A. Training Details

The training process consists of two main stages: geometric prompt learning and adapter fine-tuning.

A.1 Geometry Prompt Learning

For geometry prompt learning (Stage I), we first construct multi-geometric graph structures using k-nearest neighbors (k=10) in three complementary spaces: Euclidean, Hyperbolic (k=-1), and Spherical (k=1). The input features of dimension 1024 are processed through geometric encoders with a hidden dimension of 512, eventually producing 256-dimensional geometric embeddings for each space. During training, we employ the AdamW optimizer with an initial learning rate of 1e-3 and weight decay of 1e-5. A step-wise learning rate scheduler is applied with step size 100 and decay factor 0.5. To ensure efficient training, we use a batch size of 1024 and train for 100 epochs.

A.2 Adapter Fine-tuning

For adapter fine-tuning (Stage II), we build upon Qwen2-VL-2B as our base model. The geometry adapter is integrated every four transformer layers, consisting of three types of feed-forward networks: a standard Euclidean FFN inherited from the pre-trained weights, a Hyperbolic FFN based on the Lorentz model with learnable scale, and a Spherical FFN utilizing von Mises-Fisher distribution with learnable concentration. These experts are orchestrated through a gating mechanism with temperature τ = 0.1. We utilize LoRA for efficient fine-tuning, with rank 16, alpha 32, and dropout 0.1. The training process employs mixed-precision (bf16) and runs with a batch size of 16 per GPU for 3 epochs. We use the AdamW optimizer with a learning rate of 2e-5 and weight decay of 0.01, incorporating 500 warmup steps. The maximum sequence length is set to 512 tokens.

A.3 Data Pre-processing

For data processing, images are resized to 144×144 pixels and normalized to [0,1] range. Spectral data is encoded into 1024-dimensional features and L2 normalized. These features are then projected into different geometric spaces through manifold-specific mappings, resulting in 256-dimensional representations in each space.

B. Riemannina Geometry

A smooth manifold M is referred to as a Riemannian manifold when it possesses a Riemannian metric g. Curvature c is an important measure of the degree of geodesic bending. For each point $x \in M$, there exists a tangent space $T_xM \subseteq \mathbb{R}^d$ that surrounds x, where the metric g is applied to determine the manifold's shape. The relationship between the tangent space and the manifold is established through the use of exponential and logarithmic maps. In particular, the exponential map at point x, represented as $\exp_x^c(\cdot) : T_xM \to M$, transforms points from the tangent space into the manifold, while the logarithmic map function is the inverse function of exponential map $\log_x^c(\cdot) = (\exp_x^c(\cdot))^{-1}$.

In this paper, we use three geometric spaces of different curvature to form a Riemannian expert: Euclidean space ($c = 0$), hyperbolic space ($c < 0$), and spherical space ($c > 0$).

B.1 Euclidean space

Euclidean space is based on Euclidean coordinates. Since the curvature is zero, the geodesic remains parallel. Euclidean space can be used to describe a flat universe very well. Each galaxy is influenced by its neighbors, capturing the local structure of galaxies in the universe. The exponential mapping of Euclidean Spaces is defined as:

$$\exp_{x_p}^c(x) = \mathbf{x}_p + \mathbf{x}. \tag{1}$$

B.2 Hyperbolic space

A hyperbolic space is defined as $\mathbb{H}_c^d = \{\mathbf{x}_p \in \mathbb{R}^{d+1} : \langle \mathbf{x}_p, \mathbf{x}_p \rangle_{\mathcal{L}} = 1/c\}$, where d represents the dimension and the inner product is defined as $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}} = -x_1 y_1 + \sum_{j=2} x_j y_j$). In a hyperbolic space, The geodesic distance between the two points is:

$$d(x, y) = \frac{1}{\sqrt{-c}} \operatorname{arccosh}\left(c * \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}\right). \tag{2}$$

