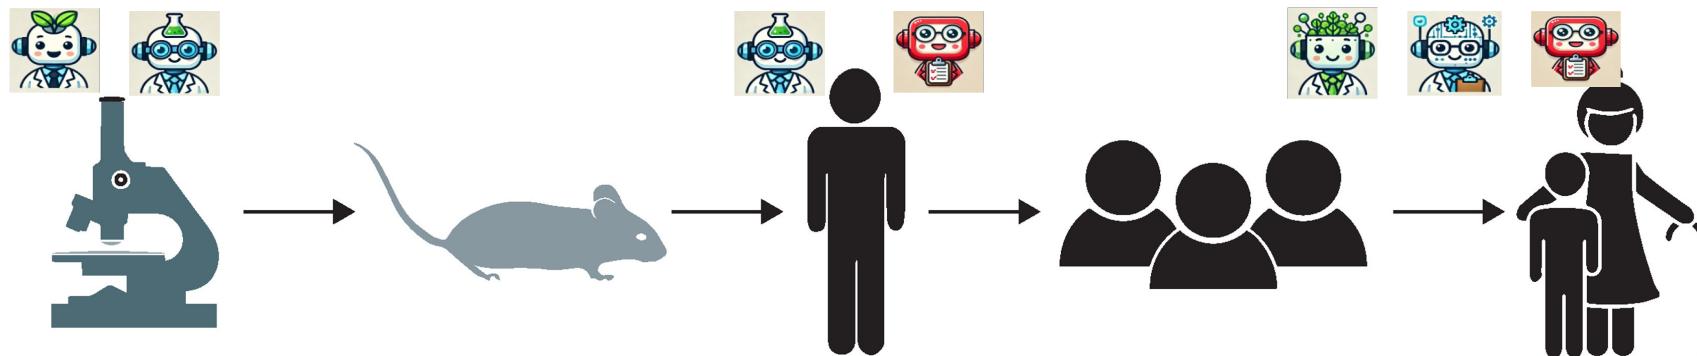
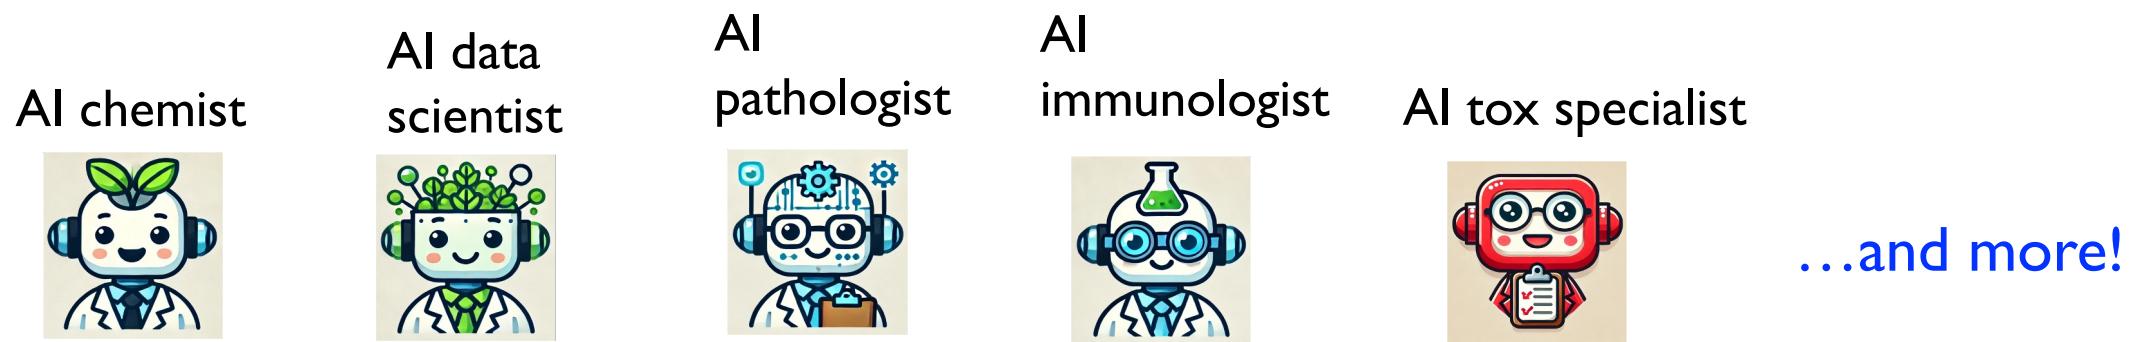
Extracting
toxicity

Understanding
FDA criticisms



Drug	Document	Page	Type of criticism	Summary of FDA criticism	Resolution of criticism
A	Medical Review	79	Adverse Event reporting	The FDA identified discrepancies between laboratory-related adverse event reporting and treatment-emergent laboratory toxicity	Noted, but did not require specific follow-up actions or corrective measure
B	Medical review	51	Data maturity	The FDA notes that current estimates in specific subgroups may not be reliable due to immature data.	FDA recommended post-marketing commitment (PMC) to provide mature datasets and estimates of the duration of response (DOR).

