Galaxy Component Decomposition Agent
星系成分分解智能体

Teaching a AI to follow an astronomer's scientific inference process
让机器学习天文学家的科学反演过程

Duan Manni
段曼妮

@Zhejiang Lab-地球与空间科学计算研究中心

Jinyi Shangguan @ Peking University



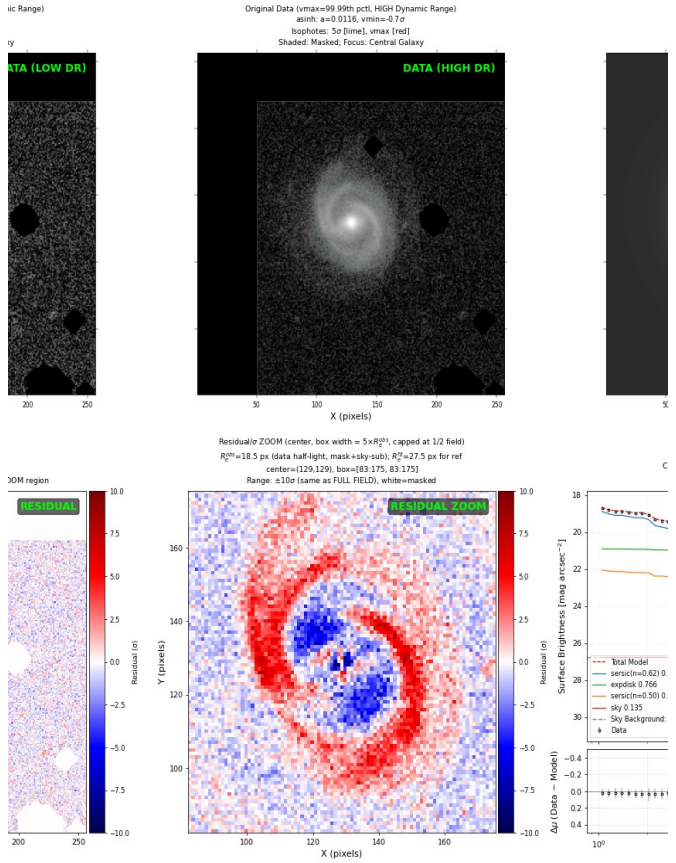

Image · Model · Residual · Radial profile / 原图 · 模型 · 残差 · 径向轮廓

What is galaxy decomposition? A galaxy image combines several physical structures. 什么是星系成分分解？一张星系图像，是多个物理结构叠加后的结果

图1 星系成分分解的核心概念：观测光分布 = 多个结构成分 + PSF/噪声

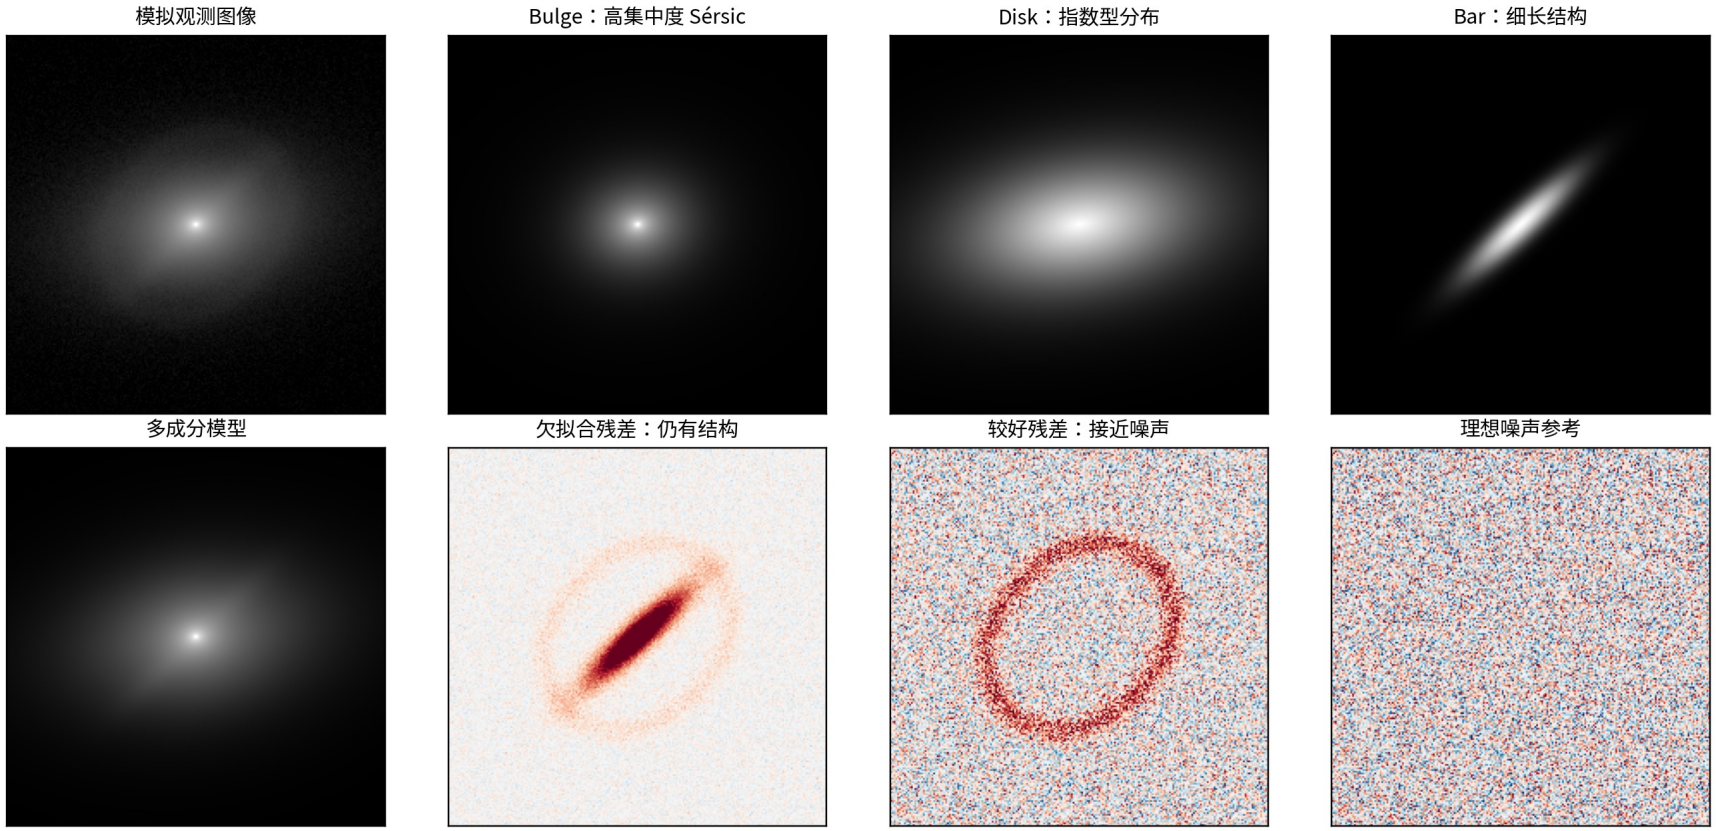

示意图由解析光度模型合成，用于说明概念，不代表真实星系。

Observed light = structural components ⊗ PSF + background + noise
观测光分布 = 结构成分 ⊗ PSF + 背景 + 噪声

Why does it matter? A small bulge bias can make the Universe look younger

为什么它重要？少算一点核球，会让整个宇宙看起来比实际更“年轻”