Since the curvature is negative, the geodesic will diverge. This helps to describe the evolution of galaxies in a universe, and thus reflects the internal hierarchy of galaxies. Its exponential map is defined as:

$$exp_{x_p}^c(x) = \cosh\left(\sqrt{-c}||\mathbf{x}||\right)\mathbf{x}_p + \sinh\left(\sqrt{-c}||\mathbf{x}||\right)\frac{\mathbf{x}}{\sqrt{-c}||\mathbf{x}||}. \tag{3}$$

B.3 Sphere space

Spherical space is defined as $\mathbb{S}_c^d = \{\mathbf{x}_p \in \mathbb{R}^{d+1} : \langle \mathbf{x}_p, \mathbf{x}_p \rangle_{\mathbb{S}} = 1/c\}$, where the inner product is the standard Euclidean inner product $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{S}} = \sum_{j=1}^{d+1} x_j y_j$.The geodesic distance between the two points is:

$$d(x, y) = \frac{1}{\sqrt{c}} \arccos\left(c\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{S}}\right). \tag{4}$$

Geodesics in spherical space are convergent. Therefore, it can reflect the global information of the galaxy. Capture the overall star map content. Its exponential map is defined as

$$exp_{x_p}^c(x) = \cosh\left(\sqrt{c}||\mathbf{x}||\right)\mathbf{x}_p + \sinh\left(\sqrt{c}||\mathbf{x}||\right)\frac{\mathbf{x}}{\sqrt{-c}||\mathbf{x}||}. \tag{5}$$

C. Special Tokens and Templates

C.1 Special Modality Tokens

To effectively integrate multi-modal features into the input sequence, we carefully selected special tokens with relatively low frequency in the pre-trained vocabulary to represent different modalities. Specifically, we use token "Ã" for spectral features, "℘" for Euclidean geometric structure, "ø" for spherical geometric structure, and "æ" for hyperbolic geometric structure. During forward propagation, these tokens' embeddings are dynamically replaced with their corresponding modal embeddings after geometry-specific projection and normalization. This design allows the model to seamlessly incorporate multi-geometric and spectral information while maintaining the pre-trained model's linguistic capabilities.

C.2 Task-specific Templates

In this section, we present the task-specific templates for two main categories of tasks: galaxy property estimation and galaxy morphology classification. For each task, we utilize special modality tokens introduced in Section to incorporate different modalities. The key modalities include Image, Spectral (where available), and Geometry information. To better leverage knowledge from pre-trained models, we customize the descriptions of special modality tokens according to each task's characteristics, enabling Galaxy Walker to align multi-modal representations better.

C.2.1 Galaxy Property Estimation

Galaxy property estimation encompasses four regression tasks:

- **Stellar Mass (M^*) Prediction**: For numerical regression of the total mass of stars in a galaxy.
- **Mass-Weighted Stellar Metallicity (Z_{MW}) Prediction**: For estimating the abundance of heavy elements in stars.
- **Mass-Weighted Galaxy Age (t_{age}) Prediction**: For determining the mass-weighted average age of stars (in Gyr).
- **Specific Star-Formation Rate (sSFR) Prediction**: For calculating the rate of star formation per unit stellar mass.

For these prediction tasks, we employ a numerical head for regression. In the templates, we use the "num" token to represent all numerical values as the model's target response.

C.2.2 Galaxy Morphology Classification

Galaxy morphology classification includes ten distinct classification tasks:

1. Smooth (SMH)
2. Disk-Edge-On (DEO)
3. Spiral-Arms (SPR)
4. Bar (BAR)
5. Bulge-Size (BLG)
6. How-Rounded (RND)
7. Edge-On-Bulge (EOB)
8. Spiral-Winding (SWP)
9. Spiral-Arm-Count (SAC)
10. Merging (MRG)

For classification tasks, we structure the templates as multiple-choice questions, with the model required to select from options labeled (a), (b), (c), etc. The specific options for each classification task are presented in Table 1.

D.Additional Experiments

D.1 The Impact of Prompt

To investigate whether our designed prompts can effectively leverage knowledge from pre-trained models to improve performance, we conduct experiments comparing three different prompt settings:

- **Concat**: Directly concatenating modality tokens (Image Token, Spectral Token, Geometry Token) with the question text without any connecting words or explanations.
- **Simple Prompt**: Adding basic connecting words to describe what each modality token represents, building upon the Concat setting.
- **Prompt with Knowledge Background**: Using our carefully designed templates from Section , which incorporate detailed explanations of how each modality token contributes to the specific task, combined with relevant domain knowledge.