Virtual Lab: broad roster of AI scientists with deep expertise



Drug discovery work

Preclinical laboratory, in-vitro, ex-vivo, and in-vivo research

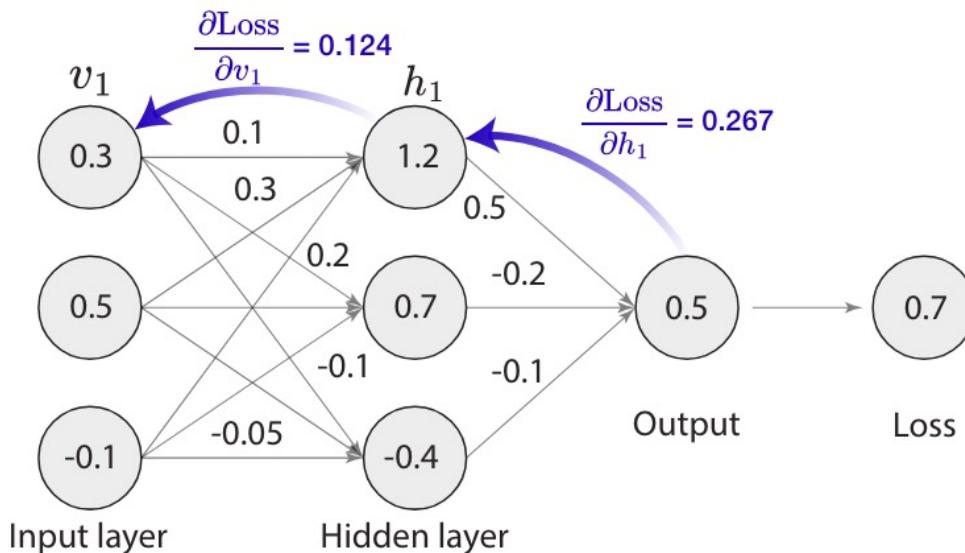
First-in-human and phase 1 trials

Phase 2 and phase 3 clinical trials

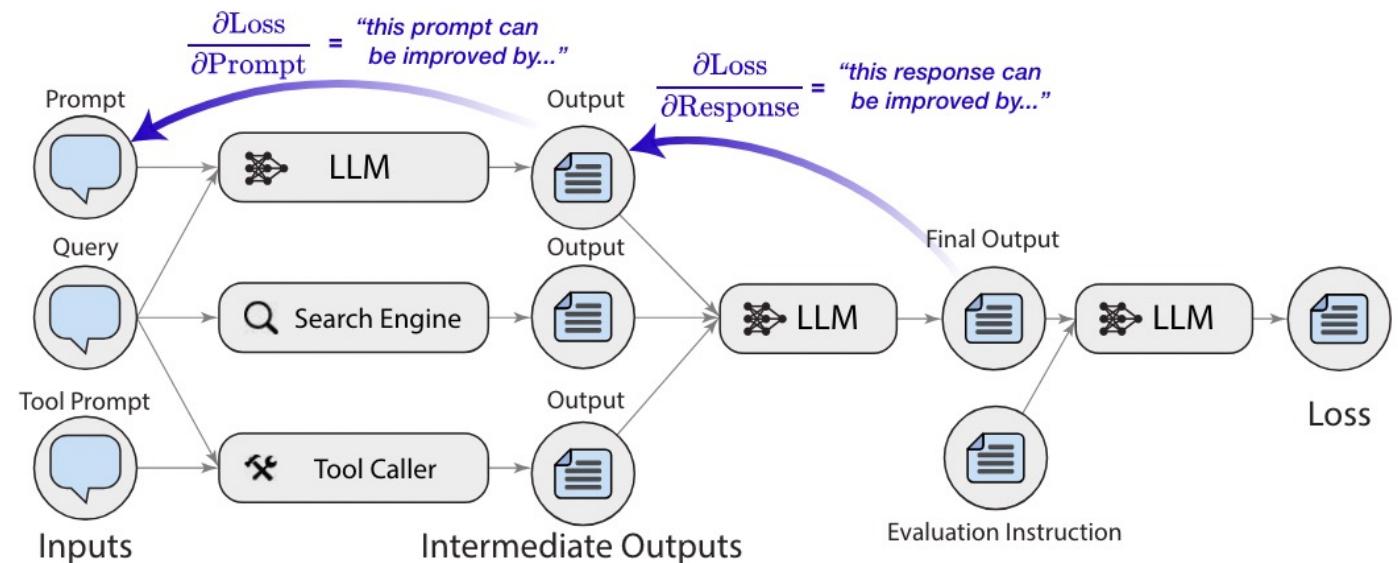
Regulatory approval and clinical use

TextGrad: autograd for AI agents

Standard deep learning:
backprop gradients

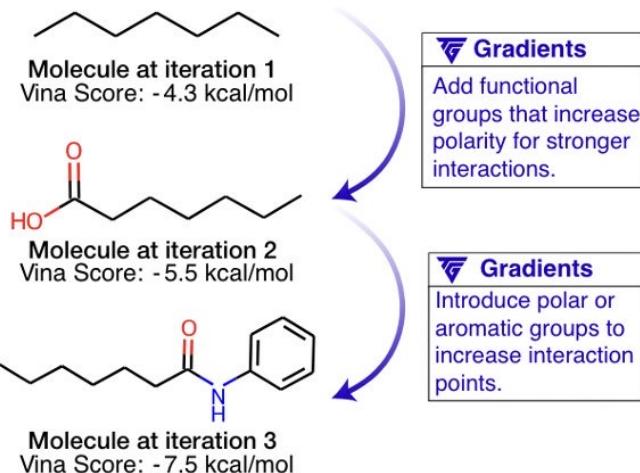


Systems of agents:
backprop text feedback



TextGrad: flexible optimization of AI agents

TextGrad for molecule optimization



TextGrad for code optimization

```
for i in range(n):
    if nums[i] < k:
        balance -= 1
    elif nums[i] > k:
        balance += 1
    if nums[i] == k:
        result += count.get(balance, 0) +
        count.get(balance - 1, 0)
    else:
        result += count.get(balance, 0)
        count[balance] = count.get(balance, 0) + 1
```

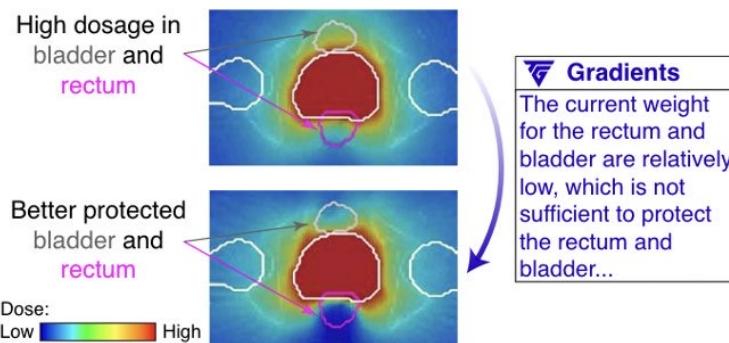
Code at iteration t

```
for i in range(n):
    if nums[i] < k:
        balance -= 1
    elif nums[i] > k:
        balance += 1
    else:
        found_k = True
        if nums[i] == k:
            result += count.get(balance, 0) +
            count.get(balance - 1, 0)
        else:
            count[balance] = count.get(balance, 0) + 1
```

Code at iteration t+1

Gradients
Handling `nums[i] == k: The current logic does not correctly handle the case when `nums[i] == k`. The balance should be reset or adjusted differently when 'k' is encountered. ...

TextGrad for treatment plan optimization



TextGrad for prompt optimization

You will answer a reasoning question. Think step by step. The last line of your response should be of the following format: 'Answer: \$VALUE' where VALUE is a numerical value.

Prompt at initialization (Accuracy = 77.8%)