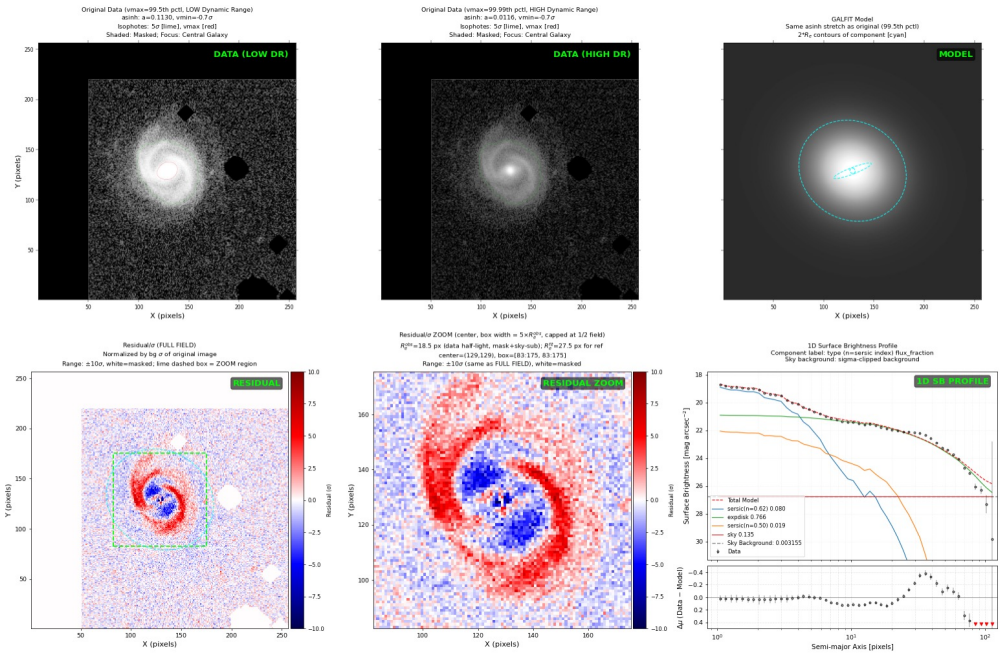

Real case: image · model · residual · radial profile / 真实案例：原图 · 模型 · 残差 · 径向轮廓

B/T **What fraction of all light comes from the central bulge?**
全部星光中，有多少来自中央核球？

If the bulge is systematically underestimated
如果核球被系统性少算

① B/T is too low for each galaxy

② More galaxies appear disk-dominated

③ mature-galaxy population underestimated

④ The pace of galaxy aging are misread

Scientific impact: the structural-evolution timeline shifts later
科学结论变化：结构演化时间线被推迟

Is this only numerical optimization? Reliable fitting checks images, residuals, and radial profiles