As shown in Figure 15, the Prompt with Knowledge Background setting consistently outperforms the other two approaches. In property estimation tasks, this setting achieves notably better results across all metrics, with particular improvements in sSFR prediction. For morphology classification tasks, while all three settings perform competitively, the Knowledge Background prompts still demonstrate advantages, especially in complex features such as Bar and Spiral Arm Count classification. These results suggest that carefully designed prompts incorporating domain knowledge can effectively help the model leverage pre-trained knowledge for better task performance.

D.2 Parameter-Efficient Training Strategy Analysis

We evaluate three different training strategies to investigate the effectiveness of our parameter-efficient approach:

- **Geometry Adapter Only**: Training only the newly added components including projection π_θ, Geometry Adapter, and Num Head weights.

CVPR
#9614

Table 1. Options for Galaxy Morphology Classification Tasks

Task	Options
Smooth	(a) Smooth (b) Featured or Disk (c) Artifact
Disk-Edge-On	(a) Yes (b) No
Spiral-Arms	(a) Yes (b) No
Bar	(a) Strong Bar (b) Weak Bar (c) No Bar
Bulge-Size	(a) Dominant (b) Large (c) Moderate (d) Small (e) None
How-Rounded	(a) Round (b) In-Between (c) Cigar-Shaped
Edge-On-Bulge	(a) Boxy (b) None (c) Rounded
Spiral-Winding	(a) Tight (b) Medium (c) Loose
Spiral-Arm-Count	(a) 1 (b) 2 (c) 3 (d) 4 (e) More than 4 (f) Can't Tell
Merging	(a) None (b) Minor Disturbance (c) Major Disturbance (d) Merger

Galaxy Walker: Stellar Mass Estimation

User: Stellar mass refers to the total mass of all the stars in a galaxy. It is a critical parameter for understanding galaxy formation and evolution and can be analyzed through multiple perspectives. Specifically, the **[Image token]** utilizes celestial image data to assess morphology and luminosity, which helps in the initial estimation of stellar mass. The **[Spectral token]** analyzes stellar spectral characteristics, such as absorption line width and radiation intensity, to directly infer mass parameters. The **[Euclidean token]** provides the object's position in flat space, aiding in the mass calculation by considering distance measurements. The **[Hyperbolic token]** describes the geometrical properties in negatively curved space, modeling more complex cosmic structures and helping to understand the distribution of massive stars in a negatively curved universe. The **[Sphere token]** uses spherical geometry in positively curved space to evaluate an object's position in the spherical coordinate system, leading to a more accurate mass estimation.

Assistant: NUM

Figure 1. Prompt template for stellar mass estimation.

- **Geometry Adapter + LoRA**: Training the Geometry Adapter components plus LoRA modules in attention and linear layers.
- **Full-Parameter Training**: Fine-tuning all model parameters.

As shown in Table 3, our Geometry Adapter + LoRA strategy achieves comparable or even superior performance to full-parameter training across most metrics. Notably, it outperforms full-parameter training in property estimation tasks, achieving better R^2 scores for all four properties ($\mathbf{M_*}$, $\mathbf{Z_{MW}}$, $\mathbf{t_{age}}$, **sSFR**). For morphology classification, the performance difference is minimal, with our parameter-efficient approach showing slight advantages in several categories (DEO, SPR, BAR). The Geometry Adapter Only setting, while using the fewest trainable parameters, still maintains strong performance, suggesting that the geometric adaptation components effectively capture domain-specific features. These results demonstrate that our parameter-efficient strategy can match or exceed the performance of full-parameter training while significantly reducing the number of trainable parameters and computational cost.

Hardware	Inference Time (s)
NVIDIA H100	0.38
NVIDIA A100	1.14
Ascend 910B	1.52

Table 2. **Inference time comparison across different hardware platforms.** Times are averaged over 100 runs with batch size 1. Ascend 910B results are measured using FP16 precision, while NVIDIA results use BF16 precision.