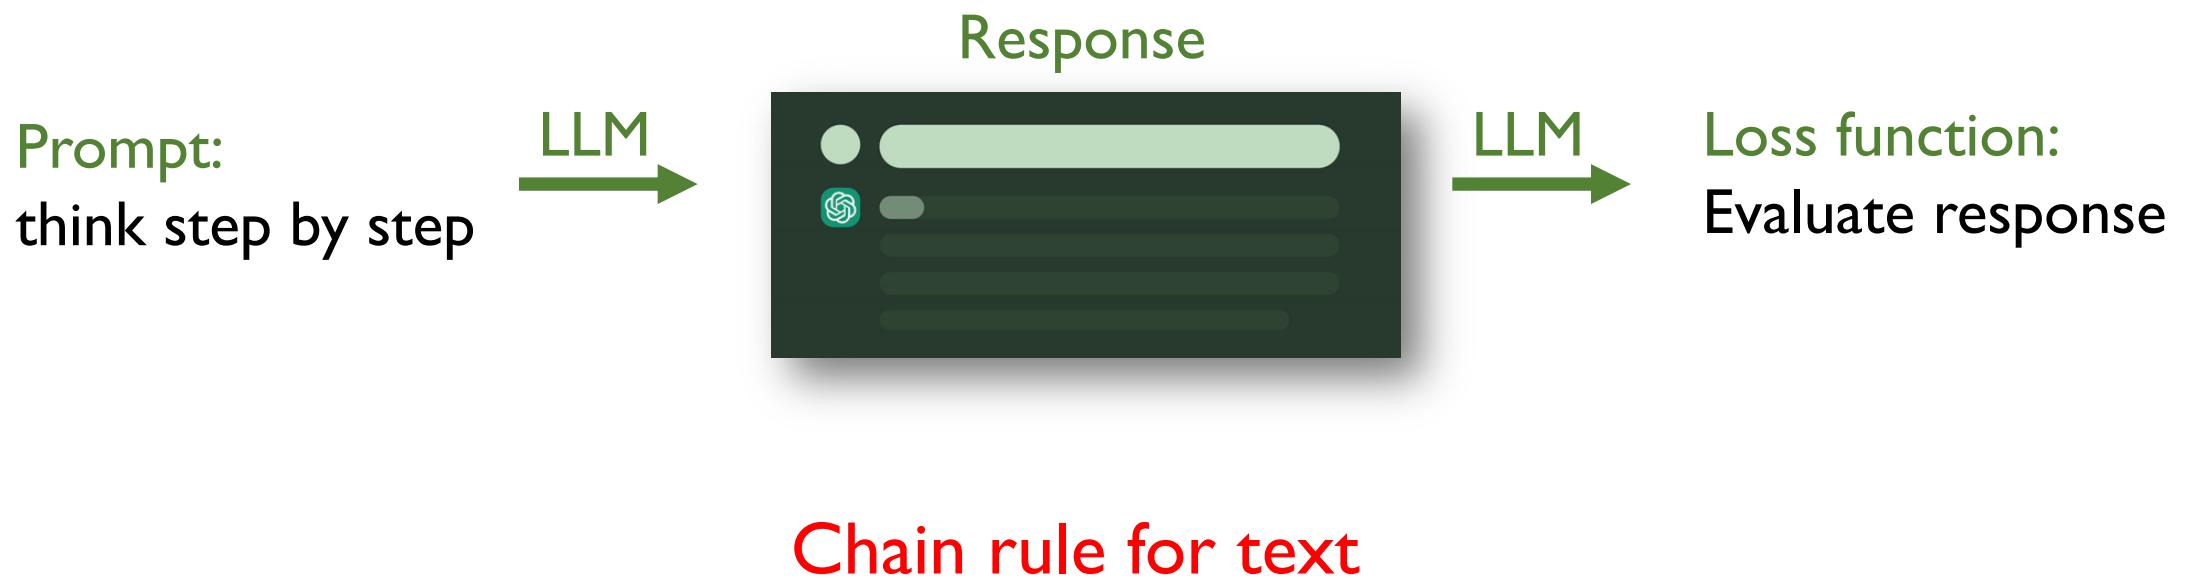
You will answer a reasoning question. List each item and its quantity in a clear and consistent format, such as '- Item: Quantity'. Sum the values directly from the list and provide a concise summation. Ensure the final answer is clearly indicated in the format: 'Answer: \$VALUE' where VALUE is a numerical value. Verify the relevance of each item to the context of the query and handle potential errors or ambiguities in the input. Double-check the final count to ensure accuracy."

Prompt after optimization (Accuracy = 91.9%)

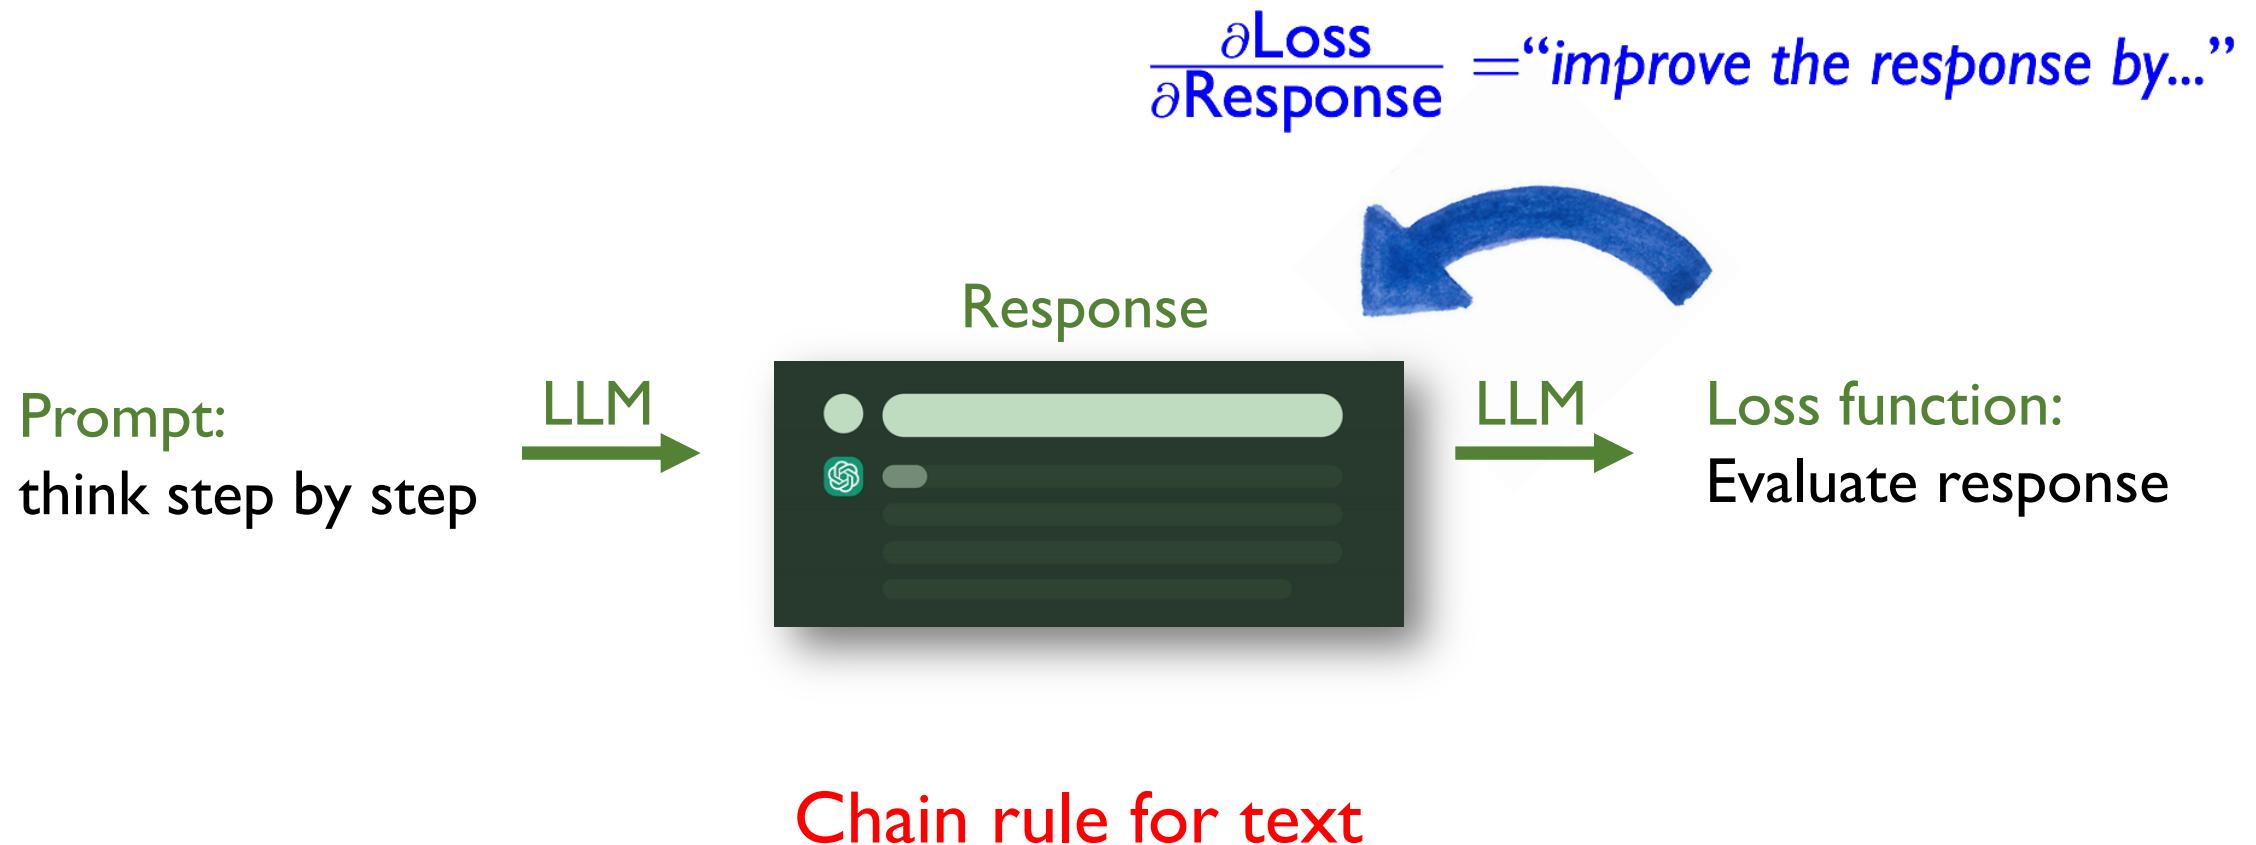
TextGrad improves LLM performance on:

1. LeetCodeHard
2. Google-proof QA
3. Treatment planning
4. Molecule opt.

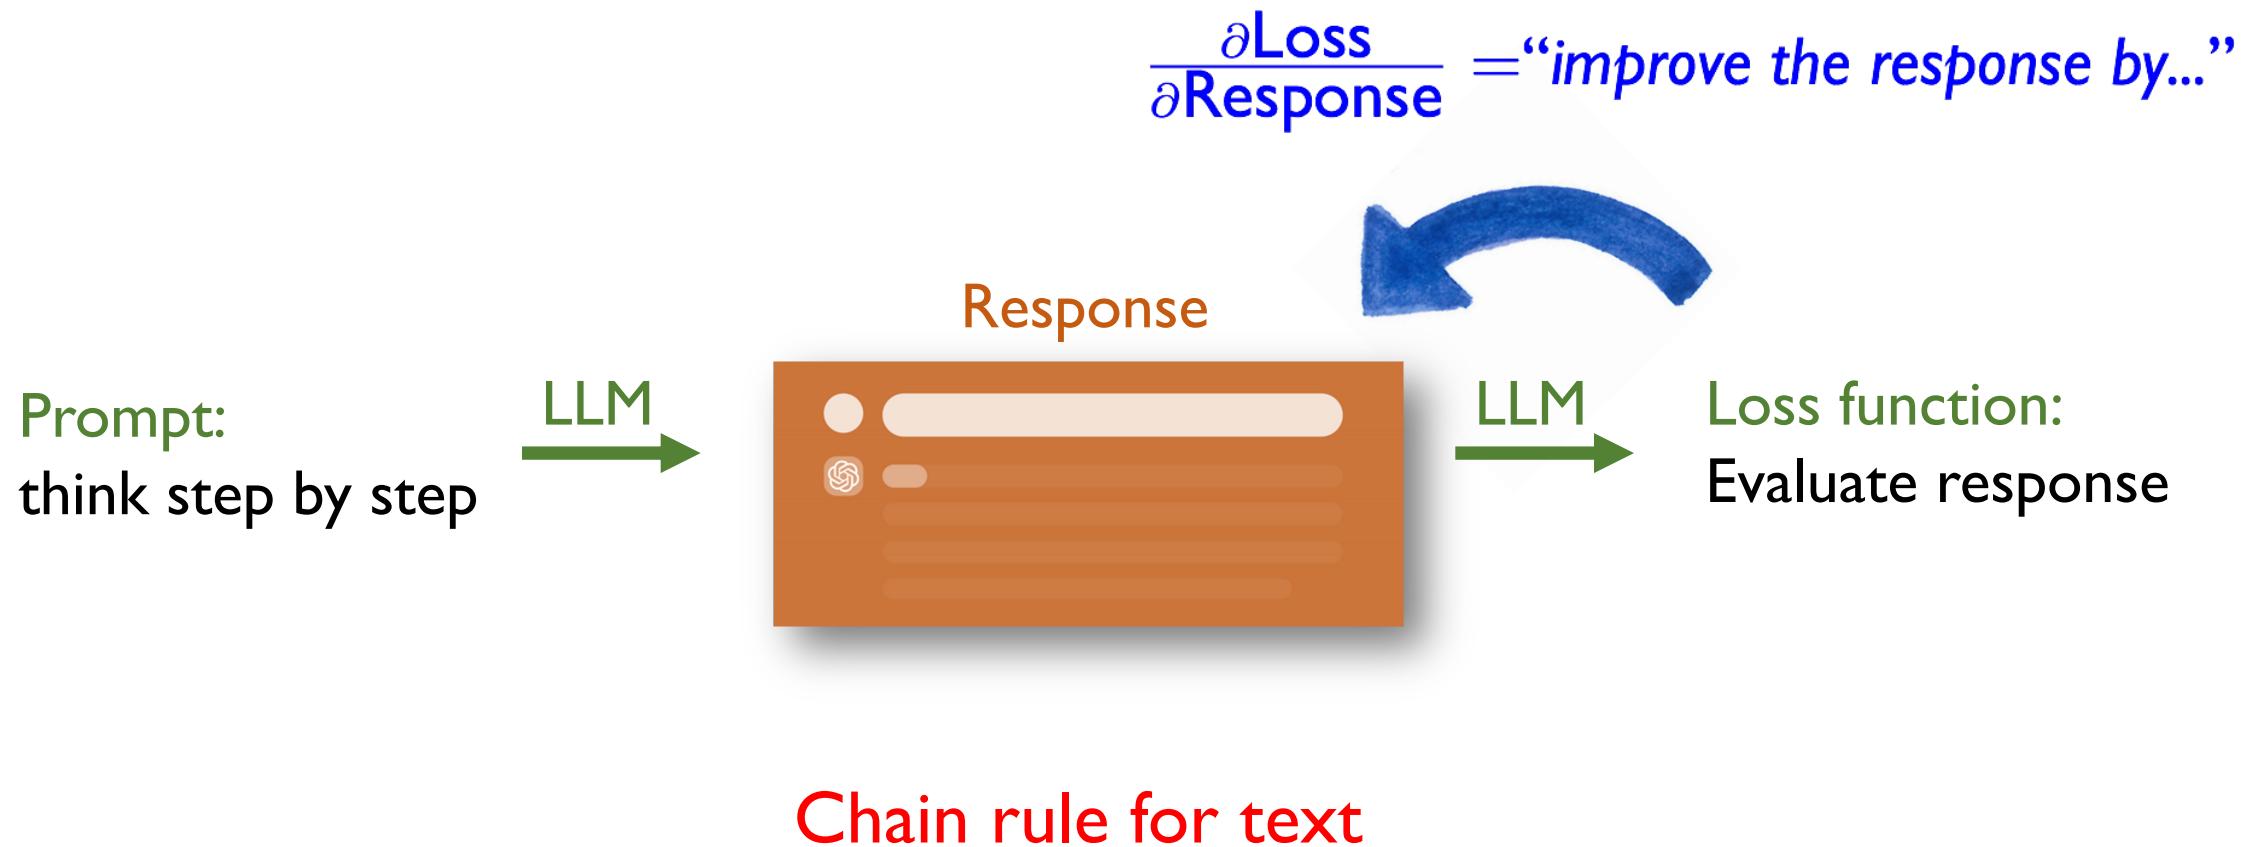
TextGrad: Automatic “Differentiation” via Text