它是个数值优化问题吗？真实拟合需要同时检查图像、残差和径向轮廓

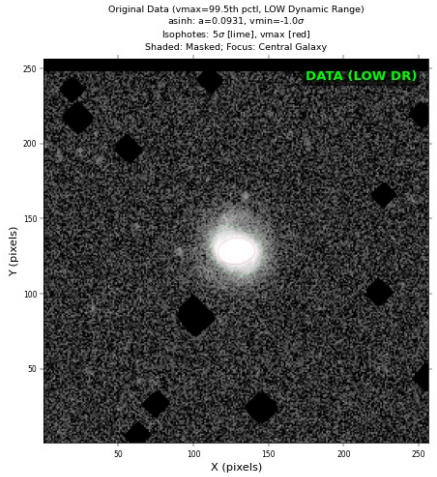

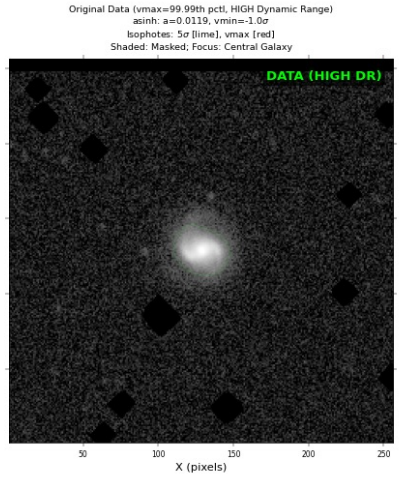



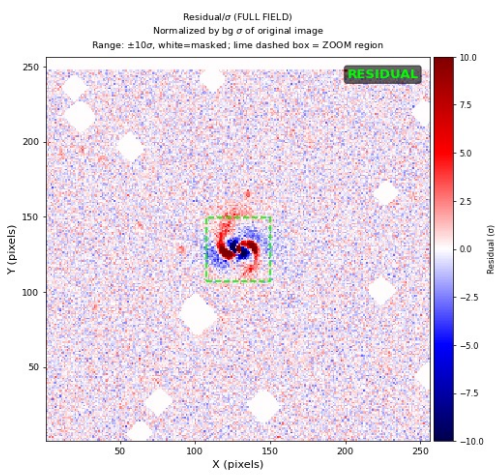

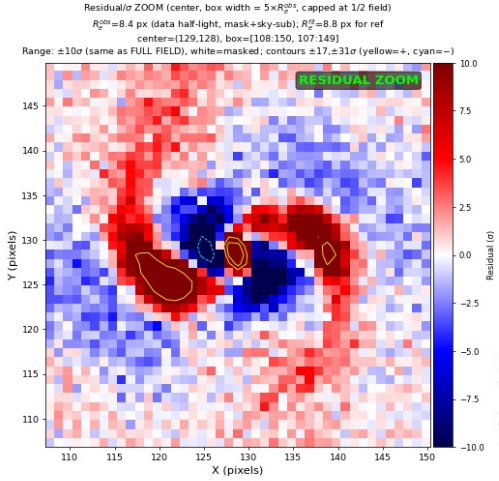

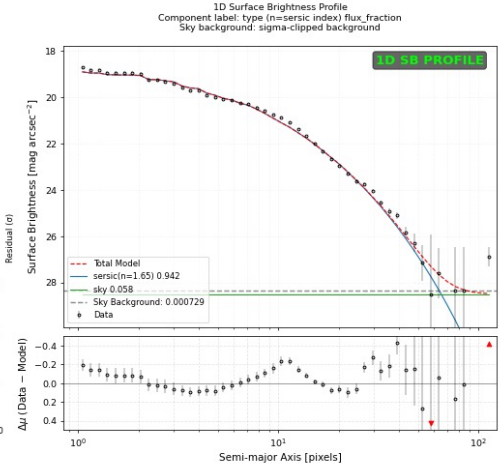

One result needs at least three kinds of evidence
同一个结果，至少有三种证据

Image / 图像

What structure is missing?
遗漏了什么结构？

Residual / 残差

Does a coherent pattern remain?
是否仍有相关模式？

Profile / 轮廓

Are parameters at a boundary or mismatched?
参数是否贴边或失配？

How does it work? Experts follow a Run-Diagnose-Decide-Refit-Stop loop

实际如何工作呢？专家实际执行的是 Run-Diagnose-Decide-Refit-Stop 闭环

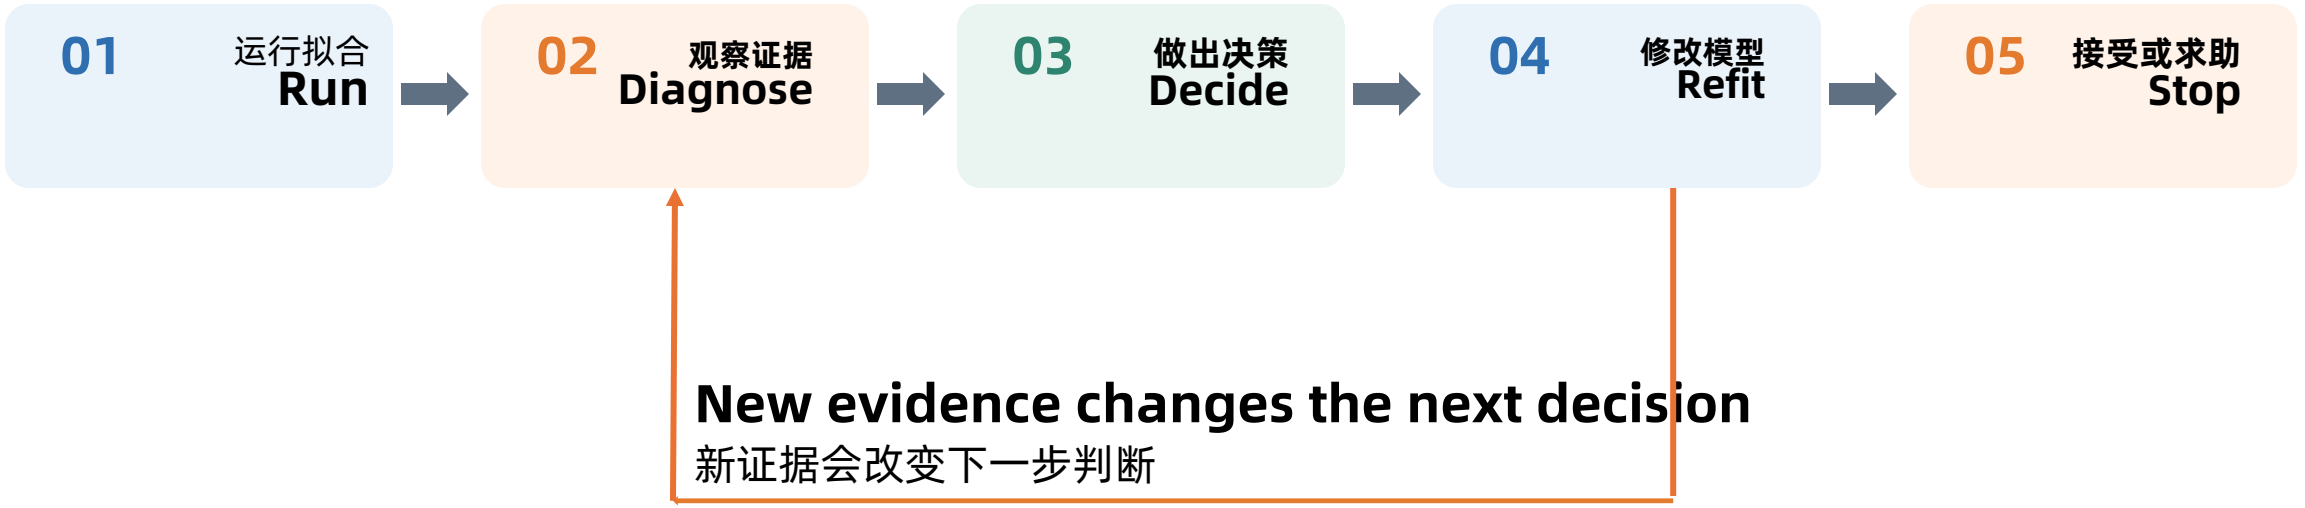

A single run is only one observation in the loop

一次运行只是闭环中的一个观察步骤

NGC 4608: a real example
NGC 4608：一个真实的例子

Same real galaxy | Every refit requires a new diagnosis
同一颗真实星系 | 每次 Refit 之后，都必须重新 Diagnose

Second diagnosis: compare the new evidence after refitting

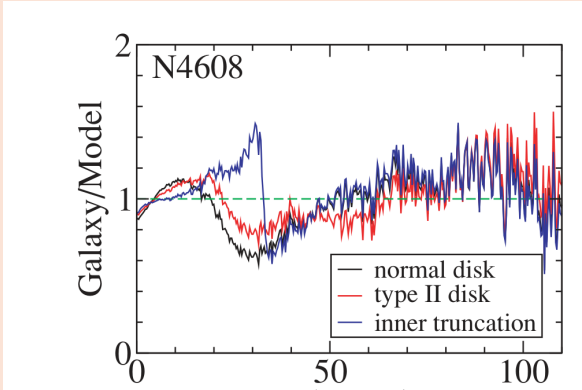

Black: standard disk
Blue: central truncation
Red: Type II
Green Dashed: observed = model

黑：普通盘 | 蓝：中心截断 | 红：Type II | 虚线：观测=模型

Final stop: accept the complete result with a Type II disk

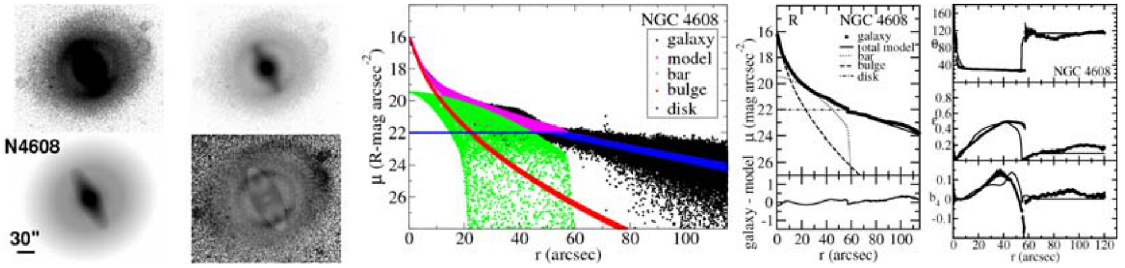

Observation · Total model · 2D residual · Component and ellipticity profiles
真实观测 · 总模型 · 二维残差 · 成分与椭率剖面

RUN

Standard exponential disk / fit completed
普通指数盘
完成拟合

DIAGNOSE$_t$

Continuous negative residual from 20″ to 50″
20″–50″
连续负残差

Observe / 观察

DECIDE

Propose two disk-structure hypotheses
提出两种
盘结构假设

Hypothesis A / abrupt central truncation

Hypothesis B / flatter inner Type II disk

REFIT

Change the disk component and rerun both hypotheses
加入 / 调整盘成分

分别重跑假设A和假设B

DIAGNOSE$_{t+1}$

A: new positive central residual
B: original negative residual weakens
假设A：中心空洞出现新正残差

假设B：原负残差减弱

Observe / 观察

STOP

Choose Type II / χ^2 improves by about 25%
选择 Type II
χ^2 改善约25%

Scientific result: the bar may capture about 13% of disk stars and grow by about 1.7×
科学结论：棒可能捕获约 13% 的盘恒星，使自身质量增长约 1.7 倍

Can decomposition scale? Prior work scaled execution, but not the decision process

能批量做分解吗？相关工作已把“运行”规模化，但尚未把决策能力规模化

	Execution automation / 执行自动化			Judgment and loop / 判断与闭环				
Representative work 代表工作	**Prepare** 准备	**Configure** 配置	**Run** 运行	**Diagnose** 诊断	**Decide** 决策	**Refit** 重拟合	**Stop** 停止	**Scale, results, and what they show** 真实规模 / 结果，以及它说明了什么
GALAPAGOS Barden+ 2012	Automatic	Preset	Automatic	Flags	—	Human	Rules	HST case: about 1,000 primary sources in 24 h on one 2.2 GHz CPU. Fixed configurations scale throughput, while results still depend on user settings. HST案例：≈1000主源 / 24 h / 单颗2.2 GHz CPU → 固定配置可扩展吞吐量，结果仍依赖用户设置
GIM2D Simard+ 2011	Automatic	3 models	Automatic	F-test	F-test	—	Rules	1,123,718 SDSS galaxies with three models and an F-test. Million-scale selection works, but the candidate models are fixed in advance. 1,123,718个SDSS星系；3套模型 + F-test → 百万级模型选择可行，但候选模型预先固定
GALFITM Häußler+ 2013	Joint	Polynomial	Automatic	Numerical	—	—	Numerical	Joint fitting across nine bands reduces scatter and failures at low signal-to-noise, but does not create new model hypotheses. 9波段联合拟合；低信噪比波段散布与失败减少 → 跨波段证据利用增强，未生成新的模型假设
S⁴G Salo+ 2015	Semi-auto	Human	Semi-auto	Human	Human	Human	Human	2,352 nearby galaxies. Final multi-component decomposition still relies heavily on human interpretation. 2352个近邻星系；最终多成分分解重度依赖人工解释 → 真实结构越复杂，专家判断越难被固定流程替代
ProFit Casura+ 2022	Automatic	4-step tree	Automatic	Rules	Rules	Swap	Flags	13,096 GAMA–KiDS galaxies. Known failure modes can be encoded, but an unsuitable model can only be flagged. 13,096个GAMA–KiDS星系；成分交换后重拟合 → 已知失败模式可编码；无合适模型时只能标记
Euclid评测 Bretonnière+ 2023	Simulation	Benchmark	Batch	Overall QA	Human	—	Benchmark	About 1.5 million simulated galaxies and five methods. Real non-analytic shapes produce roughly 3× larger errors. ≈150万模拟星系、5种方法；真实非解析形态误差≈3× → 亿级执行仍受模型失配与系统误差限制

