

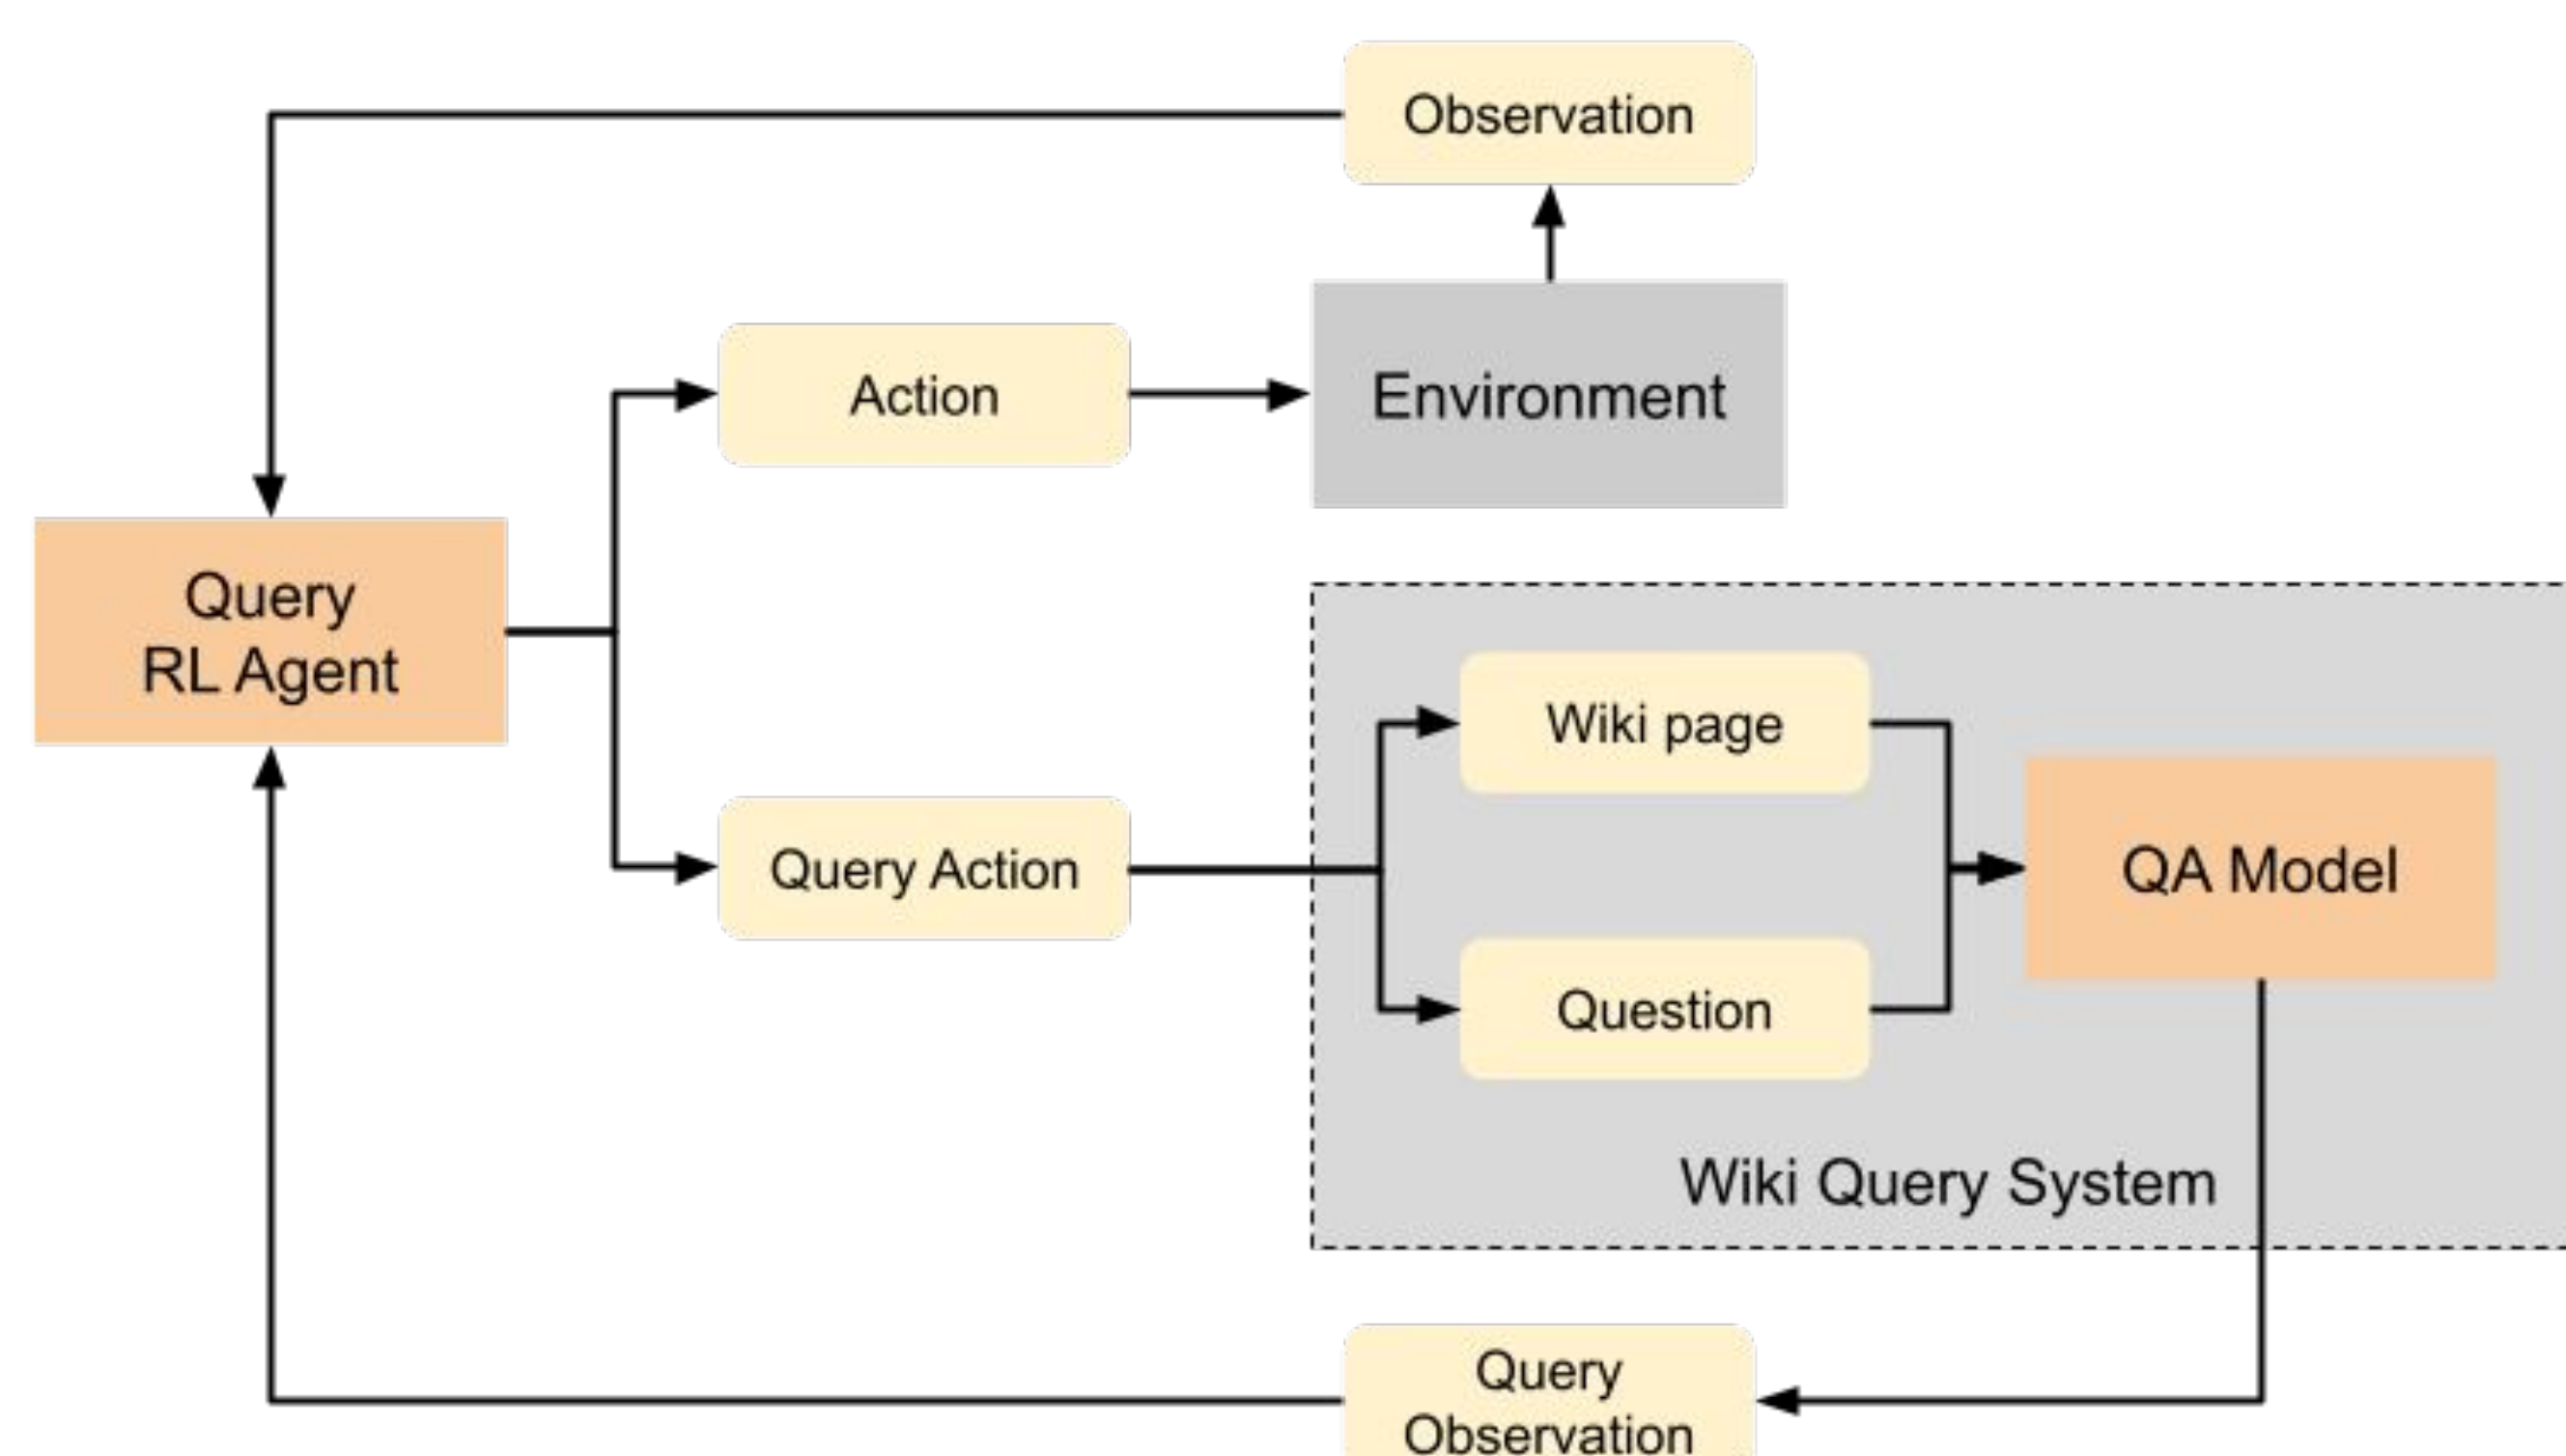
## Objectives

- Extract environment knowledge from external wiki-style knowledge base
- Minimize calls to expensive resources such as language models

## Challenges

- Unstructured, non-task-specific text
  - Noisy, difficult to extract knowledge
- Limited amount of natural text
  - Requires zero/few-shot learning
- Pretrained language models (LMs) are slow at inference time

## Query Agent




## Environment




- Agent attacks with fire or cold weapon
- Monsters have fire or cold resistance
- Rewarded for slaying monsters