TextGrad: Automatic “Differentiation” via Text



TextGrad: Automatic “Differentiation” via Text



TextGrad: Automatic “Differentiation” via Text

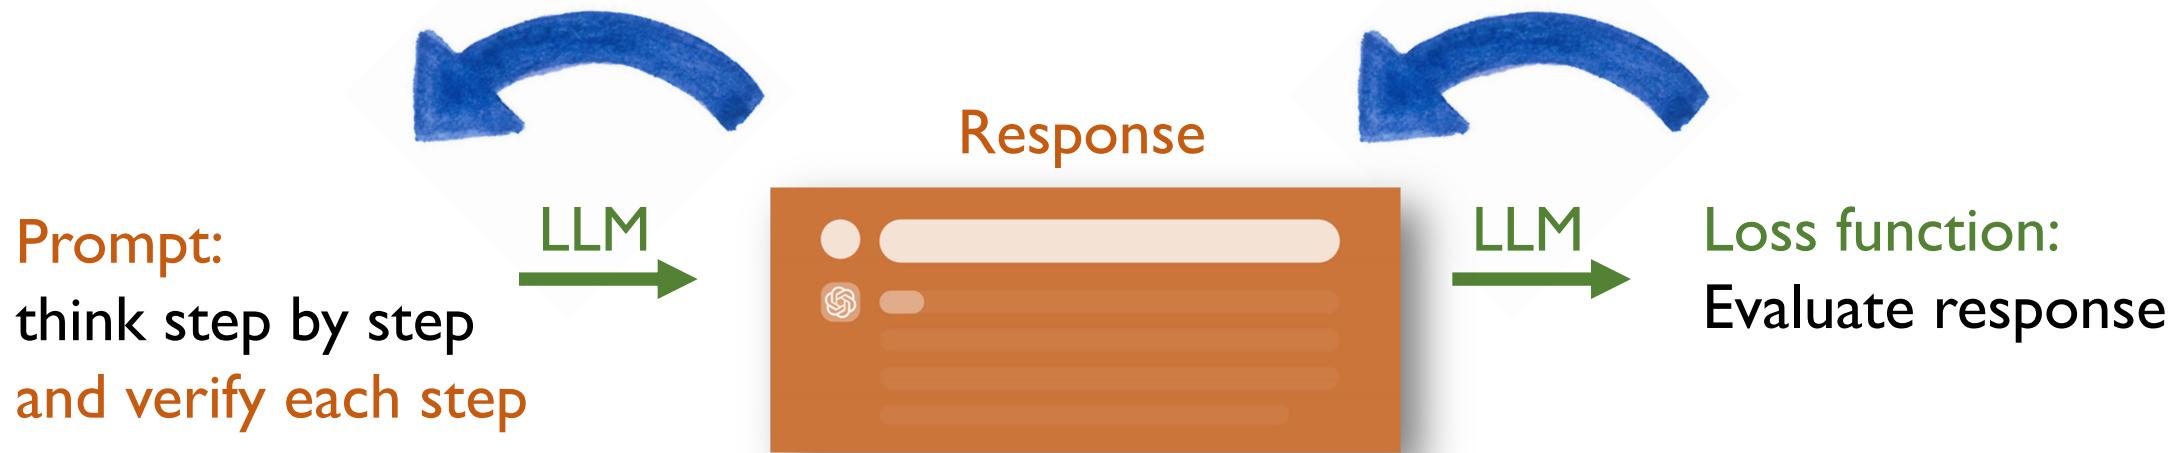
$\frac{\partial \text{Response}}{\partial \text{Prompt}} = \text{"improve the prompt by..."} \quad \frac{\partial \text{Loss}}{\partial \text{Response}} = \text{"improve the response by...”}$



Chain rule for text

TextGrad: Automatic “Differentiation” via Text

$\frac{\partial \text{Response}}{\partial \text{Prompt}} = \text{“improve the prompt by...”}$ $\frac{\partial \text{Loss}}{\partial \text{Response}} = \text{“improve the response by...”}$



Chain rule for text

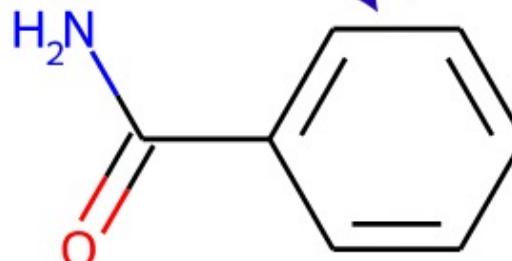
TextGrad designs new druglike molecules



Gradients: Introduce functional groups that can form hydrogen bonds or hydrophobic interactions.



