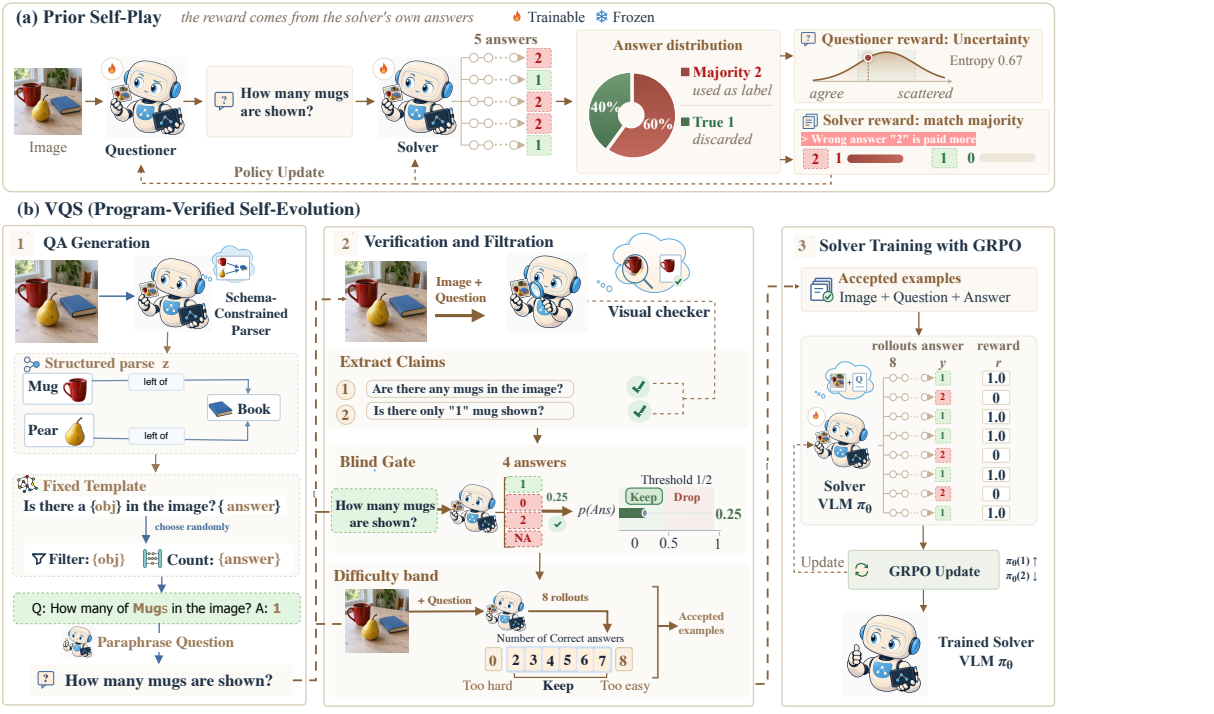

(a) Prior Self-Play
the reward comes from the solver's own answers
Trainable
Frozen
Image
Questioner
How many mugs are shown?
Solver
5 answers
2
1
2
2
1
Answer distribution
Majority 2
used as label
True 1
discarded
40%
60%
Questioner reward: Uncertainty
Entropy 0.67
agree
scattered
Solver reward: match majority
> Wrong answer "2" is paid more
2 1
1 0
Policy Update
(b) VQS (Program-Verified Self-Evolution)
1 QA Generation
Schema-Constrained Parser
Structured parse z
Mug
left of
Book
Pear
left of
Fixed Template
Is there a {obj} in the image? {answer}
choose randomly
Filter: {obj}
Count: {answer}
Q: How many of Mugs in the image? A: 1
Paraphrase Question
How many mugs are shown?
2 Verification and Filtration
Image + Question
Visual checker
Extract Claims
1 Are there any mugs in the image?
2 Is there only "1" mug shown?
Blind Gate
4 answers
How many mugs are shown?
1
0
2
NA
0.25
p(Ans)
Threshold 1/2
Keep
Drop
0.25
0
0.5
1
Difficulty band
+ Question
8 rollouts
Number of Correct answers
0 2 3 4 5 6 7 8
Too hard
Keep
Too easy
Accepted examples
3 Solver Training with GRPO
Accepted examples
Image + Question + Answer
rollouts answer reward
8 y r
1 1.0
2 0
1 1.0
1 1.0
2 0
1 1.0
2 0
1 1.0
Solver VLM π_θ
Update
GRPO Update
$\pi_\theta(1) \uparrow$
$\pi_\theta(2) \downarrow$
Trained Solver VLM π_θ