CONTINUOUS THOUGHT SEGMENTATION WITH EXPERT-ORIENTED UNCERTAINTY PRIORS FOR LUNAR TERRAIN UNDERSTANDING

Anonymous authors
Paper under double-blind review

ABSTRACT

Semantic understanding of lunar terrain presents a fundamentally ill-posed representation learning challenge: absolute ground truth is prohibitively scarce, and morphological features are deeply entangled with varying spatial resolutions and transient observational conditions. Conventional segmentation networks, bounded by single-pass feed-forward mappings and standard confidence calibration, fail to resolve these observation-dependent ambiguities and systematically misrepresent the epistemic uncertainty inherent in expert interpretation. To address this, we introduce a novel dual-branch architecture that couples a Continuous Thought Machine (CTM) with a decoupled, expert-oriented uncertainty modeling module. To dynamically adapt to resolution and observational shifts, the CTM employs an internal temporal reasoning mechanism. By iteratively refining visual representations across multiple latent "thought steps," it constructs robust, scale-aware feature manifolds rather than rigid, static mappings. Concurrently, to navigate severe label ambiguity, our decoupled branch learns fine-grained spatial uncertainty maps without degrading the primary segmentation objective. Bootstrapped from model-derived epistemic cues, this branch establishes a robust machine-derived uncertainty prior. This decoupled design lays the architectural groundwork to accommodate future expert annotations, explicitly paving the way to bridge conventional prediction entropy with human perceptual priors. Ultimately, our framework establishes an interpretable, multi-step reasoning paradigm that jointly achieves high-fidelity segmentation and expert-aware uncertainty representation in heavily observation-dependent, label-ambiguous regimes.

1 INTRODUCTION

The lunar surface has undergone prolonged impact cratering, volcanism, and tectonic evolution, producing diverse landforms such as impact craters, volcanic craters, fracture systems, rilles, and wrinkle ridges. Automated recognition and fine-grained segmentation of these structures are essential for lunar geological analysis, mission planning, and resource exploration (MAH, 2019). Advances in remote sensing have enabled the integration of complementary data sources: optical imagery captures surface texture, while digital elevation models (DEMs), slope maps, and hillshade maps characterize topographic geometry (Smith et al., 2010). Nevertheless, lunar terrain segmentation remains challenging due to substantial scale variation, morphological similarities, and observation-dependent degradation. Reliable terrain-understanding systems must therefore provide accurate segmentation alongside insights into their decision processes and uncertain regions (Robinson et al., 2010).

While Convolutional and Transformer-based networks (e.g., SegFormer (Xie et al., 2021), SAM (Kirillov et al., 2023)) have substantially improved visual representation learning, they largely rely on a static, single-pass prediction paradigm. This fails to explicitly model how decisions evolve when interpreting complex visual structures. Continuous reasoning offers an alternative by introducing an internal temporal dimension where representations evolve through successive computational steps (Chen et al., 2019). The Continuous Thought Machine (CTM) implements this principle through neuron-level temporal dynamics (Darlow et al., 2026). However, its application to dense prediction remains underexplored for complex scientific imagery.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

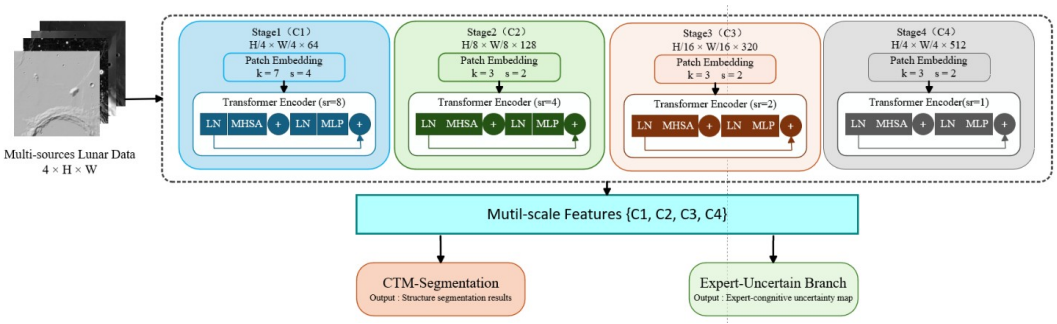


Figure 1: Overall Architecture of the Dual-Branch Multi-Output Segmentation Model.