## Wiki

- 3,000+ total pages
- 393 monsters
- 9 resistance and 9 attack types



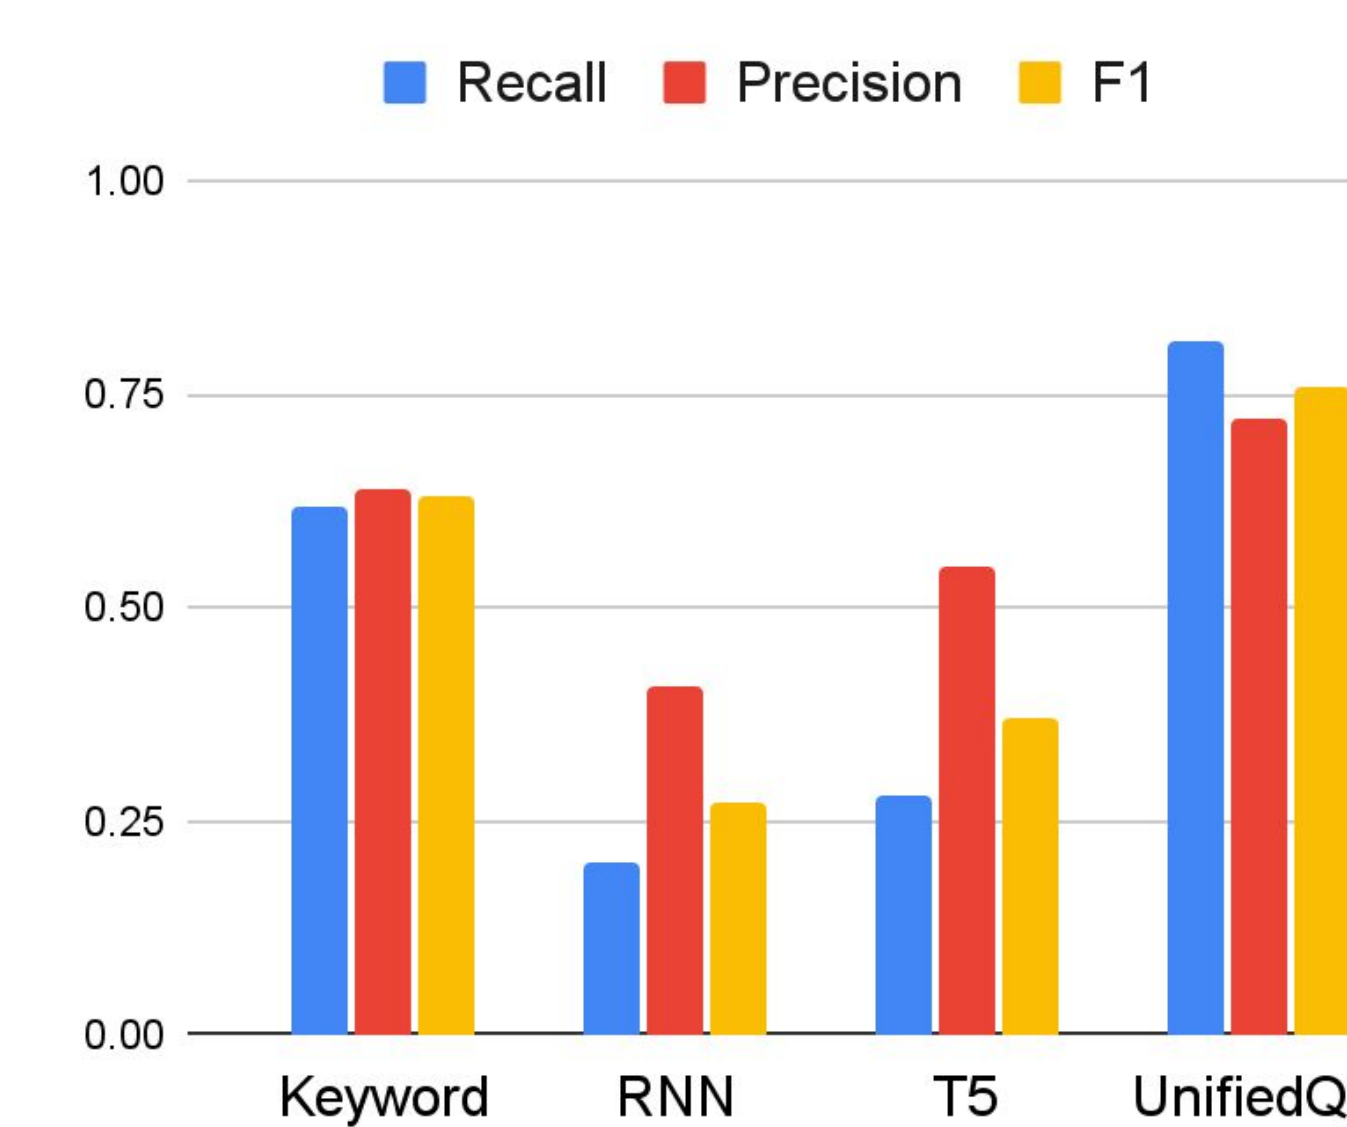
“Balrogs have [...] resistances to fire and poison.”