D.3 Inference Time Analysis

To evaluate the practical deployment potential of GalaxyWalker, we conduct inference time benchmarks across different hardware platforms. We measure the average inference time per sample using batch size 1, with BF16 precision on NVIDIA GPUs and FP16 precision on Ascend hardware.

The results demonstrate that GalaxyWalker achieves practical inference speeds across all tested platforms. The NVIDIA H100 shows superior performance with an average inference time of 0.38 seconds per sample, while the A100 and Ascend 910B maintain reasonable inference speeds at

Galaxy Walker: Mass-Weighted Stellar Metallicity Estimation

User: Mass-weighted stellar metallicity measures the abundance of elements heavier than hydrogen and helium in a galaxy's stars, weighted by their mass. This aids in understanding the galaxy's chemical evolution and can be analyzed through multiple perspectives. Specifically, the **[Image token]** helps observe the color and brightness variations of celestial objects, providing initial metallicity estimates. The **[Spectral token]** offers a detailed analysis of spectral lines, such as the strength and shift of metal lines, to directly infer the mass-weighted metallicity. The **[Euclidean token]** provides precise coordinates in flat space, aiding in the calculation of metallicity distribution within stars by using distance information. The **[Hyperbolic token]** describes the geometrical properties in negatively curved space, modeling complex star cluster structures and giving geometrical background support for metallicity distribution. The **[Sphere token]** employs spherical geometry in positively curved space to understand the distribution of celestial objects within the spherical coordinate system, leading to comprehensive metallicity estimation.
Assistant: NUM

Figure 2. Prompt template for mass-weighted stellar metallicity estimation.

Galaxy Walker: Mass-Weighted Galaxy Age Estimation

User: Mass-weighted galaxy age refers to the average age of stars within a galaxy, weighted by their mass, providing insights into the galaxy's formation history. This can be analyzed through multiple perspectives. Specifically, the **[Image token]** assesses morphology and color via celestial images to estimate the age distribution of stellar populations in the galaxy. The **[Spectral token]** primarily uses spectral analysis, such as examining the spectral energy distribution and absorption line changes, to determine the overall age of the galaxy. The **[Euclidean token]** provides the galaxy's coordinates in flat space, assisting in refining age estimation based on distance and position. The **[Hyperbolic token]** describes complex geometrical backgrounds in negatively curved space, aiding in the detailed understanding of mass-weighted age composition. The **[Sphere token]** utilizes positively curved space in spherical geometry to assist in distribution analysis and age estimation of different regions within the galaxy.
Assistant: NUM

Figure 3. Prompt template for mass-weighted galaxy age estimation.

1.14 and 1.52 seconds respectively. These results suggest that GalaxyWalker is suitable for real-world applications across various hardware configurations.

Galaxy Walker: Specific Star-Formation Rate Estimation

User: The specific star-formation rate (sSFR) is the rate of star formation per unit stellar mass in a galaxy, indicating how actively the galaxy is forming stars relative to its existing stellar mass. It can be analyzed through multiple perspectives. Specifically, the **[Image token]** helps analyze star-forming regions, morphology, and density variations via celestial images for initial estimation of the star-formation rate. The **[Spectral token]** provides detailed spectral analysis, especially the intensity and distribution of emission lines, to measure the current star-formation rate. The **[Euclidean token]** offers precise positioning in flat space, aiding in inferring the star-formation rate based on distance and velocity information. The **[Hyperbolic token]** describes geometrical properties in negatively curved space, modeling complex cosmic environments and star cluster structures for supporting star-formation rate estimation. The **[Sphere token]** utilizes positively curved space in spherical geometry to understand the distribution of star formation within the spherical coordinate system, assisting in specific rate determination.

Assistant: NUM

Figure 4. Prompt template for specific star-formation rate estimation.