Uncertainty estimation presents a related challenge. Predictive entropy primarily quantifies the model's own uncertainty, which does not necessarily reflect the subjective uncertainty experienced by geological experts dealing with incomplete morphological structures Kendall & Gal (2017); Gal & Ghahramani (2016). Incorporating expert feedback into uncertainty learning is therefore critical for developing scientific AI systems that better reflect domain-specific interpretation (Amershi et al., 2014).

To address these challenges, we propose a multisource lunar terrain-understanding framework integrating CTM-based continuous reasoning with a decoupled uncertainty-learning branch. The CTM iteratively reads visual features as its internal state evolves, generating temporally evolving segmentation predictions to form an observable internal reasoning trajectory. In parallel, the independent uncertainty branch learns spatial ambiguity representations from model-derived priors, designed for subsequent refinement through expert corrections. Together, these components yield an interpretable, expert-oriented terrain segmentation system.

2 Data

Our dataset comprises four complementary global cylindrical-projection lunar observation layers: optical imagery, Digital Elevation Model (DEM), slope maps, and hillshade maps. These layers have a spatial resolution of 100 m per pixel and have undergone spatial registration and preprocessing, and combined with expert-annotated vector labels to construct the dataset. Following the label locations, the global images were divided into 1024×1024 patches for training and validation. The dataset exhibits a pronounced long-tailed distribution, substantial scale variation, and ambiguous boundaries. Some classes occur only in a limited number of regions. These characteristics make the dataset ideal for evaluating dynamic reasoning processes and uncertainty estimation in complex scientific imagery. Detailed patch-generation procedures are provided in Appendix A.

3 Methodology : Continuous-Thought Segmentation with Expert-Aware Uncertainty Modeling

3.1 Overall Framework

We propose a dual-branch framework consisting of a continuous-thought segmentation branch and an expert-oriented uncertainty modeling branch (Figure 1). The four complementary modalities are first processed by a SegFormer-based encoder. The extracted multiscale features are then shared by two independent branches.

The segmentation branch applies a Continuous Thought Machine (CTM) over the baseline network to generate multiple predictions across sequential thought ticks (Abdić et al., 2015). A Quality Head estimates the relative quality of these states for inference selection. Meanwhile, the Expert Uncertainty Branch independently utilizes the same visual features to predict spatial uncertainty without affecting segmentation capacity, providing a foundation for future human-in-the-loop calibration (Christiano et al., 2023; Kohl et al., 2019).

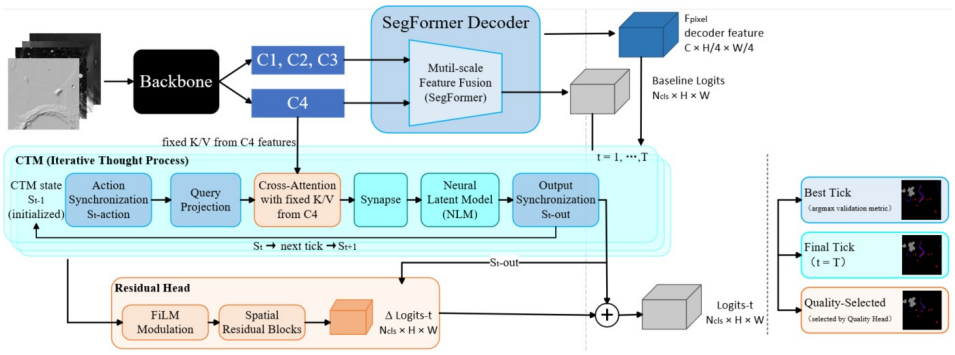


Figure 2: Architecture of the Continuous Thought-Based Model for Terrain Segmentation.

3.2 Residual Segmentation Driven by Continuous Thought

Instead of modifying SegFormer's basic capability, we use it to provide stable visual representations and baseline predictions, building an additional continuous-thought process with CTM to refine the final segmentation through residual updates.

3.2.1 CTM-Driven Segmentation

The base model extracts multiscale features and a baseline segmentation from the four-channel input. The high-level features remain fixed for the CTM. During one forward pass, CTM executes multiple consecutive thought ticks. Unlike conventional recurrent methods, each tick forms a new synchronization representation from the continuously evolving internal neural state, repeatedly retrieving visual information relevant to the current decision.