Automatic execution 自动执行

Preset rules / limited QA 预设规则 / 有限QA

Human judgment 人工判断

Next step: study evidence-driven sequential diagnosis and decision intelligence on top of mature solvers

未来核心：在成熟求解器之上，研究证据驱动的序贯诊断与决策智能

What are we going to build?
我们打算做什么？

Build an agent to learn the reasoning process that an astronomer uses around an existing tool.

What intelligent work do experts perform? They repeatedly propose and test hypotheses against evidence
回顾下，专家做了哪些“智能”工作？他们围绕证据反复提出和验证假设

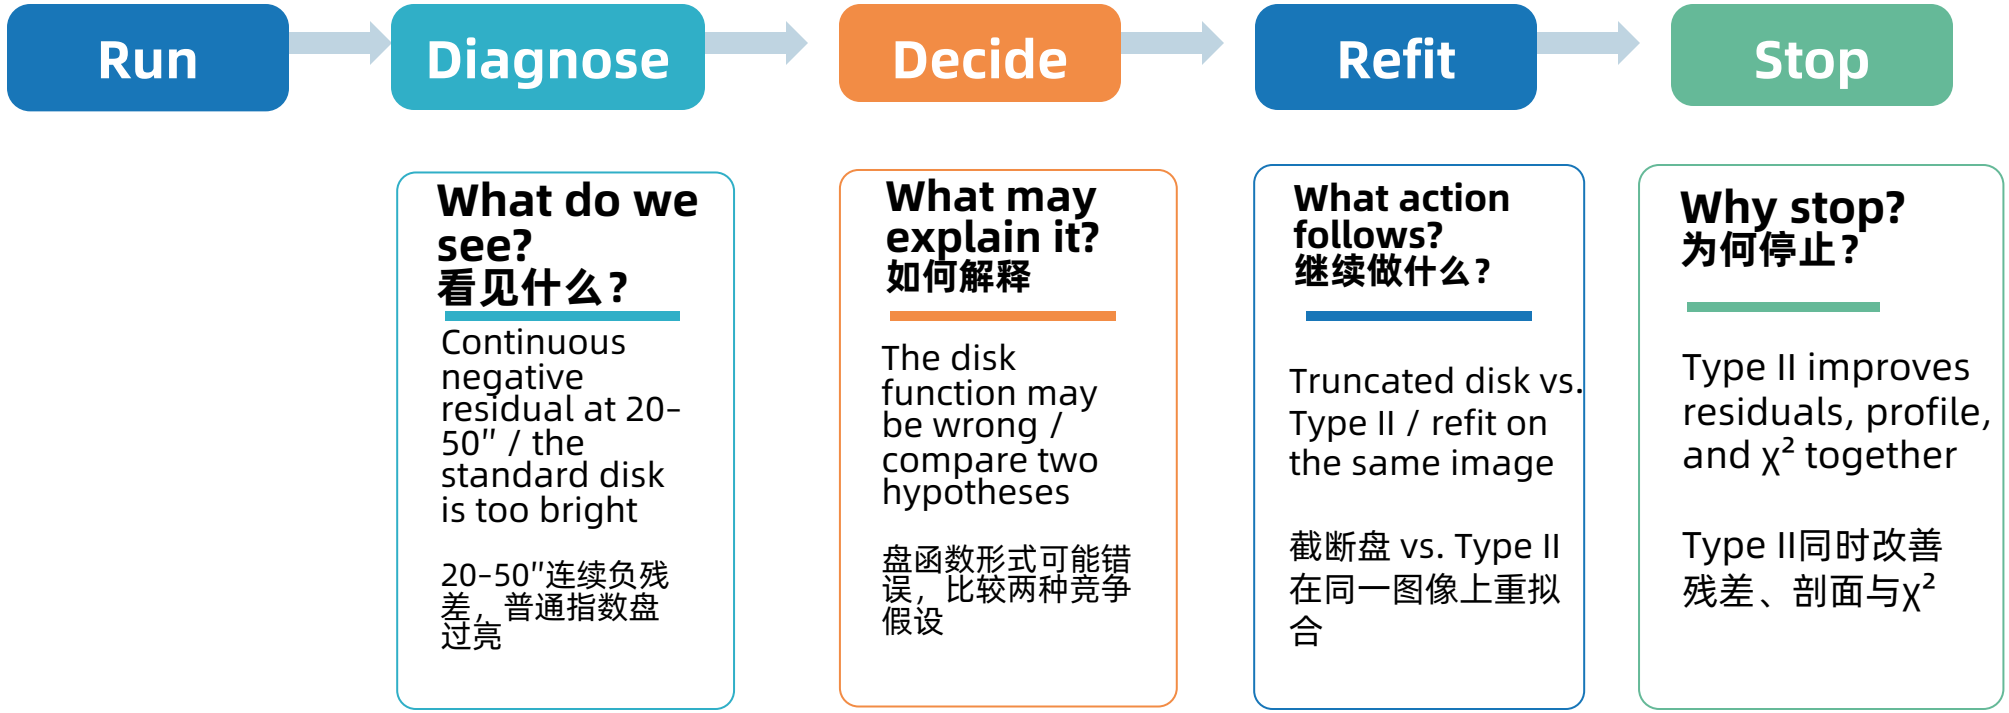

The real automation target is how judgment updates across iterations

真正需要自动化的不是一次软件调用，而是多轮反馈中如何更新判断

Step 0: make the process runnable, observable, and reviewable
我们的第0步：让过程可研究，先把专家闭环变成可运行、可观察、可复核的环境

Run → **Diagnose** → **Decide** → **Refit** → **Stop**

Machine execution
机器执行
MCP wraps GALFIT for reliable runs and refits

Model observation
模型观察
A VLM reads the image, model, residual, and profile

Model proposals and expert correction
模型提议&专家校正
Confirm observations, causes, and the next action round by round