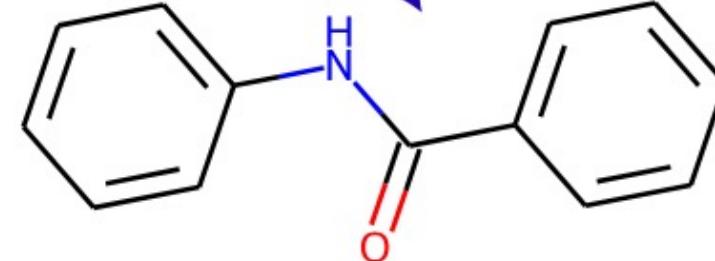
Vina: -4.2 kcal/mol
QED: 0.44



Vina: -5.5 kcal/mol
QED: 0.59



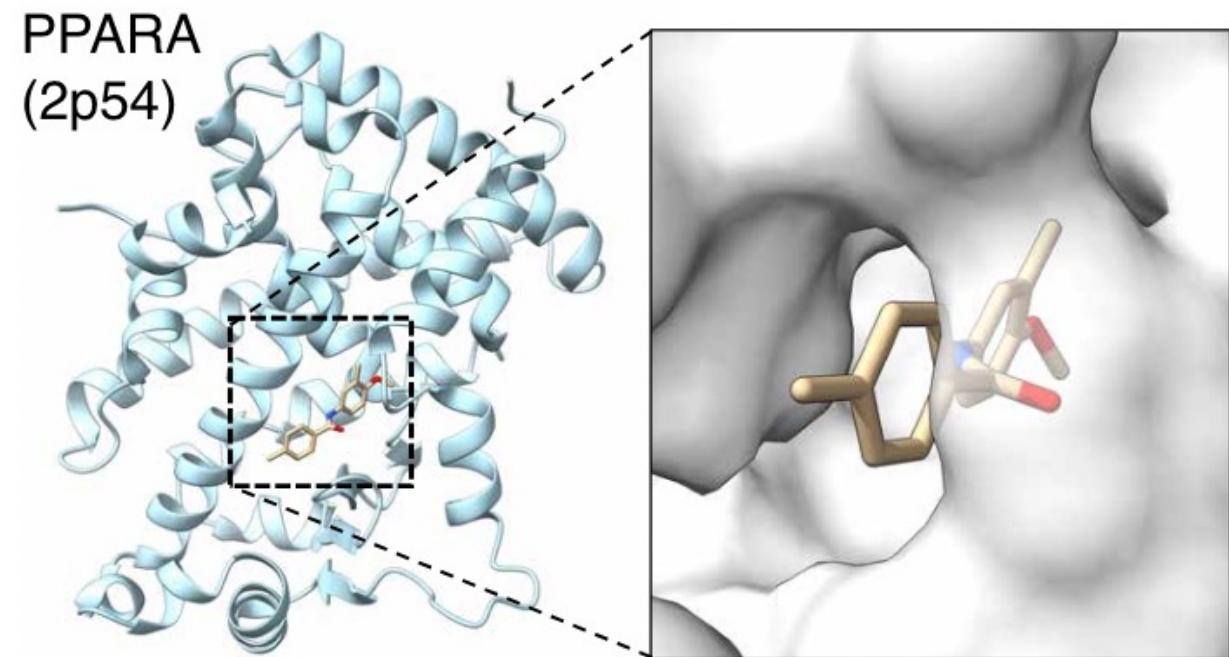
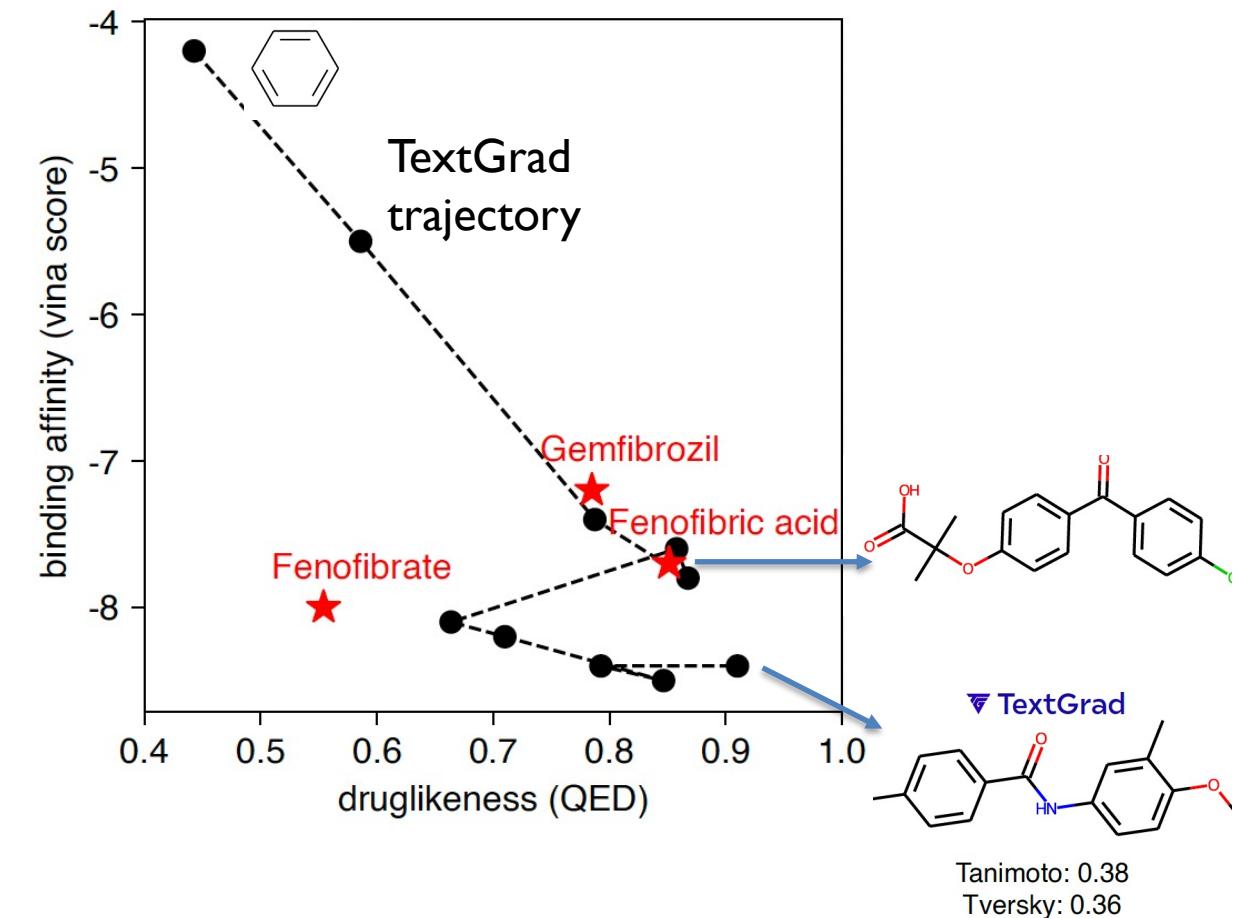
Gradients: Add hydrophobic groups or aromatic rings to enhance interactions ... while maintaining a balance of hydrophobic and hydrophilic properties.