CTM updates its internal state through neural activity histories and Neuron-Level Models, generating the output synchronization for the current tick. To convert this global state into a dense segmentation, we use a conditionally modulated residual-prediction structure. After local convolutional processing, the final output is written as:

$$z_t = z_{\text{base}} + \Delta Z_t$$

where z_{base} is the baseline prediction and ΔZ_t is the CTM residual at thought step t. The final 1×1 convolution is zero-initialized, ensuring CTM residuals start at zero. During training, the residual branch gradually learns to correct the baseline. The resulting multistep segmentation is a sequence of visual interpretations driven by internal thought states, not merely progressive mask refinement (Figure 2).

3.2.2 Adaptive Result Selection Based on Quality Prediction

Since interpretation difficulty varies, different samples require different reasoning depths. Selecting a fixed intermediate tick fails to fully exploit CTM's potential. We introduce a Quality Head to predict the sample-level IoU gain of each thought tick relative to the baseline:

$$G_t = \text{IoU}_t - \text{IoU}_{\text{base}}$$

As shown in Figure 3, the Quality Head processes the current CTM state and residual features with detached gradients, ensuring its training does not alter the main segmentation pathway. During inference, it selects the result with the largest predicted gain:

$$t^* = \arg\max_t G_t$$

If all ticks predict negative gains, the model falls back to the baseline. This adaptive mechanism prevents stable baseline predictions from being forcibly overwritten when further reasoning provides no benefit.

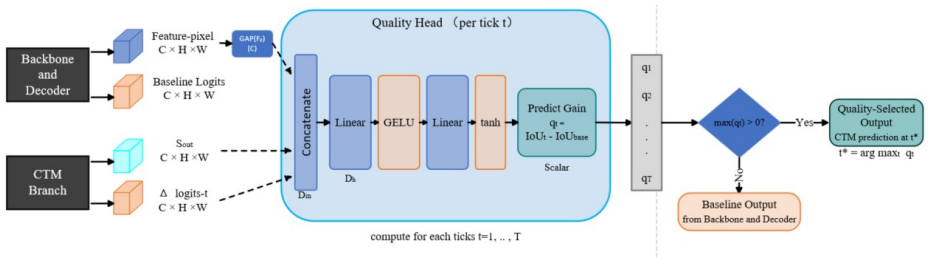


Figure 3: Schematic Illustration of the Quality Selected Head Architecture and Its Interaction with the CTM Segmentation Model.

3.3 Expert-Cognitive Uncertainty Modeling Branch

To quantify model reliability beyond prediction probabilities, we introduce an independent Expert Uncertainty Branch to learn spatial uncertainty representations. The branch uses multiscale visual features from the SegFormer backbone to characterize uncertainty arising from local ambiguity and class confusion.

The lightweight uncertainty decoder projects each of the four backbone feature maps to 64 channels, resizes them to the C_1 spatial resolution, and concatenates them. Following a 1×1 fusion layer and two residual spatial blocks, a final convolution outputs single-channel uncertainty logits (converted via Sigmoid during inference). Given the initial absence of human annotations, we employ the trained CTM model to generate cold-start pseudo-labels by fitting prediction entropy and spatial ambiguity across final reasoning steps (detailed in Appendix B).

3.4 Training Strategy

We employ a staged training strategy to preserve branch independence and allow cold-start supervision to imitate the segmentation model's inference behavior.

3.4.1 Training the CTM Segmentation Branch

We first freeze the CTM module to pretrain the baseline segmentation model, then unfreeze CTM to learn residual refinements. Segmentation is jointly supervised by BCE and a low-weight Dice loss (to address class imbalance). Since optimal reasoning depth varies, we do not force all samples to converge at a fixed tick. Instead, the CTM segmentation objective averages the final-tick loss and the lowest intermediate-tick loss ($L_{seg} = 0.5L_{best} + 0.5L_{final}$). Meanwhile, the Quality Head is trained independently using Smooth L1 loss to predict the IoU gain.