Trajectory capture
轨迹沉淀
Save state, evidence, judgment, action, and result

CSST类型匹配率
93%

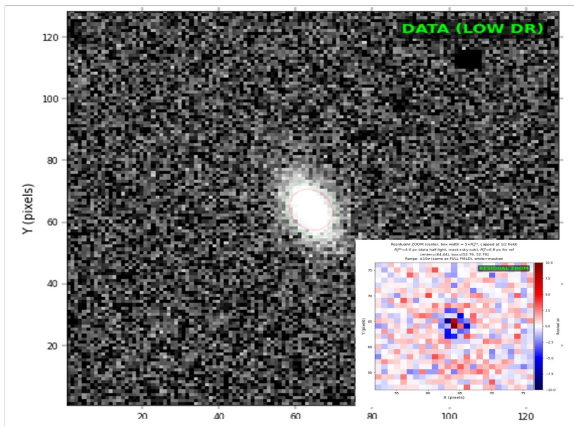

CSST 仿真数据 : Disk + Bulge

Gadotti（SDSS）类型匹配率
70%

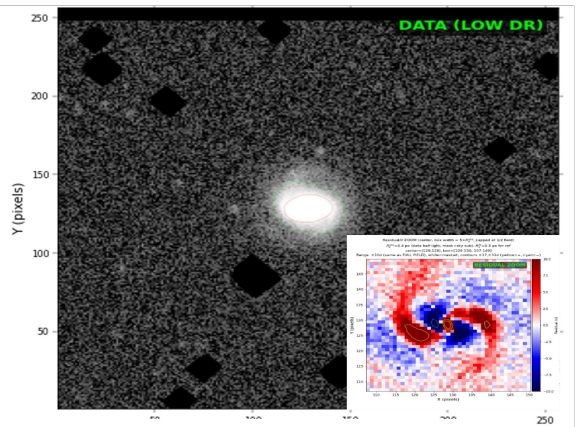

Gadotti: Disk + Bulge + Bar (+ Lens)

JWST类型匹配率
50%

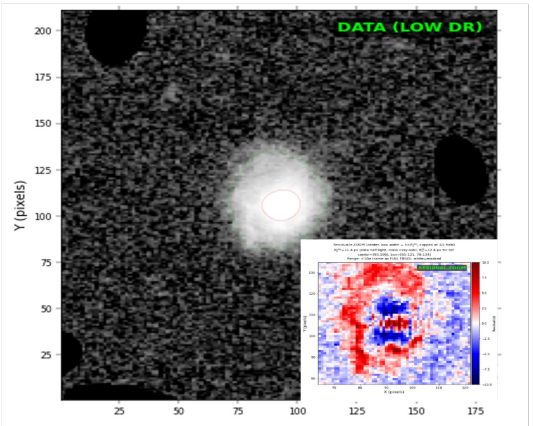

JWST: Disk(+Lop) (+Bar) (+Bulge)

数据集难度依次增大，CSST仿真数据干扰少; Gadotti 弱Bar 识别有难度； JWST会有更多拟合收敛

Agent system first turns an implicit expert process into a verifiable research object

智能体的第一价值不是“智能”，而是把隐性的专家过程变成可验证的研究对象

Real MCP runs show that expert experience cannot be copied into a complete rulebook
MCP真实运行表明：专家经验不是一本可以抄完的规则手册

Run → Diagnose → Decide → Refit → Stop

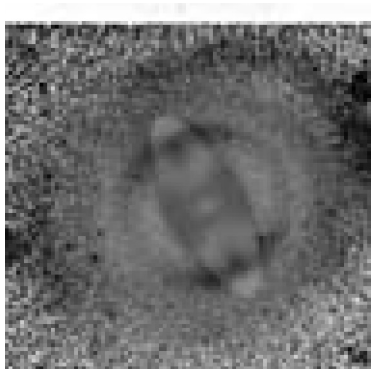

The same residual symptom
同一个残差症状

- **Coupled conditions / 规则条件耦合**
- **Subjective visual scale / 尺度直觉主观**
- **Depends on the science goal / 依赖科学目标**

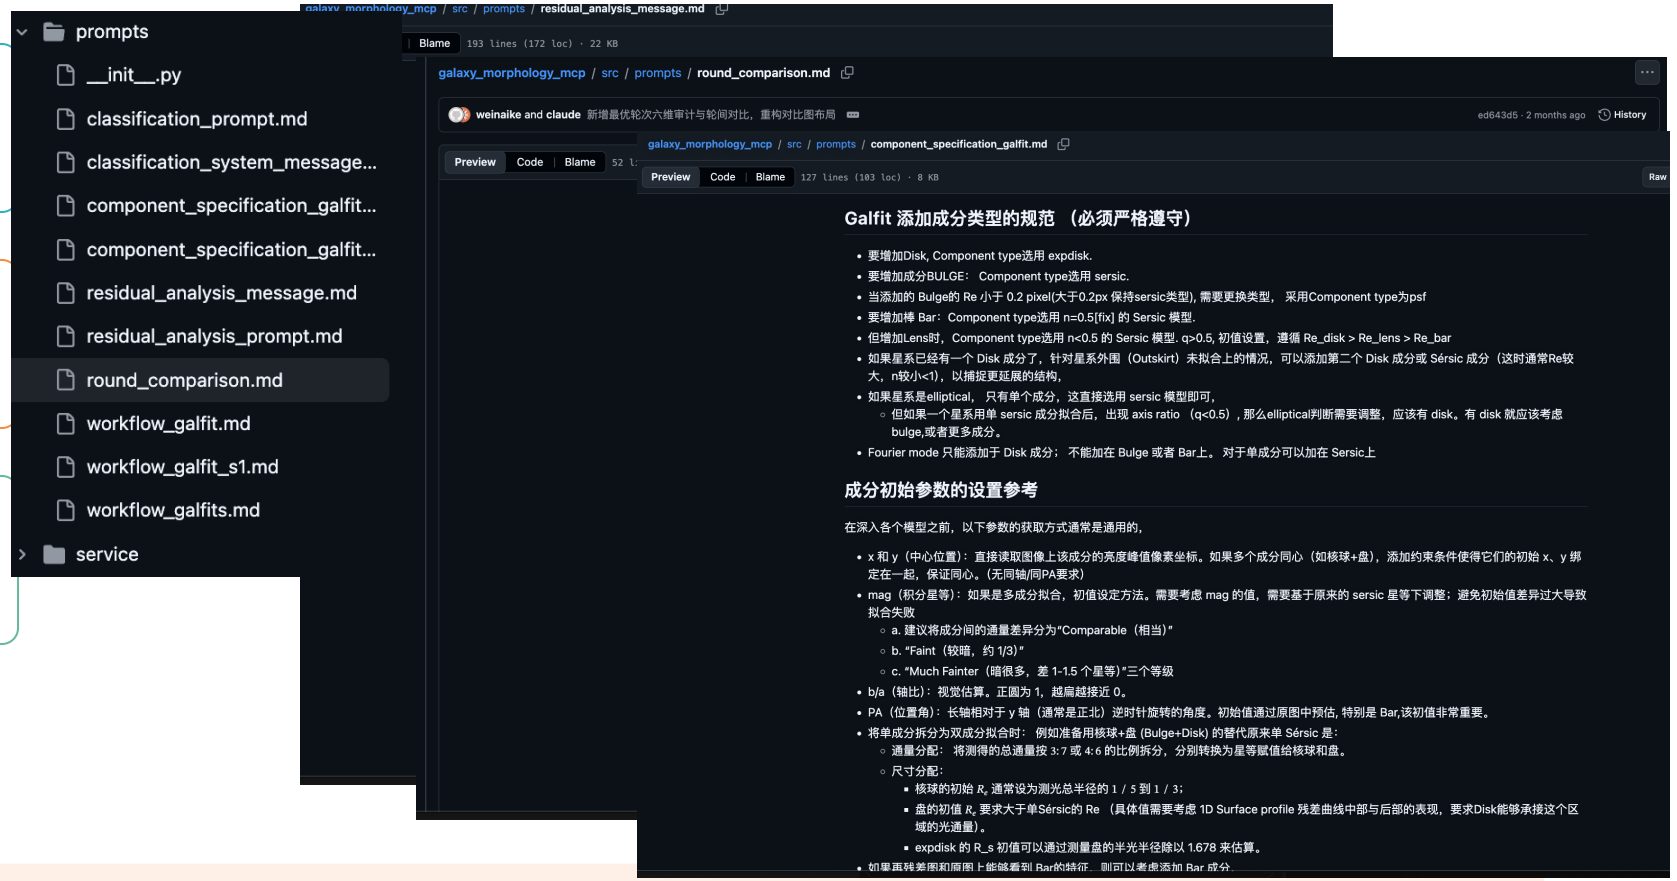

Rules define safety boundaries and known constraints, but cannot cover contextual judgment
规则适合定义安全边界与已知约束，却不适合穷尽情境化判断