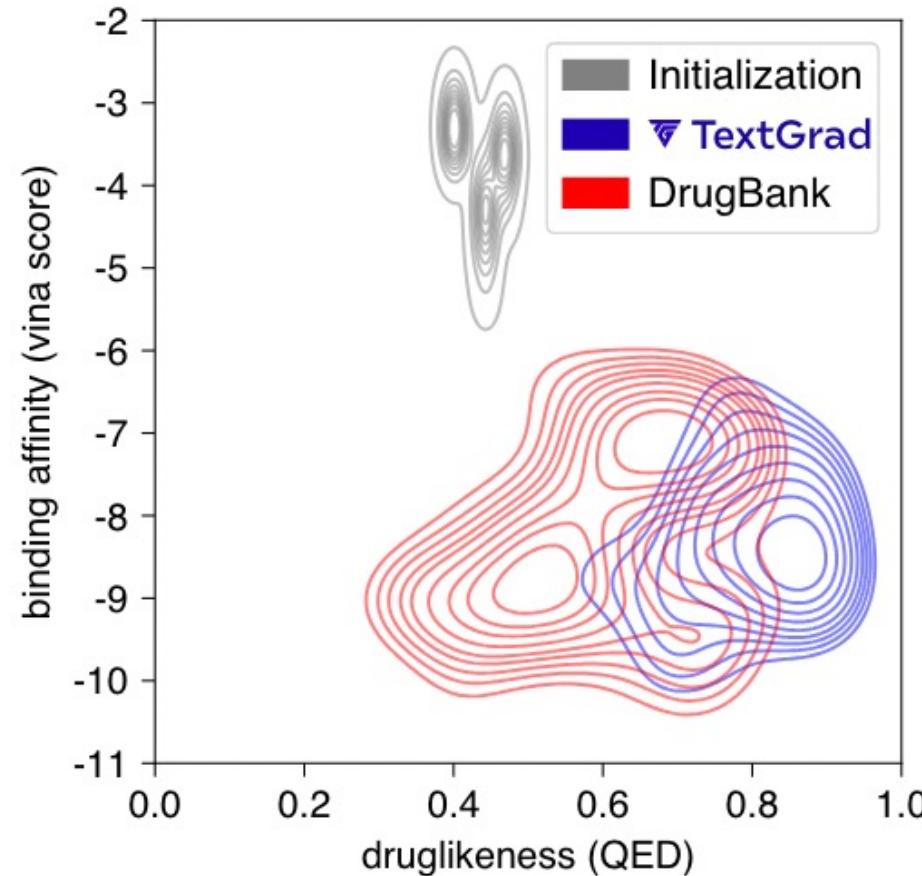
Vina: -7.5 kcal/mol
QED: 0.79

Molecule optimization

TextGrad optimizes desirable molecular properties

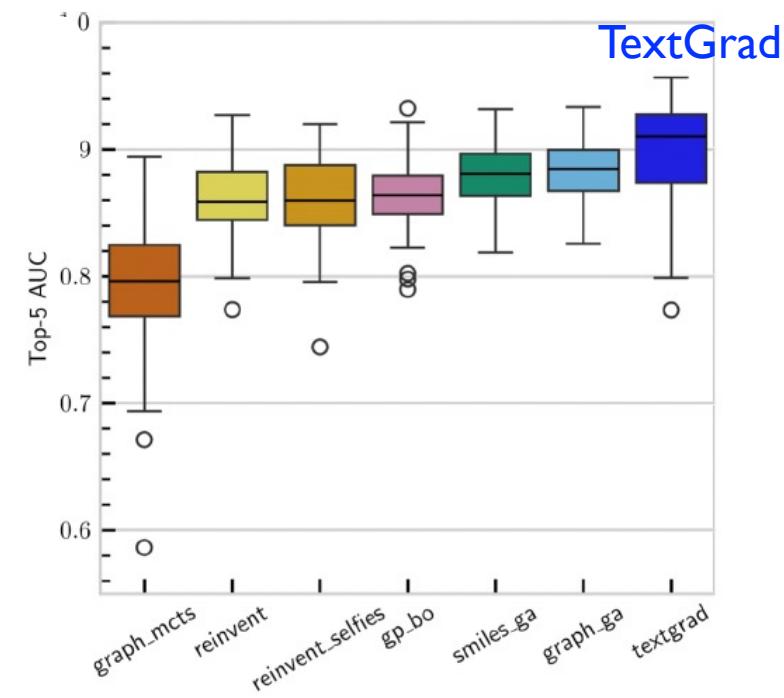
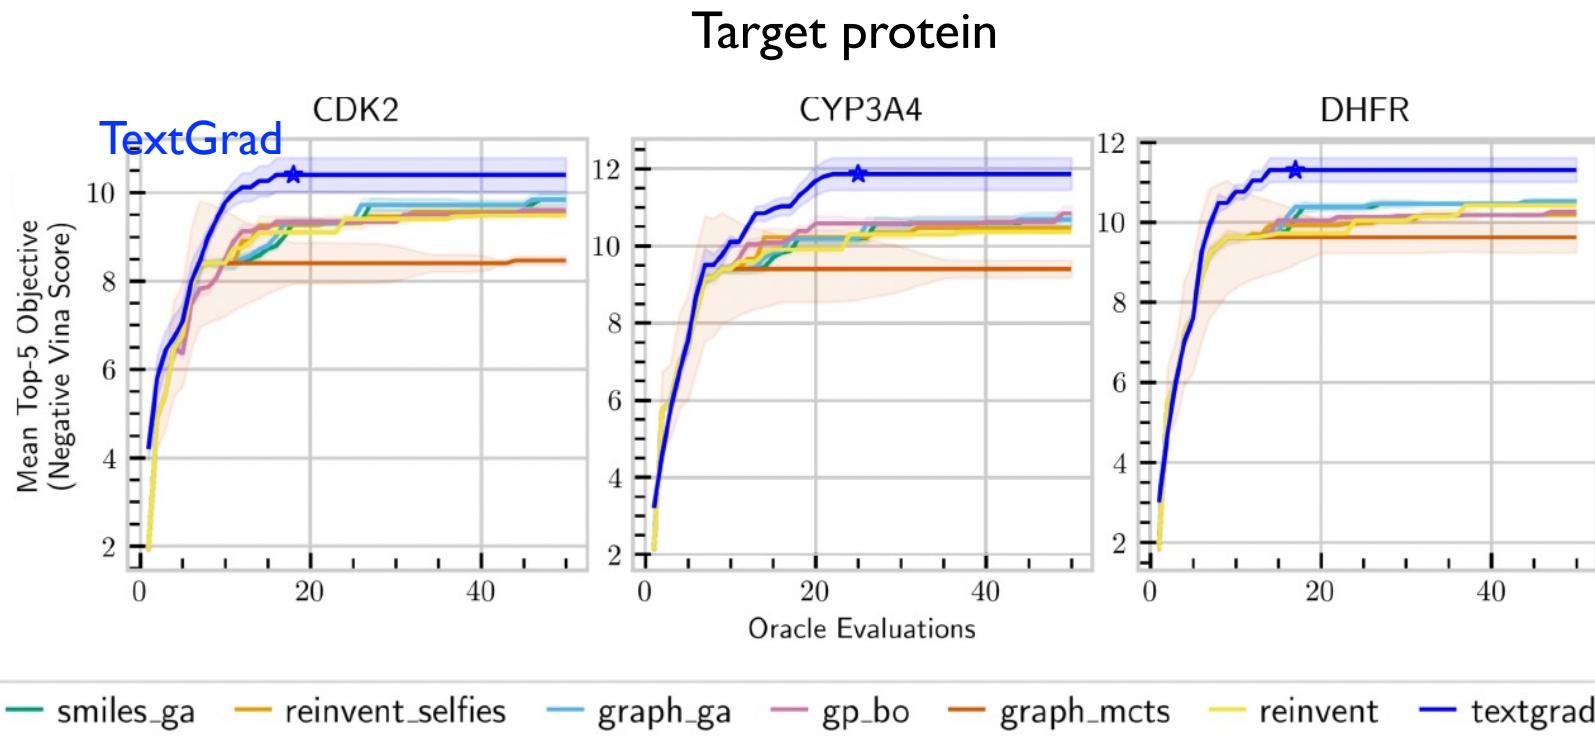


TextGrad designs desirable and novel small molecules



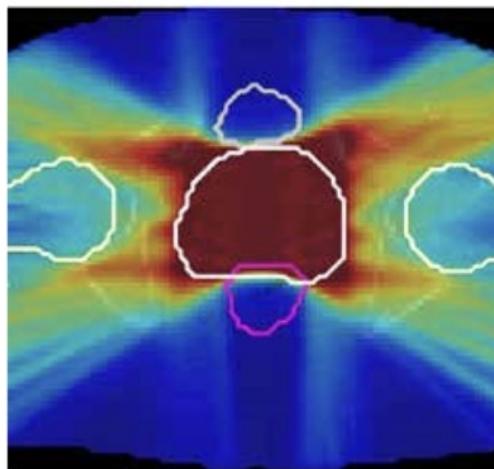
Optimizing binders to 58 target proteins

TextGrad improves over previous molecule optimizers



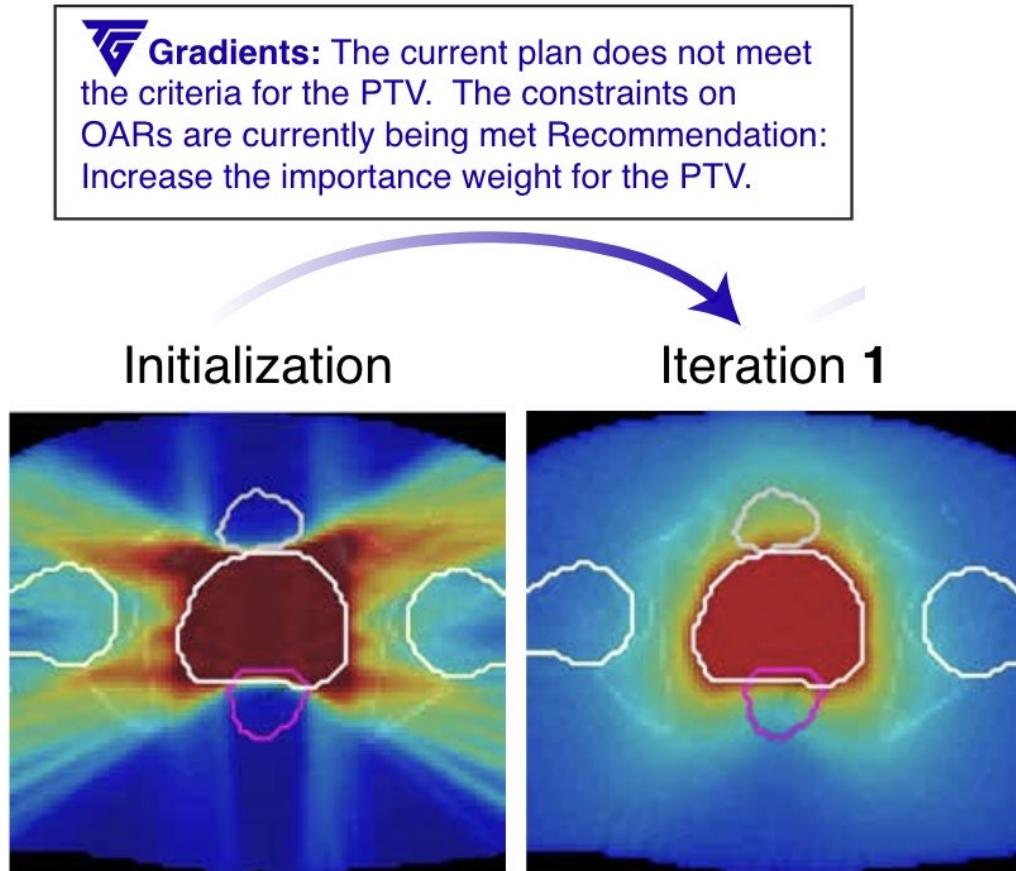
TextGrad improve radiation oncology treatment planning