3.4.2 Uncertainty-Branch Cold Start and Architecture for Future Expert Calibration

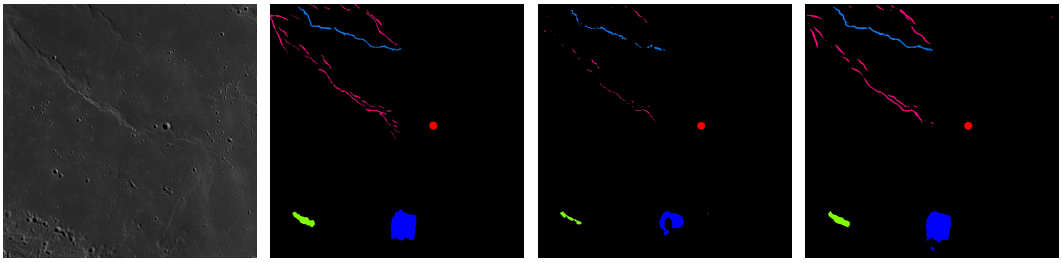
Initially, the model's estimated uncertainty distribution serves as a cold-start pseudo-label to pretrain the Uncertainty Decoder, with the segmentation backbone frozen. Consistency is constrained via a weighted combination of pixel-level regression, relative ranking, and spatial smoothness losses. These pseudo-labels merely provide a robust machine-derived prior. As expert interaction data accumulates, training will proceed to a calibration stage where manual supervision overrides pseudo-labels. This explicitly defines a pathway to gradually shift from machine-derived uncertainty toward expert-calibrated representations.

4 Experiments and Results

Detailed training configurations are described in Appendix C.3. Our experimental evaluation specifically adopts SegFormer as the singular baseline to construct a strictly controlled setting. Rather than pursuing state-of-the-art benchmarking across disparate architectures, our primary objective is to investigate the relative reasoning delta introduced by the CTM mechanism when built upon a standard

Table 1: Comparison of segmentation performance between the SegFormer baseline and the proposed CTM-based methods.

Model	Train mIoU	Valid mIoU	Valid Recall
SegFormer Baseline	56.19%	48.52%	63.47%
CTM Final Tick	76.25%	70.24%	79.43%
CTM with Quality Head	75.82%	71.54%	79.68%



(a) Image (b) Mask (c) Segformer-results (d) CTM-results

Figure 4: Comparison of a group of results from the SegFormer baseline and the CTM-enhanced segmentation model. From left to right, the four columns show the input imagery, the ground-truth (GT) mask, the SegFormer baseline prediction, and the segmentation result produced by our model with the CTM mechanism.

hierarchical Transformer. This isolation allows us to exclusively attribute performance gains to the multi-step continuous thought process.

4.1 CTM-Segmentation Results

Table 1 summarizes segmentation performance. The final CTM thought tick achieves an mIoU of 70.24%, substantially improving over the baseline SegFormer (48.52%). With the Quality Head for adaptive selection, the mIoU further increases to 71.54%, accompanied by a similar improvement in Recall.

Figure 4 presents qualitative comparisons. CTM produces refined predictions in regions with complex structures, improving boundary continuity and recovering weakly activated landforms.

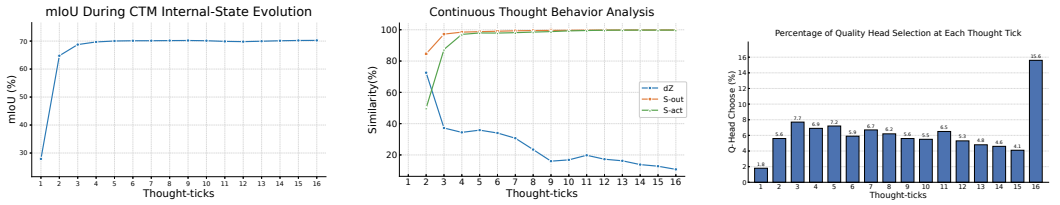
4.2 Analysis of Continuous Thought and the Quality Head

To verify that CTM develops meaningful internal dynamics rather than repeatedly producing identical outputs, we analyze segmentation changes, latent state dynamics, and Quality Head behavior (Figure 5).

Performance improves rapidly during the early ticks and saturates later. Both predictive performance and latent states exhibit large variations initially, stabilizing as reasoning proceeds. Notably, the final tick is selected for only 15.6% of samples, indicating that effective reasoning depth is highly sample-dependent.