Learn two abilities from expert trajectories: what to do next, and when the result is good enough

向专家轨迹学习两种能力：下一步怎么做？改善了吗以及做到什么程度可以停？

Run → Diagnose → Decide → Refit → Stop

Ability 1 | Decision

能力一 | 决策能力

Choose the most informative next action from evidence, history, constraints, and budget

根据证据、历史尝试、物理约束和预算，选择最有辨识力的下一动作。

Diagnose → Decide → Refit

Path 1 / 路径1

Knowledge-Guided Search

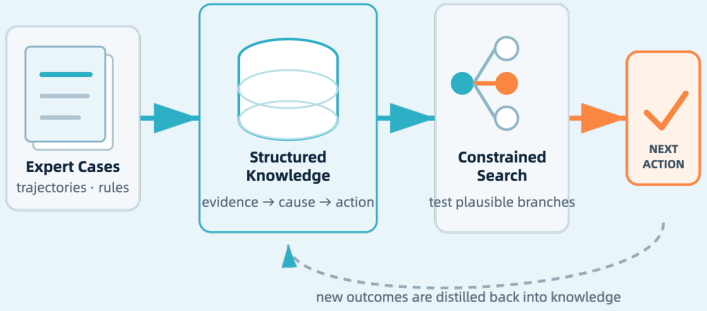

Path 2 / 路径2

Reinforcement Policy Learning

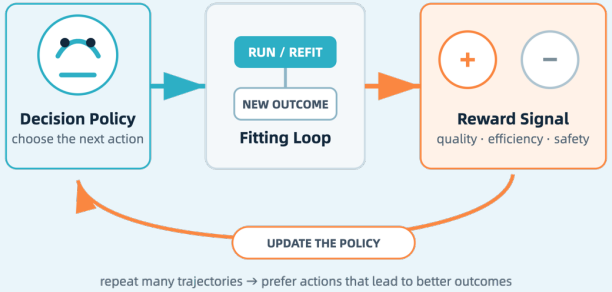

Ability 2 | Judgment

能力二 | 判断能力

Judge improvement and credibility, then accept, continue, or ask a human

理解是否改善、结构是否可信，以及应当接受、继续还是转人工。

Diagnose → Stop

VLM-Guided Judgment

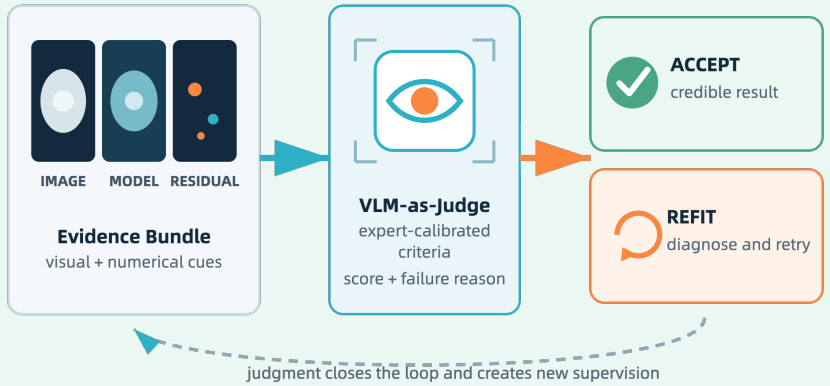

Decision learning, Path 1: rules become a knowledge base for search
决策能力学习路径1：规则不再充当最终答案，而是成为搜索可利用的知识起点

Run → **Diagnose** → **Decide** → Refit → Stop

Real expert trajectories / 真实专家轨迹

Evidence → cause hypothesis / action → outcome / failure → correction
证据 → 病因假设
动作 → 执行结果
失败 → 修正理由

Structured knowledge (85 items) / 结构化知识（85条）

Morphology and scale	Bar evidence
Residual patterns	Models and constraints
Nearby-source context	Quality and provenance

Feed the search / 进入搜索器

Limit legal actions / generate paths / set priorities / prune low-value branches
限定合法动作
生成候选路径
设置优先级
剪除低价值分支

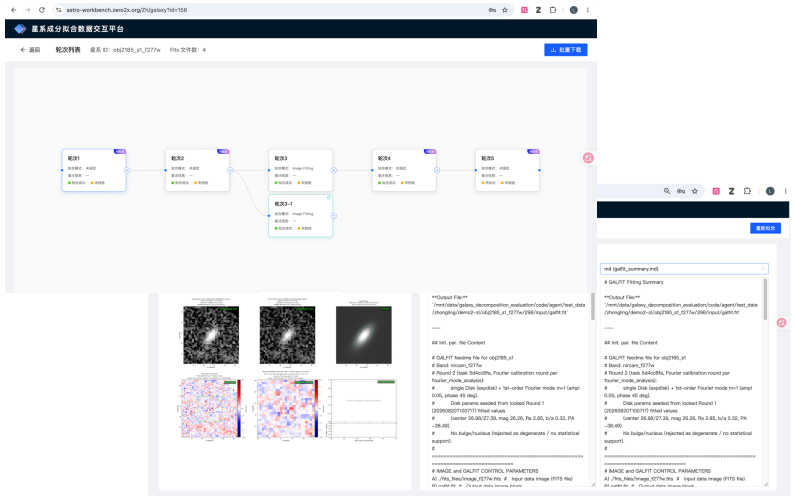



KILOGAS

> 84%

Knowledge extraction gives exploration a better starting point and clear boundaries
知识萃取不是写完一本规则书，而是给探索提供更好的起点和边界

Decision learning, Path 1: Beam Search tests expert-style hypotheses with real tool feedback

决策能力学习路径1：Beam Search用真实工具反馈验证多个专家式假设

Run → Diagnose → Decide → Refit → Stop

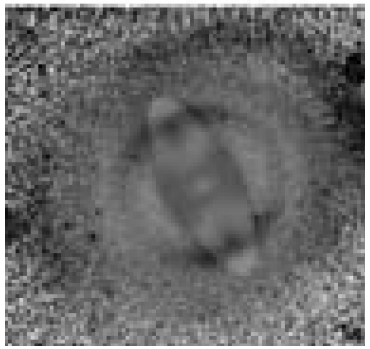

Current evidence / 当前证据

H1: unsuitable disk function
H1 盘函数不合适
Use a Type II disk
改用Type II盘

H2: PSF mismatch
H2 PSF不匹配
Refit with another PSF
更换PSF重拟合

H3: missing component
H3 成分遗漏
Add a candidate component
增加候选成分

MCP-GALFIT
获得真实拟合结果

MCP-GALFIT
获得真实拟合结果

MCP-GALFIT
获得真实拟合结果

Compare and prune / 比较与剪枝

Keep promising paths, expand again, stop at a quality threshold or budget limit
保留更有希望的路径
继续展开
达到停止条件
或预算上限

Search widens exploration and generates richer high-quality trajectories for later learning
搜索扩展了一次能够尝试的宽度，也为后续学习产生更丰富的高质量轨迹

KILOGAS

> 84%

Decision learning, Path 2: learn expert trajectories and internalize the policy in an LLM
决策能力学习路径2: 通过对专家轨迹的学习， 将策略内化到LLM中

Run → Diagnose → **Decide** → **Refit** → Stop

Expert and search trajectories / 专家 / 搜索轨迹

High-quality state-action sequences
优质状态—动作序列

SFT initialization / SFT初始化

Learn a strong start, legal actions, and common patterns
学习正确的起步
学习合法动作与常见模式

Environment interaction / 环境交互

MCP executes multi-round GALFIT runs
MCP执行多轮GALFIT

GRPO optimization / GRPO优化

Learn which multi-step policy produces the best final result
学习哪条策略最终更优
按最终质量更新策略