Initialization



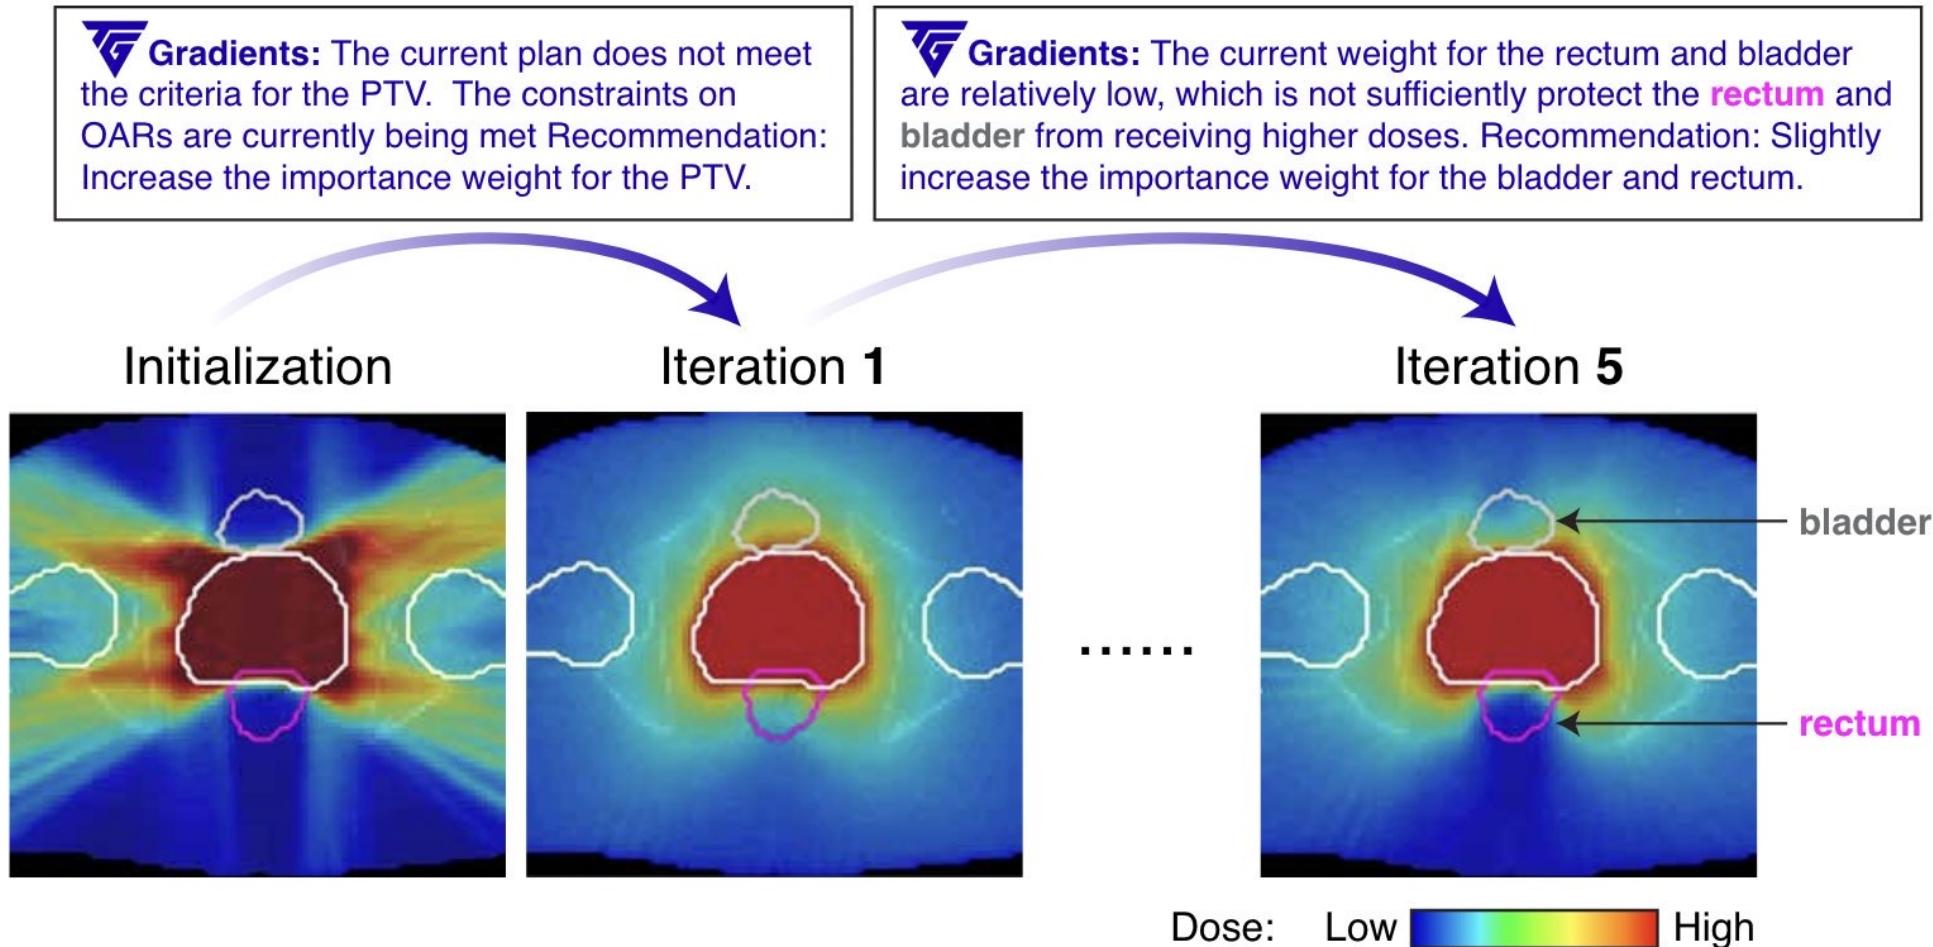
Goal: increase radiation in target region while minimizing radiation outside

TextGrad improve radiation oncology treatment planning



Goal: increase radiation in target region while minimizing radiation outside

TextGrad improve radiation oncology treatment planning



Goal: increase radiation in target region while minimizing radiation outside

TextGrad improve radiation oncology treatment planning