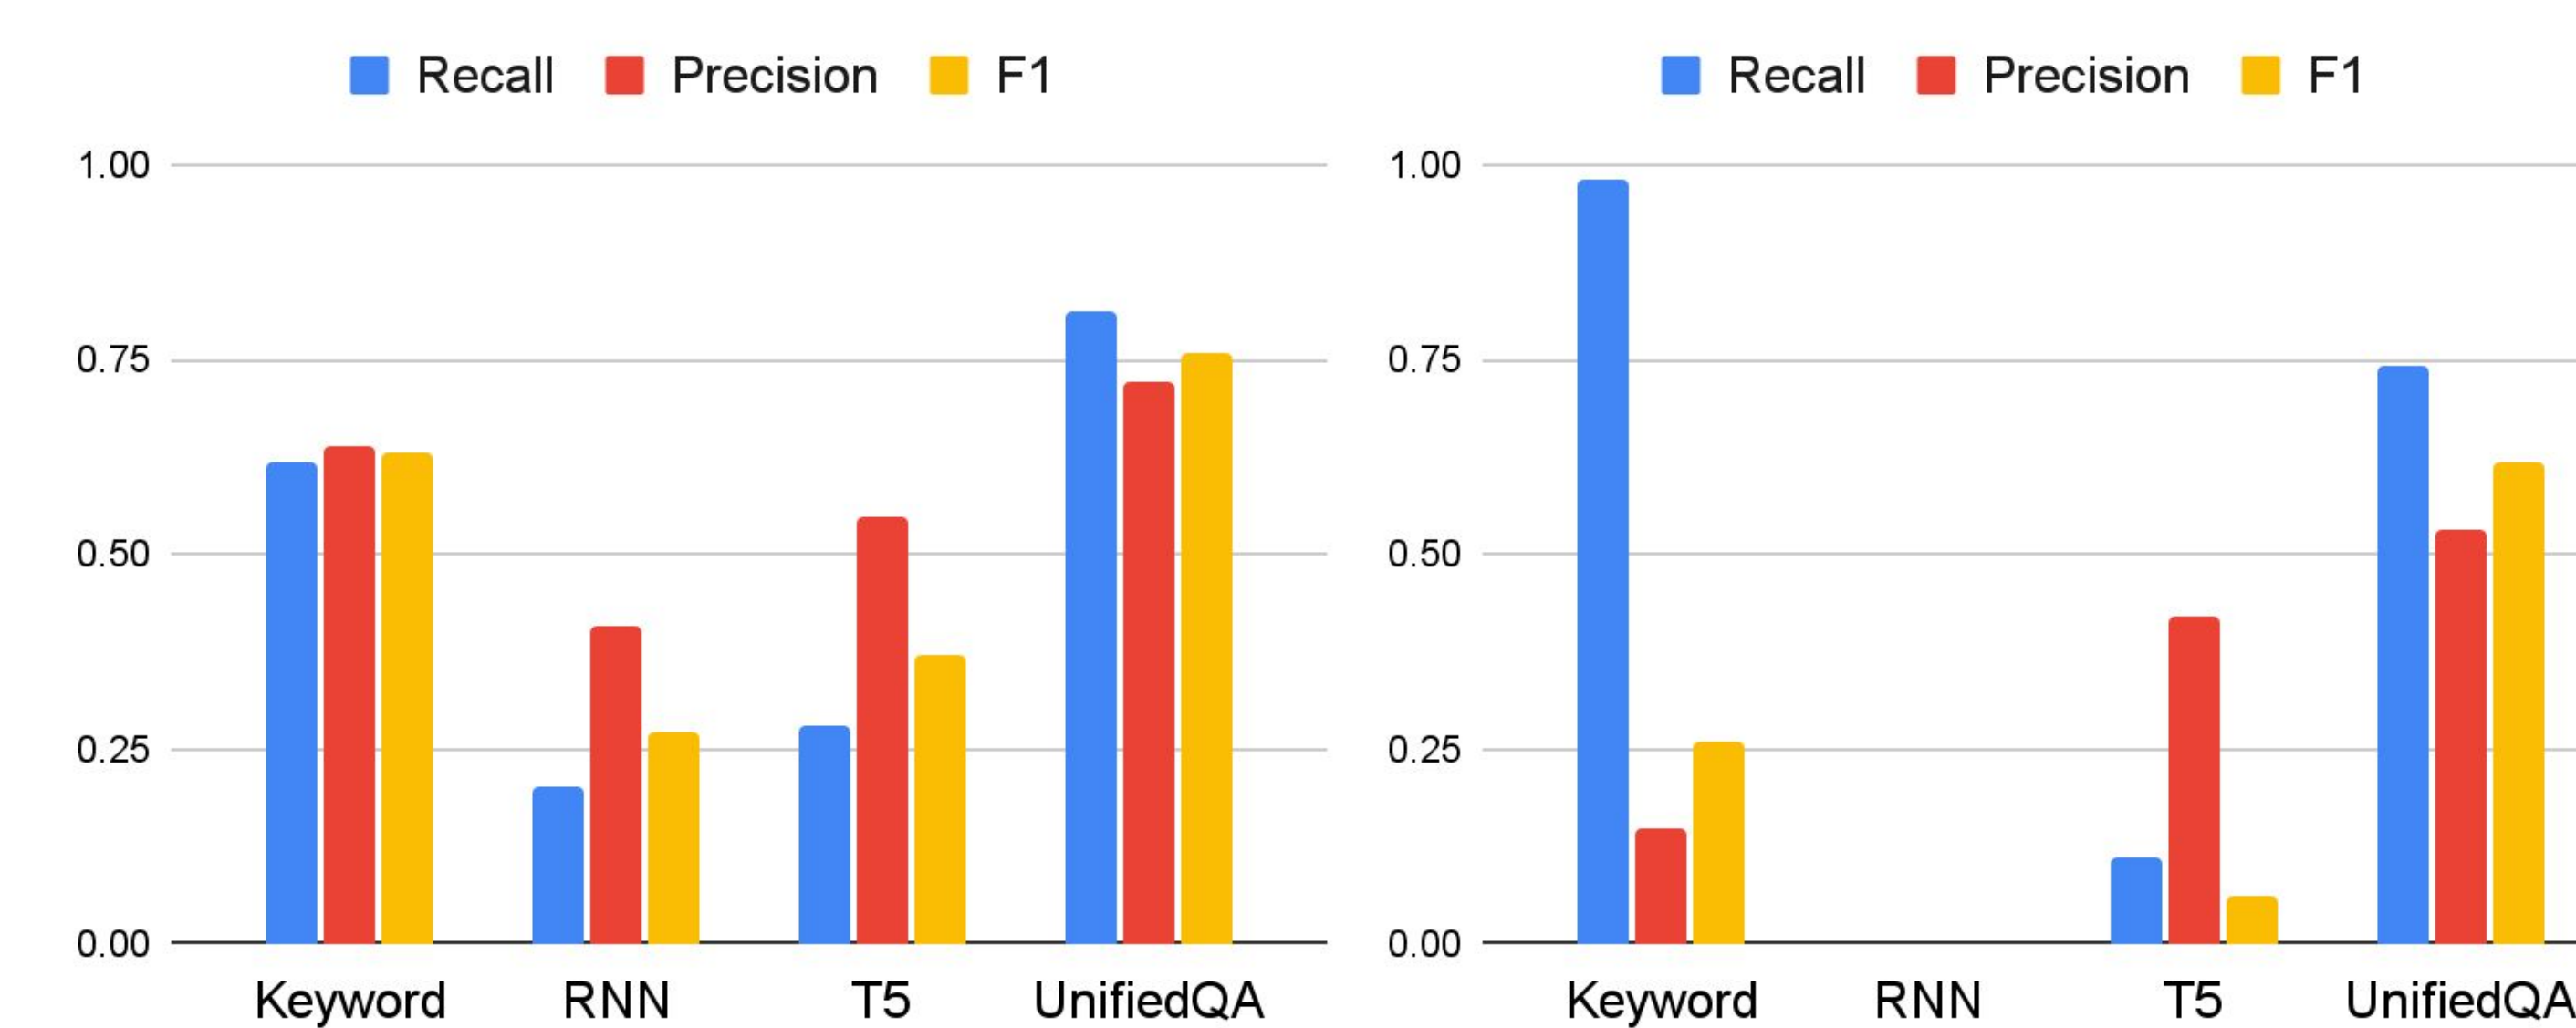
“Ice trolls are immune to cold, and one of their attacks deals cold damage.”