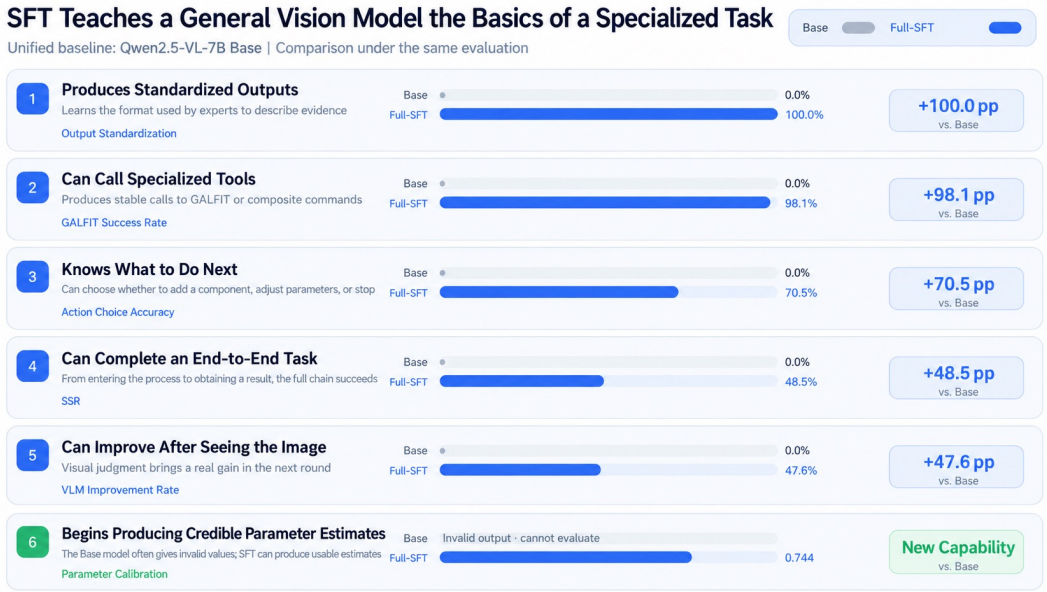

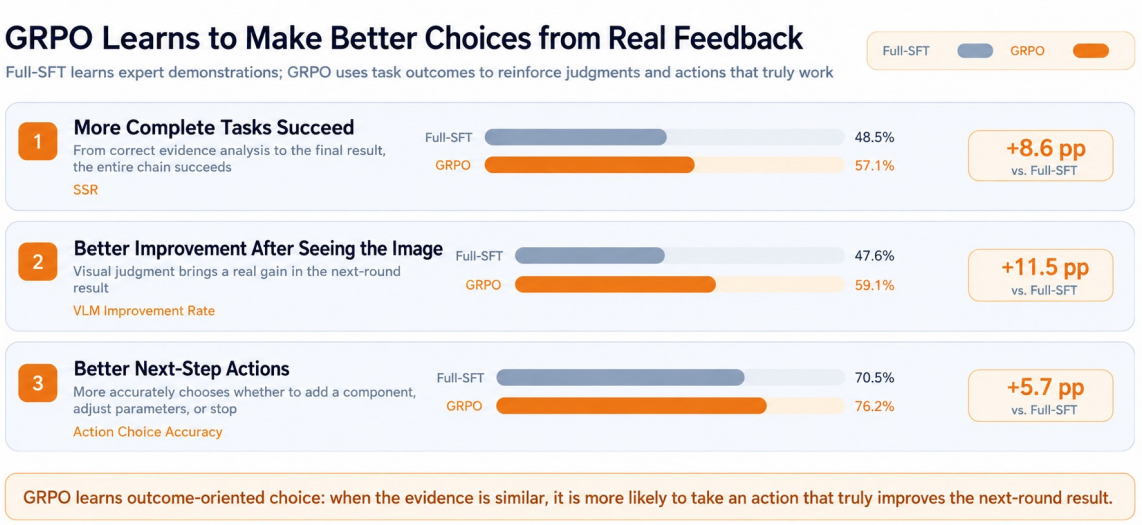

Challenge for Path 2: final quality is complex, and credit is hard to assign across steps
决策能力学习路径2的挑战：最终结果评价复杂，也难以把责任分配给中间每一步

Run → Diagnose → Decide → Refit → Stop

Round 1
Add a bar?
加棒？

Round 2
Change the disk?
改盘？

Round 3
Change the PSF?
换PSF？

Round 4
Tighten constraints?
收紧约束？

Final evaluation / 最终评价
Residual
visual quality
physical credibility
science goal

残差,视觉质量,物理可信
科学目标

Early actions may only show their value several rounds later
早期动作可能数轮后才显效

Challenge1: Long-term credit / 长程归因

Challenge2 Structured credit / 结构化归因

One action contains components, initial values, bounds, and constraints
一次动作含成分、初值、边界与约束

Action space	Round 1	Round 2	Round 3	Round 4	Round 5
Add/remove component	Action 1	·	·	·	·
Change disk/bulge model	·	·	Action 3	·	·
Adjust PSF/mask	·	Action 2	·	·	·
Change bounds	·	·	·	Action 4	·
Reinitialize	·	·	·	·	Action 5

Multi-objective evaluation / 多目标评价

Numerical, visual, physical, and scientific goals may disagree
数值、视觉、物理与科学目标并不总一致

Estimated final reward +1.0
最终估计的奖励 +1.0

How should it be assigned?

The challenge is whether reward represents scientific quality and assigns credit correctly
难点是reward能否代表科学质量并正确归因

Judgment learning: VLM-as-Judge learns an expert's scale of visual credibility
判断能力学习：VLM-as-Judge学习专家“看起来是否可信”的尺度

Run → **Diagnose** → Decide → Refit → Stop

Example A
Heavy central residual
中心残差过重
Continue

Example B
Clear improvement and physically reasonable
改善明显且物理合理
Accept

Example C
Conflicting evidence / new data domain
证据冲突 / 新数据域
Human review

VLM-as-Judge

Infer the hidden evaluation scale from a few comparisons and examples
从少量比较和示例中归纳隐含评价尺度

Improve / degrade / unchanged / accept / continue / human review
改善 / 退化 / 持平
接受 / 继续 / 转人工

What to learn / 跟专家学会：

Match expert judgments
跟随专家尺度

Stable repeated judgments
重复判断稳定

Calibration across data domains
跨数据域校准

Escalate when uncertain
不确定时转人工

The Judge provides a repeatable signal, while code still rejects hard physical violations
Judge提供可重复的评价信号；硬性物理约束与参数越界仍由程序否决

All learning returns to the Agent loop: search finds paths, training turns paths into capability, and the Judge provides a shared standard