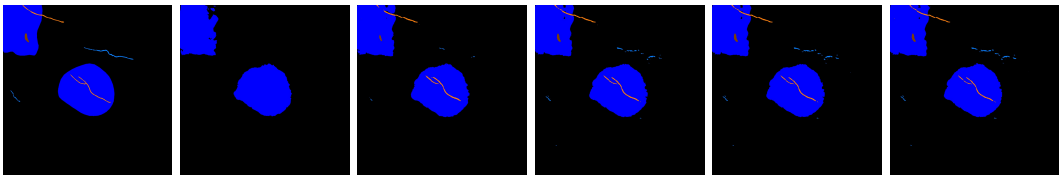
Crucially, from a planetary geology perspective, the intermediate thought ticks of CTM exhibit a progressive reasoning pattern highly analogous to a human expert's visual survey logic. Rather than merely smoothing statistical boundaries, the temporal evolution reflects a geologically meaningful macro-to-micro morphological decomposition, as visualized in Figure 6.

For example, when analyzing complex terrain systems containing nested and small-scale structures, we observe that early reasoning steps(e.g. Figure 6b) tend to delineate high-contrast, large-scale geological features, such as lunar domes. As the internal states evolve, intermediate reasoning steps progressively identify secondary structures embedded within larger formations, successfully capturing fine-scale lunar rilles(Figure 6c). In later stages, the model further refines structural boundaries



(a) Changes in Segmentation (b) Internal-State Dynamic (c) Quality-Based Selection

Figure 5: Analysis of the CTM Internal States.



(a) GT-mask (b) t1 (c) t4 (d) t6 (e) t12-selected (f) t16-final

Figure 6: Visualization of a representative continuous reasoning process. Several representative thought ticks are selected, including those with significant changes, Quality Head-selected ticks, and the final tick. The visualization shows that the segmentation results of nested terrains are progressively refined throughout the multi-step reasoning process.

by resolving ambiguous, low-contrast edge features(Figure 6d). This temporal disentanglement of features suggests that CTM implicitly aligns its iterative computation with the hierarchical organization of geological structures, providing a scientifically interpretable perspective on the dense prediction process.

Furthermore, we evaluate the selection accuracy of the Quality Head. A selection is considered fully correct when the selected tick achieves the highest mIoU among all reasoning steps. If the mIoU gap between the selected tick and the optimal tick is within 0.5%, we regard it as a near-optimal selection. The accuracies under these two criteria are 38.1% and 96.7%, respectively. The high near-optimal selection accuracy is mainly because the segmentation results are rapidly optimized and nearly converged during early reasoning stages (Figure 5a). Consequently, the mIoU remains relatively stable in later thought steps, suggesting that the model has already formed a sufficiently refined understanding of the terrain structures.

4.3 Cold-Start Analysis of the Uncertainty Branch

We evaluate whether the independent uncertainty decoder can learn the spatial uncertainty prior while the segmentation pathway remains frozen (Table 2). For the detailed calculation method of the evaluation criteria and visible rtesults, please refer to Appendix E and Appendix D.2.

The results demonstrate effective recovery of the machine-derived uncertainty distribution. This confirms the branch successfully establishes a stable initialization prior, ready for subsequent refinement via expert feedback.

5 Conclusion and Discussion

In this work, we advance the representation learning paradigm for ill-posed, label-ambiguous scientific domains by introducing a dynamic, dual-branch architecture for lunar terrain understanding. Our primary contribution lies in transcending the limitations of static, single-pass feedforward networks. By integrating a Continuous Thought Machine (CTM) modulated by an adaptive Quality Head, we formulate dense prediction as an explicit, multi-step latent reasoning trajectory that dynamically allocates computation depth based on morphological complexity. Furthermore, we address the critical gap between model predictive entropy and human cognitive uncertainty by designing a decoupled Expert Uncertainty Branch. This module successfully establishes a robust

Table 2: Quantitative Analysis of Training During the Cold-Start Stage of the Uncertainty Branch

Stage	MAE	Spearman	Pearson
Train	0.009871	0.798735	—-
Valid	0.009920	0.802779	0.840097

machine-derived spatial prior without entangling or degrading the primary visual features, explicitly laying the architectural groundwork for future human-in-the-loop (HITL) calibration. Ultimately, this framework not only yields high-fidelity segmentation in heavily degraded observational regimes but also provides a scalable blueprint for developing transparent, expert-aware AI systems across the broader scientific landscape.