## Analysis

Resistance Prediction

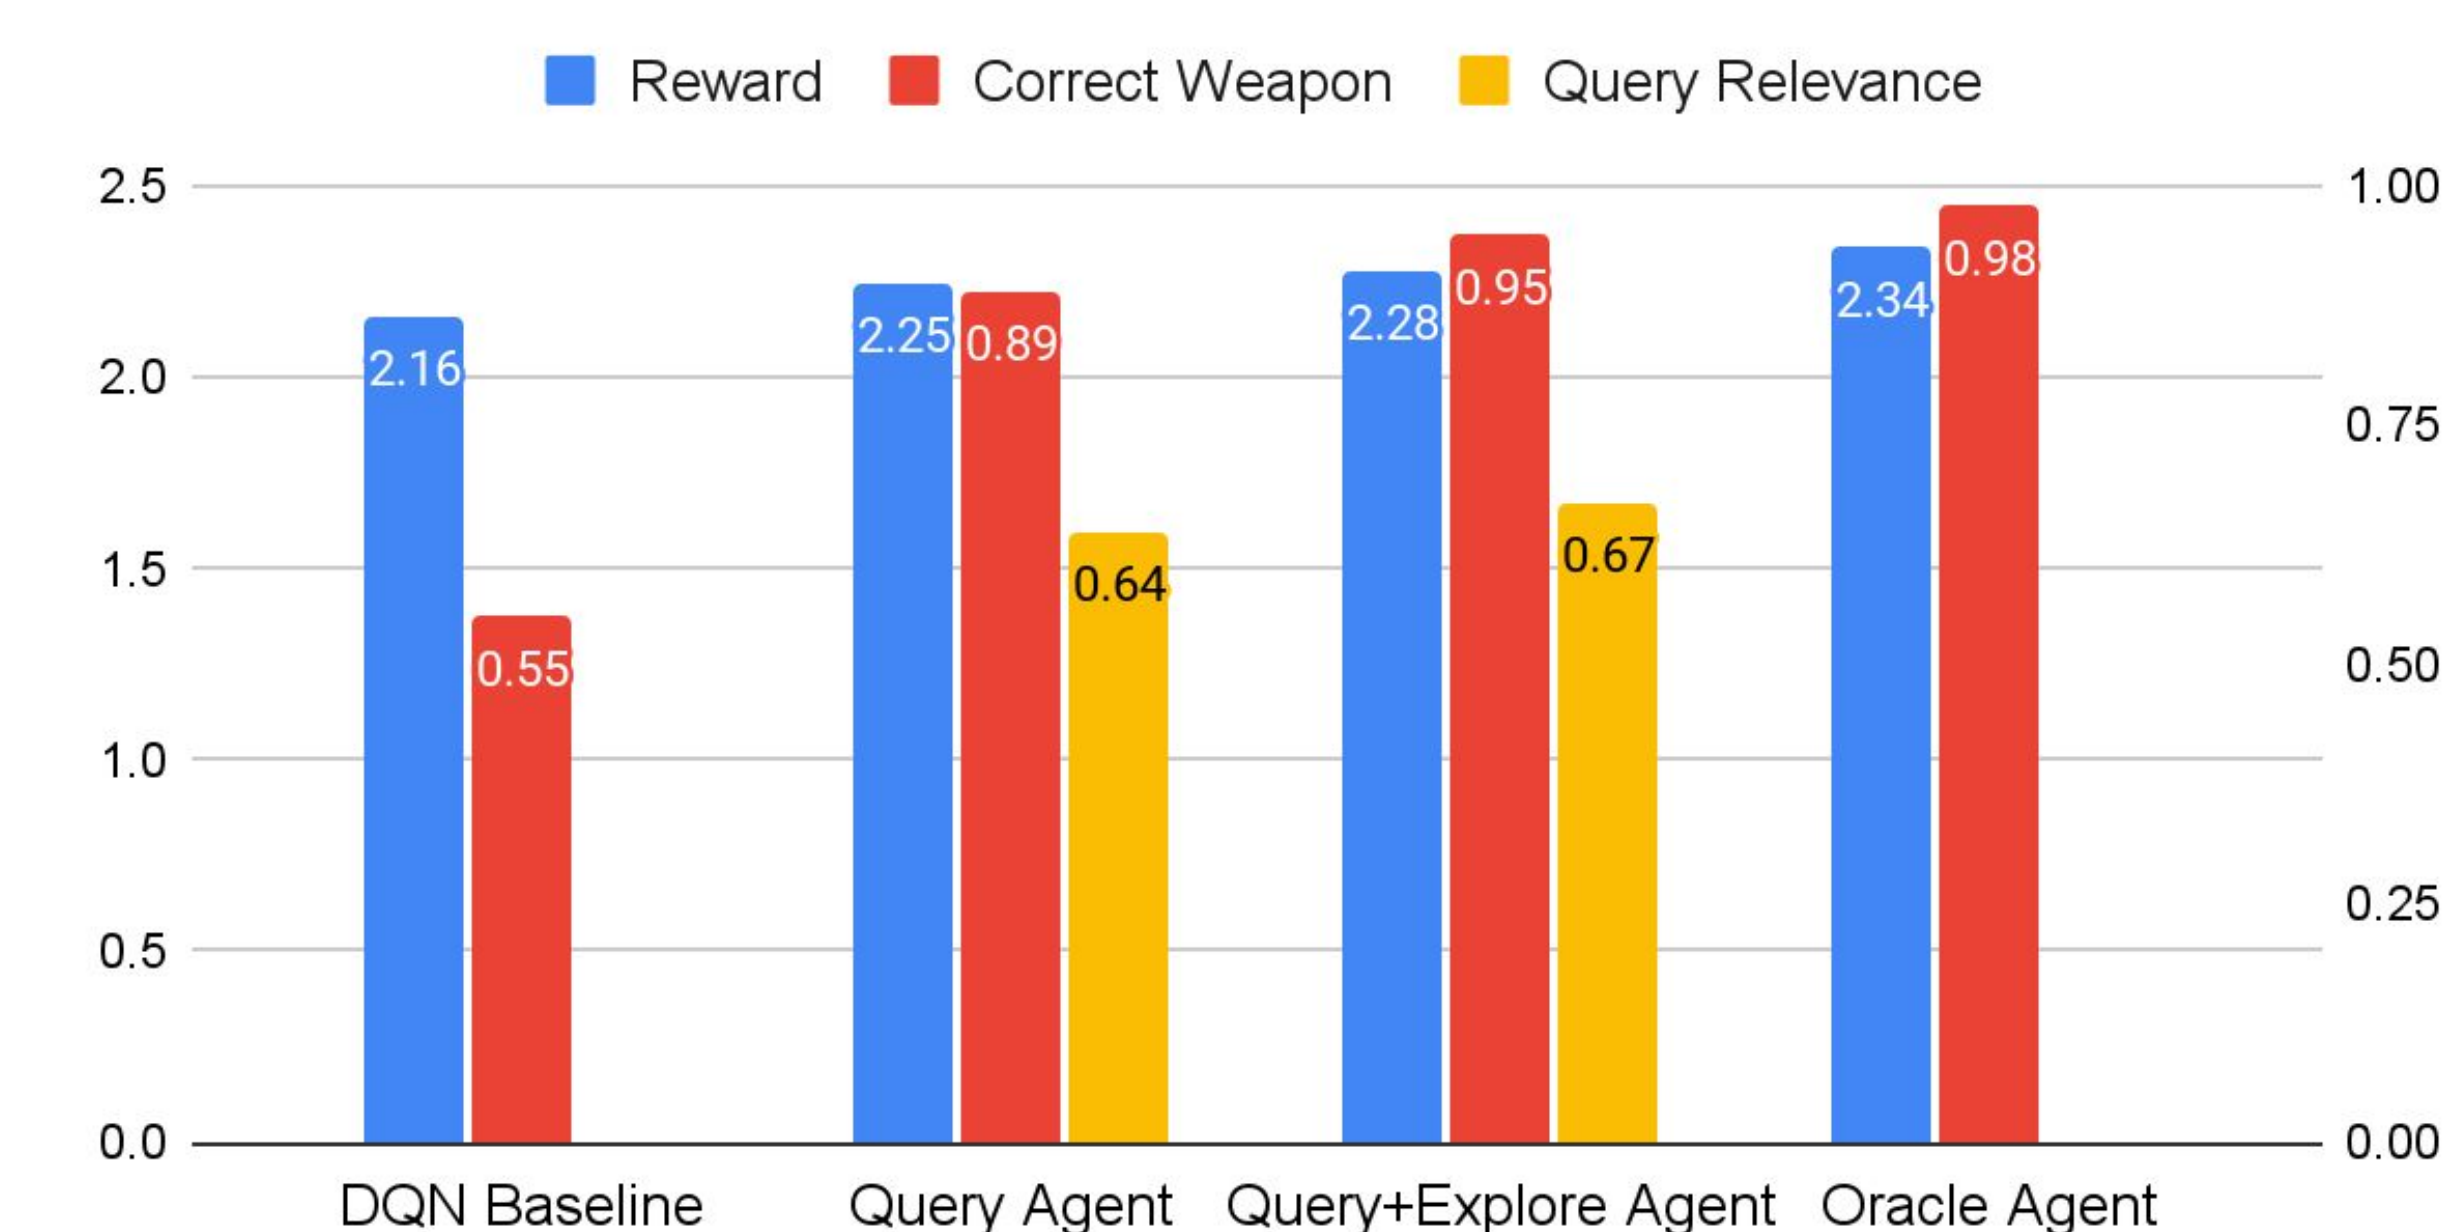


Attack Prediction



- LMs are necessary for generalization

RL Results



- Learns to query the ambiguous states
- Query-specific exploration increases convergence speed by x4

## Summary

- Nethack for testing NL informed RL
- QA models for zero-shot knowledge
- Learning *when* to query improves inference latency