Galaxy Walker: Galaxy Smoothness Classification

User: The morphological class of a galaxy can be analyzed using multiple tokens: the morphological classification of galaxies, such as spiral, elliptical, or irregular, can be directly observed by analyzing their images. Specifically, the **[Image token]** utilizes celestial images to assess the galaxy's overall shape and structural features, helping to classify it as smooth, featured, or an artifact. The **[Euclidean token]** offers the galaxy's precise coordinates in flat space, allowing for spatial analysis and comparison with known morphological classes. The **[Hyperbolic token]** provides insights into negative curvature space, aiding in the understanding of complex structures that might influence the galaxy's morphology. The **[Sphere token]** uses spherical geometry to interpret the galaxy's appearance in positively curved space, helping to refine its classification. **Please choose from these options:(a) Smooth (b) Featured or Disk (c) Artifact.**

Assistant:

[The choice of true label]

Figure 5. Prompt template for galaxy smoothness classification.

Galaxy Walker: Disk-Edge-On Classification

Prompt: Determining if a galaxy is disk-edge-on can be analyzed using multiple tokens: edge-on disk galaxies are characterized by their flat, edge-like appearance when observed. This can be directly identified from images. Specifically, the **[Image token]** offers visual information on the galaxy's edge-on appearance, which is indicative of a disk-edge-on orientation. The **[Euclidean token]** gives the galaxy's precise coordinates in flat space, assisting in spatial orientation analysis. The **[Hyperbolic token]** models the galaxy's structure in a negatively curved space, helping to understand any distortions that confirm its disk-edge-on nature. The **[Sphere token]** uses spherical geometry to analyze the galaxy's orientation in positively curved space. **Please choose from these options: (a) Yes, it is a disk-edge-on galaxy (b) No, it is not a disk-edge-on galaxy.**

Assistant: [The choice of true label]

Figure 6. Prompt template for disk-edge-on classification.

Galaxy Walker: Spiral Arms Classification

Prompt: Determining if a galaxy has spiral arms can be analyzed using multiple tokens: spiral-arm galaxies typically exhibit distinct spiral patterns in images. Specifically, the **[Image token]** provides visual information to identify the presence and patterns of spiral arms. The **[Euclidean token]** provides the galaxy's coordinates in flat space, aiding in spatial relationship analysis of spiral structures. The **[Hyperbolic token]** models the galaxy in a negatively curved space, providing geometric context for the spiral arms' formation. The **[Sphere token]** uses spherical geometry to interpret the distribution and winding of spiral arms in positively curved space. **Please choose from these options:(a) Yes, it is a spiral-arms galaxy (b) No, it is not a spiral-arms galaxy.**

Assistant: [The choice of true label]

Figure 7. Prompt template for spiral arms classification.

Galaxy Walker: Bar Type Classification

Prompt: Determining the type of bar in a galaxy can be analyzed using multiple tokens: bar structures in galaxies can be directly observed through images, revealing their length and strength. Specifically, the **[Image token]** offers visual information to observe and classify the bar's strength in the galaxy. The **[Euclidean token]** gives the galaxy's position in flat space, assisting in the spatial analysis of the bar. The **[Hyperbolic token]** helps model the galaxy's structure in negatively curved space, which aids in understanding the bar type. The **[Sphere token]** uses spherical geometry to analyze the distribution of stellar masses in the bar, refining its classification. **Please choose from these options: (a) Strong Bar (b) Weak Bar (c) No Bar.**
Assistant: [The choice of true label]

Figure 8. Prompt template for bar type classification.

Galaxy Walker: Bulge Size Classification

Prompt: Determining the bulge size of a galaxy can be analyzed using multiple tokens: the size of a galaxy's bulge can be observed in images by its prominence. Specifically, the **[Image token]** provides visual information to assess the bulge's prominence in the galaxy. The **[Euclidean token]** offers the galaxy's coordinates in flat space, assisting in spatial analysis of the bulge's physical size. The **[Hyperbolic token]** models the galaxy in negatively curved space, which helps in understanding the bulge size in a broader context. The **[Sphere token]** uses spherical geometry to analyze the distribution and density of stars within the bulge. **Please choose from these options:(a) Dominant Bulge (b) Large Bulge (c) Moderate Bulge (d) Small Bulge (e) No Bulge.**
Assistant: [The choice of true label]

Figure 9. Prompt template for bulge size classification.