AI USE STATEMENT

In this work, we utilized generative AI to assist with web searches for information and the analysis of prior research. AI tools also aided in the translation of the paper.

However, we did not use AI to assist with tasks such as code generation, data analysis, structuring the logic, or creating figures.

REFERENCES

Parametric study of water vapor and water ice particle plumes based on dsmc calculations: Application to the enceladus geysers. *Icarus*, 319:729–744, 2019. ISSN 0019-1035. doi: https://doi.org/10.1016/j.icarus.2018.10.022. URL https://www.sciencedirect.com/science/article/pii/S0019103518302173.

Irman Abdić, Lex Fridman, Erik Marchi, Daniel E Brown, William Angell, Bryan Reimer, and Björn Schuller. Detecting road surface wetness from audio: A deep learning approach, 2015. URL https://arxiv.org/abs/1511.07035.

Saleema Amershi, Maya Cakmak, W. Bradley Knox, and Todd Kulesza. Power to the people: The role of humans in interactive machine learning. *AI Magazine*, 35(4):105–120, 2014. doi: https://doi.org/10.1609/aimag.v35i4.2513. URL https://onlinelibrary.wiley.com/doi/abs/10.1609/aimag.v35i4.2513.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations, 2019. URL https://arxiv.org/abs/1806.07366.

Paul Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences, 2023. URL https://arxiv.org/abs/1706.03741.

Luke Darlow, Ciaran Regan, Sebastian Risi, Jeffrey Seely, and Llion Jones. Continuous thought machines. *Advances in Neural Information Processing Systems*, 38:21548–21594, 2026.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning, 2016. URL https://arxiv.org/abs/1506.02142.

Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision?, 2017. URL https://arxiv.org/abs/1703.04977.

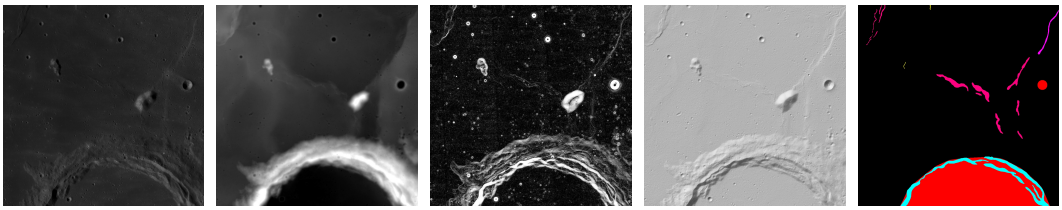
Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *2023 IEEE/CVF international conference on computer vision (ICCV)*, pp. 3992–4003. IEEE, 2023.

Simon A. A. Kohl, Bernardino Romera-Paredes, Clemens Meyer, Jeffrey De Fauw, Joseph R. Ledsam, Klaus H. Maier-Hein, S. M. Ali Eslami, Danilo Jimenez Rezende, and Olaf Ronneberger. A probabilistic u-net for segmentation of ambiguous images, 2019. URL https://arxiv.org/abs/1806.05034.

Mark S Robinson, SM Brylow, M ea Tschimmel, D Humm, SJ Lawrence, PC Thomas, Brett W Denevi, E Bowman-Cisneros, J Zerr, MA Ravine, et al. Lunar reconnaissance orbiter camera (lroc) instrument overview. *Space science reviews*, 150(1):81–124, 2010.

David E. Smith, Maria T. Zuber, Gregory A. Neumann, Frank G. Lemoine, Erwan Mazarico, Mark H. Torrence, Jan F. McGarry, David D. Rowlands, James W. Head III, Thomas H. Duxbury, Oded Aharonson, Paul G. Lucey, Mark S. Robinson, Olivier S. Barnouin, John F. Cavanaugh, Xiaoli Sun, Peter Liiva, Dan-dan Mao, James C. Smith, and Arlin E. Bartels. Initial observations from the lunar orbiter laser altimeter (lola). *Geophysical Research Letters*, 37(18), 2010. doi: https://doi.org/10.1029/2010GL043751. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2010GL043751.

Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers, 2021. URL https://arxiv.org/abs/2105.15203.



(a) Image data (b) Dem data (c) Slope data (d) Hillshade data (e) Mask-label

Figure 7: Illustration of a Multisource Lunar Dataset and Multiterrain Segmentation Labels.

A DATA

We construct a terrain-understanding dataset from multi-source lunar remote-sensing data to evaluate the model's ability to segment complex lunar landforms and model uncertainty. The multi-source modalities, including the DEM and optical imagery, were derived from [e.g., Lunar Reconnaissance Orbiter (LRO) LOLA and LROC WAC/NAC] observations. The global datasets were unified to a consistent spatial resolution of [e.g., 100 meters/pixel] to ensure morphological features are represented at a geologically meaningful scale. The dataset comprises complementary global lunar observation layers in a cylindrical projection, including optical imagery, a digital elevation model (DEM), a slope map, and a hillshade map. Optical imagery provides surface texture and morphological information, the DEM describes variations in terrain elevation, and the slope and hillshade maps further emphasize geomorphic structures and spatial geometric features. All data layers are spatially registered and processed using a unified pipeline to ensure consistency in location and resolution, and they are jointly provided to the model for lunar terrain semantic segmentation, A set of sample data is shown in Figure 7. The observational data and expert-defined geological vector labels were provided by our institutional collaborators.

To construct the deep-learning dataset, we use label-mask-guided patch sampling. Specifically, we first merge all annotated terrain regions and transform the vector annotations into a local Azimuthal Equidistant (AEQD) coordinate system; we then calculate the mask coverage at candidate locations according to the ground footprint of each patch, select as the patch center the location that covers the largest amount of as-yet-unsampled terrain, and iteratively update the covered area until the remaining valid annotations are insufficient to form another patch. This procedure produces a set of 1024×1024 pixel patches, after which the annotations within each patch are rasterized to generate masks. The dataset contains 14 representative lunar landform classes; because a single lunar region may contain multiple landforms simultaneously—for example, a central peak may lie within an impact crater—we encode the 14 classes as separate binary segmentation channels, producing a 14-channel binary mask in which each channel corresponds to one landform class.

Because lunar landform types differ substantially in their spatial distributions and formation mechanisms, the dataset exhibits pronounced class imbalance. Table 3 reports the occurrence of each class in the dataset samples. Because a sample may contain multiple landform classes, the occurrence rate of each class is calculated independently, and the rates therefore do not sum to 100%. The dataset is characterized by a pronounced long-tailed distribution, large variations in landform scale, and ambiguous boundaries. Large, frequent landforms such as the "Impact crater" class occur relatively often, whereas some fracture-related, tectonic, and specialized landform classes appear in only a small number of regions. These properties make the dataset suitable not only for evaluating segmentation accuracy, but also for studying dynamic reasoning and uncertainty estimation in complex scientific imagery.

Our multi-source remote sensing observations usually suffer from significant distribution discrepancies caused by different sensors, physical meanings, and acquisition conditions. The raw data may contain inconsistent value ranges and abnormal responses, making direct fusion challenging. Therefore, we adopted a percentile-based robust normalization strategy. For each modality, extreme values were first removed by clipping the data within the 2nd–98th percentile range, followed by normalization to obtain scale-consistent inputs. This preprocessing procedure improves numerical stability during training and enhances the generalization capability across heterogeneous remote sensing datasets.

Table 3: Distribution of 14 imbalanced geological feature classes across the dataset. The Class Proportion column represents the percentage of samples containing each geological feature relative to the total number of samples. Since a single sample typically contains multiple classes, the sum of class proportions exceeds 100%.

Structure	Class Proportion(%)	Structure	Class Proportion(%)
Impact crater	96.48	Graben	18.60
Volcanic crater	17.10	Crater floor fracture	9.44
Lunar dome	12.83	Rille	19.53
Lobate scarp	16.09	Crater chain	16.25
Shallow fracture	21.85	Wrinkle ridge	8.91
Impact fracture	11.58	Central peak	23.14
Impact ejecta	13.53	Degraded crater	13.55

B Cold-Start Labels Based on Predictive Entropy

The primary difficulty for the Expert Uncertainty Branch is the lack of manual uncertainty annotations at the beginning of training. We therefore use the trained CTM segmentation model to generate a machine uncertainty prior as cold-start supervision for the uncertainty decoder. Cognitive-uncertainty data produced by experts during later use are then used to continue adapting the model, progressively bringing its output closer to expert cognition.