所有的学习会再次汇聚在Agent框架中，搜索发现路径，训练把路径变成能力，Judge提供共同评价

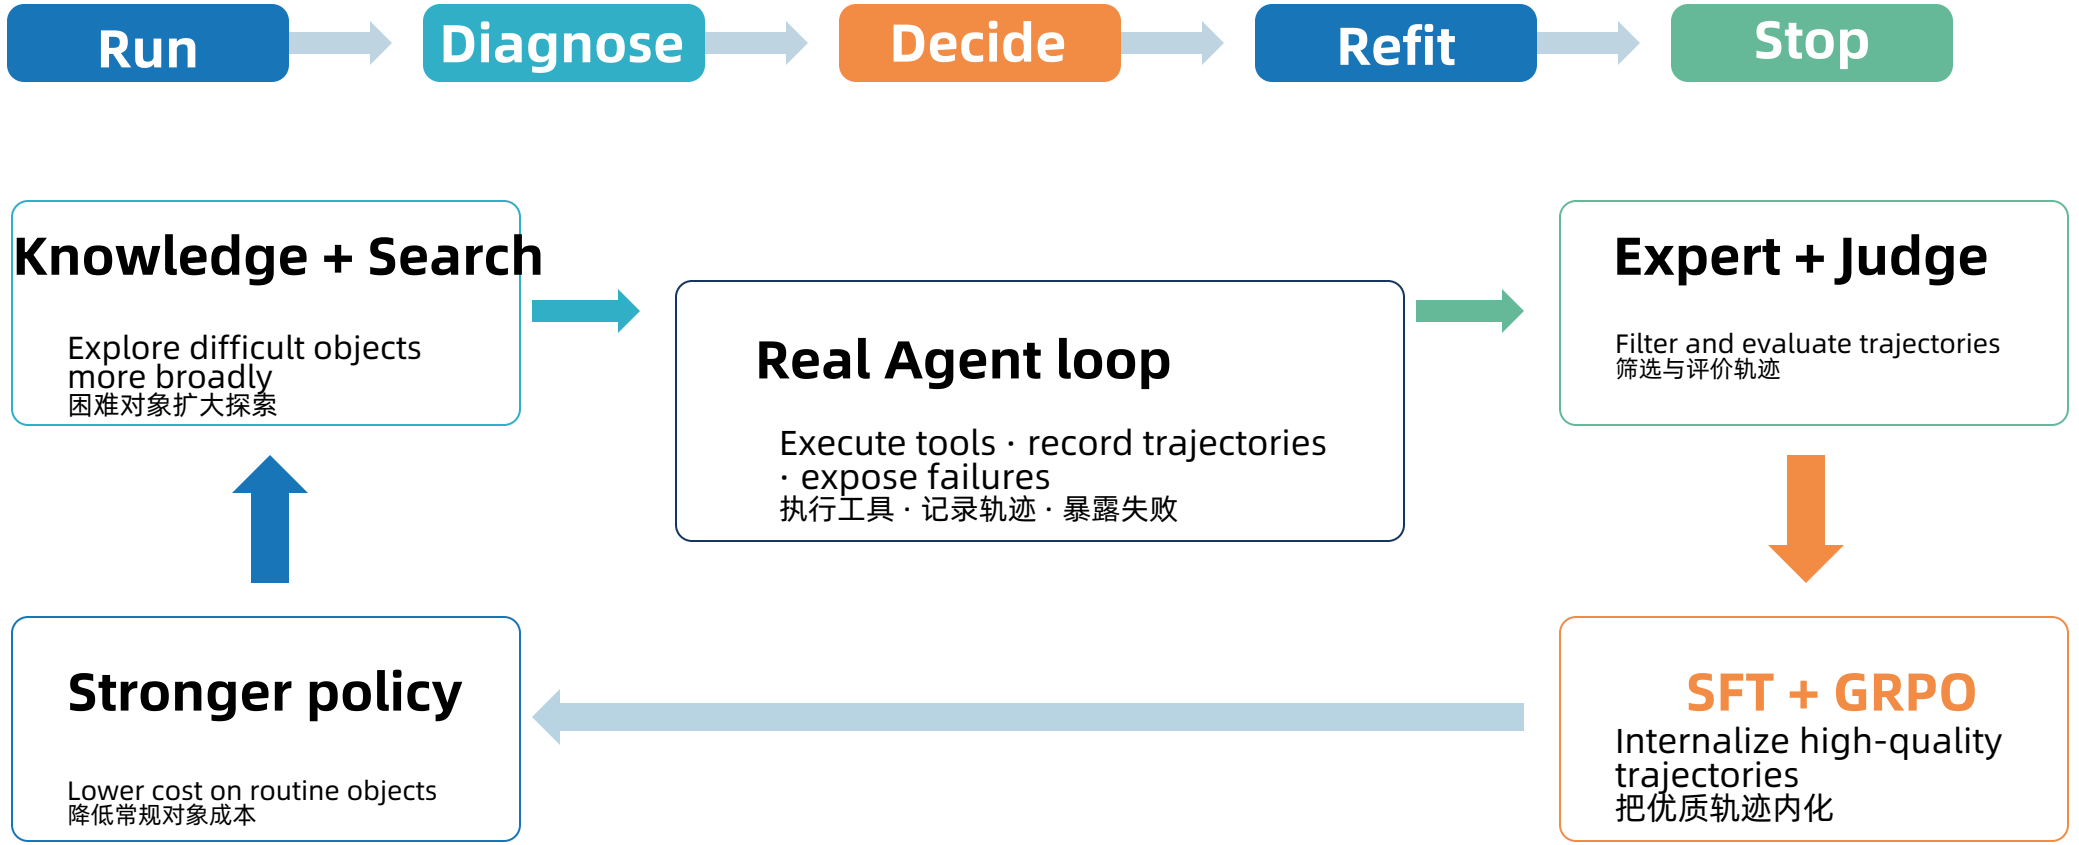

Knowledge-guided search expands experience; policy learning reduces its cost; the Judge aligns quality

知识搜索扩大经验边界；Plan B压缩经验成本；Judge让两者共享质量尺度

Summary：Real scientific problems test AI against tools, data, and expert standards
真实科学问题，让科学智能在工具、数据与专家标准中接受检验

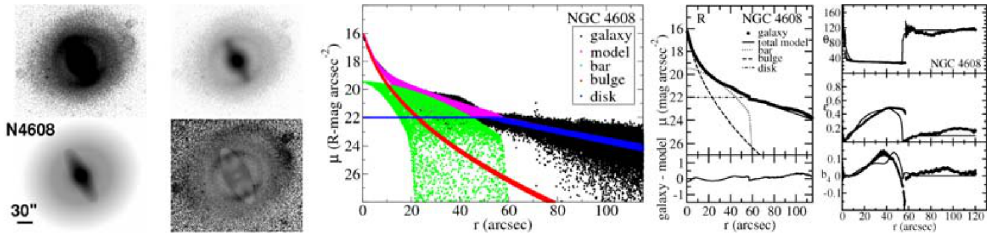

Engineering
工程

MCP makes tools, data, and multi-round work reliable and reviewable
MCP把工具、数据和多轮过程变成可靠、可复核的环境

Methods
方法

Search, decision learning, and the Judge grow from real failures and expert trajectories
搜索、决策学习与Judge从真实失败和专家轨迹中产生

Science
科学

Evaluation rests on real data, physical constraints, and credible conclusions
评价落在真实数据、物理约束与可信结论，而非模型是否会说

Useful scientific reasoning emerges only when models face tool feedback, data variation, and expert standards

只有进入真实问题，接受工具反馈、数据差异和专家标准的“摩擦”，
模型才可能形成真正可用的科学推理能力。

Thanks!

[Reference]
- Barden et al. (2012), GALAPAGOS: From Pixels to Parameters, https://doi.org/10.1111/j.1365-2966.2012.20619.x
- Simard et al. (2011), A Catalog of Bulge+disk Decompositions and Updated Photometry for 1,123,718 Galaxies in the SDSS, https://doi.org/10.1088/0067-0049/196/1/11
- Häußler et al. (2013), MegaMorph—multiwavelength measurement of galaxy structure: complete Sérsic profile information from modern surveys, https://doi.org/10.1093/mnras/sts633
- Vika et al. (2013), MegaMorph—multiwavelength measurement of galaxy structure: physically meaningful bulge–disc decomposition of galaxies near and far, https://doi.org/10.1093/mnras/stt1320
- Salo et al. (2015), The Spitzer Survey of Stellar Structure in Galaxies (S4G): Multi-component Decomposition Strategies and Data Release, https://doi.org/10.1088/0067-0049/219/1/4
- Casura et al. (2022), Galaxy And Mass Assembly (GAMA): Structural decomposition of KiDS galaxies, https://doi.org/10.1093/mnras/stac2077
- Euclid Collaboration: Bretonnière et al. (2023), Euclid preparation XXVI: Towards structural parameters for billions of galaxies, https://doi.org/10.1051/0004-6361/202245042