Galaxy Walker: Galaxy Roundness Classification

Prompt: Determining the shape of a galaxy can be analyzed using multiple tokens: the shape of a galaxy can be directly observed by analyzing its images. Specifically, the **[Image token]** provides visual information to classify the galaxy as round, in-between, or cigar-shaped. The **[Euclidean token]** gives the galaxy's precise coordinates in flat space, aiding in the geometric analysis of its shape. The **[Hyperbolic token]** models the galaxy's structure in negatively curved space, providing a complex geometric context for its shape classification. The **[Sphere token]** uses spherical geometry to analyze the galaxy's three-dimensional shape in positively curved space. **Please choose from these options: (a) Round (b) In-Between (c) Cigar-Shaped.**
Assistant: [The choice of true label]

Figure 10. Prompt template for galaxy roundness classification.

Galaxy Walker: Edge-On Bulge Classification

Prompt: Determining the type of bulge in an edge-on galaxy can be analyzed using multiple tokens: the type of bulge in an edge-on galaxy can be identified by observing its images, which show whether it is boxy or rounded. Specifically, the **[Image token]** gives visual information to identify and classify the bulge as boxy, rounded, or absent in an edge-on galaxy. The **[Euclidean token]** offers the galaxy's coordinates in flat space, assisting in the spatial analysis of the bulge. The **[Hyperbolic token]** models the galaxy in negatively curved space, helping to understand the bulge type in a broader geometrical context. The **[Sphere token]** uses spherical geometry to analyze the three-dimensional distribution of stars within the bulge. **Please choose from these options:(a) Boxy Bulge (b) No Bulge (c) Rounded Bulge.**
Assistant: [The choice of true label]

Figure 11. Prompt template for edge-on bulge classification.

Galaxy Walker: Spiral Winding Classification

Prompt: Analyzing how tightly wound the spiral arms of a galaxy are can be done using multiple tokens: the tightness of spiral arms can be directly observed in images, showing the winding patterns clearly. Specifically, the **[Image token]** provides visual information to determine the tightness of the spiral arms. The **[Euclidean token]** gives the galaxy's precise coordinates in flat space, aiding in the spatial analysis of spiral arm winding. The **[Hyperbolic token]** models the galaxy's structure in negatively curved space, helping to understand the geometric properties affecting spiral arm tightness. The **[Sphere token]** uses spherical geometry to analyze the three-dimensional winding of the spiral arms. **Please choose from these options:(a) Tight Winding (b) Medium Winding (c) Loose Winding.**
Assistant: [The choice of true label]

Figure 12. Prompt template for spiral winding classification.

Galaxy Walker: Spiral Arm Count Classification

Prompt: Determining the number of spiral arms in a galaxy can be analyzed using multiple tokens: the number of spiral arms in a galaxy can be directly counted from images. Specifically, the **[Image token]** provides visual information to count and identify the number of spiral arms. The **[Euclidean token]** gives the galaxy's precise coordinates in flat space, aiding in the spatial analysis of the spiral arms. The **[Hyperbolic token]** models the galaxy's structure in negatively curved space, providing a geometric context for the number of spiral arms. The **[Sphere token]** uses spherical geometry to analyze the three-dimensional distribution of spiral arms. **Please choose from these options:(a) 1 Spiral Arm (b) 2 Spiral Arms (c) 3 Spiral Arms (d) 4 Spiral Arms (e) More than 4 Spiral Arms (f) Can't Tell.**
Assistant: [The choice of true label]

Figure 13. Prompt template for spiral arm count classification.

Galaxy Walker: Galaxy Merging State Classification

Prompt: Determining the merging state of a galaxy can be analyzed using multiple tokens: the merging state of a galaxy can be observed through signs of disturbance or merging in images. Specifically, the **[Image token]** provides visual information to observe signs of merging or disturbances. The **[Euclidean token]** offers the galaxy's coordinates in flat space, aiding in assessing merging stages from spatial data. The **[Hyperbolic token]** models the galaxy in negatively curved space, helping to understand the geometric properties affecting the merging state. The **[Sphere token]** uses spherical geometry to analyze the three-dimensional interactions of merging galaxies**. Please choose from these options: (a) No Merging (b) Minor Disturbance (c) Major Disturbance (d) Merger.**
Assistant: [The choice of true label]

Figure 14. Prompt template for galaxy merging state classification.

Training Strategy	**Property Estimation (R^2 Score)**				**Morphology Classification (F1 Score)**									
	$\mathbf{M_*}$	$\mathbf{Z_{MW}}$	$\mathbf{t_{age}}$	**sSFR**	**SMH**	**DEO**	**SPR**	**BAR**	**BLG**	**RND**	**EOB**	**SWP**	**SAC**	**MRG**
Geometry Adapter Only	0.89	0.67	0.50	0.81	0.74	0.95	0.94	0.68	0.81	0.80	0.85	0.77	0.62	0.75
Geometry Adapter + LoRA	**0.91**	**0.69**	**0.52**	**0.84**	0.76	**0.97**	**0.96**	**0.71**	**0.83**	**0.82**	**0.87**	**0.79**	**0.64**	**0.77**
Full-Parameter Training	0.90	0.68	0.51	0.82	**0.77**	0.96	0.95	0.69	0.82	0.81	0.86	0.78	0.63	0.76

Table 3. **Comparison of different training strategies.** Results show that our parameter-efficient approach (Geometry Adapter + LoRA) achieves comparable or even superior performance to full-parameter training while requiring significantly fewer trainable parameters. The best results for each metric are shown in **bold**.

CVPR
#9614

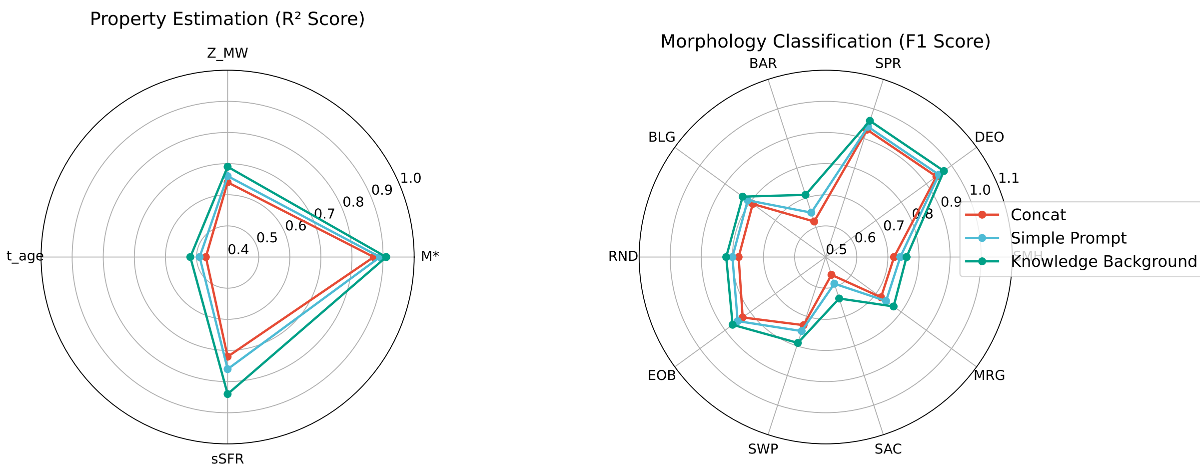

Figure 15. **Performance comparison of different prompt settings.** The radar plots show the performance of three prompt settings across property estimation tasks (left) and morphology classification tasks (right). The Knowledge Background prompts consistently outperform simpler approaches, demonstrating the effectiveness of incorporating domain knowledge into prompts. The improvement is particularly notable in property estimation tasks, where the Knowledge Background setting achieves superior performance in all metrics, with the most significant gains in sSFR estimation. For morphology classification, while the margins are smaller, the Knowledge Background setting still shows consistent advantages, especially in complex features like Bar and Spiral Arm Count classification.