For the four-channel remote-sensing input, the segmentation model outputs independent sigmoid probabilities for 14 potentially overlapping landform classes.Because the task is a multi-label segmentation problem, we use a normalized Bernoulli entropy for each class:

$$h_c(x,y) = -\frac{p_c(x,y)\log p_c(x,y) + (1-p_c(x,y))\log\left(1-p_c(x,y)\right)}{\log 2}$$

Many of the 14 classes have very low responses at the same pixel location. Directly averaging the entropy across all classes would therefore dilute genuine local uncertainty. We consequently aggregate the highest-responding class entropies:

$$H(x,y) = \frac{1}{K}\sum_{c\in \mathrm{TopK}(x,y)} h_c(x,y)$$

This top-K aggregation is appropriate for the multilabel setting in which lunar landforms may overlap.

In addition to the predictive entropy used for the final deployed output, we use CTM's multistep reasoning process to construct temporal disagreement.Specifically, the frozen baseline output and the predictions from the final K_t = 6 CTM ticks are treated as multiple predictive states. Let $p_{m,c}$ denote the probability of class c in state m; the mean probability is:

$$\bar{p}_c(x,y) = \frac{1}{M}\sum_{m=1}^{M} p_{m,c}(x,y)$$

The class-level CTM disagreement is defined as:

$$d_c(x,y) = h_c(\bar{p}_c(x,y)) - \frac{1}{M}\sum_{m=1}^{M} h_c(p_{m,c}(x,y))$$

This quantity measures the disagreement between the baseline and the multiple CTM reasoning states. We apply the same top-K class aggregation to the disagreement and obtain:

$$D(x,y) = \frac{1}{K}\sum_{c\in \mathrm{TopK}(x,y)} d_c(x,y).$$

Finally, cold-start uncertainty supervision is obtained by weighting predictive entropy and CTM disagreement.

The resulting distribution is a class-agnostic, single-channel uncertainty map that represents the overall decision uncertainty at each pixel without distinguishing specific landform classes. Importantly, we do not regard the fitted entropy map as ground-truth expert uncertainty. Its role is to provide an initial machine uncertainty prior with a reasonable spatial structure, allowing the uncertainty decoder to be initialized stably before sufficient expert data become available.

C Model Training Strategy and Loss Function

C.1 CTM SegFormer Loss Function and Model Training

During one forward pass, CTM generates multiple consecutive thought ticks, each corresponding to an independent segmentation prediction. Every prediction is supervised with a combination of BCE loss and Dice loss: BCE constrains pixel-level classification for each class, while Dice loss further alleviates the imbalance in foreground-pixel counts among landform classes. The segmentation loss for thought tick t is written as:

$$L_t = L_{\mathrm{BCE}}^t + \lambda_D L_{\mathrm{Dice}}^t$$

Because different samples do not necessarily require the same reasoning depth, we do not require all samples to reach their best state at a fixed intermediate tick. Instead, we search across multiple thought ticks for the prediction with the lowest segmentation loss for the current sample. To prevent excessive reliance on a particular intermediate state and to ensure stable output after the complete continuous-thought process, we also impose supervision on the final thought tick. The main CTM segmentation objective therefore combines the best intermediate prediction with the final prediction:

$$L_{seg} = 0.5L_{best} + 0.5L_{final}$$

This training scheme allows different samples to form their best decisions at different reasoning depths while constraining the final thought tick to retain effective segmentation capability, thereby preventing substantial degradation late in the continuous reasoning process.

The Quality Head is trained with Smooth L1 loss and is gradient-isolated from the main segmentation path. Its training therefore does not modify SegFormer, CTM, or the residual prediction module. It constructs the ground-truth gain from the difference between the current tick and the baseline in dataset-level IoU and learns to predict this gain. Its objective can be written as:

$$L_Q = \mathrm{SmoothL1}\Big(\hat{G}_t, G_t\Big)$$

The module is used primarily to determine whether a thought tick genuinely outperforms the baseline and to support result selection during inference.

The final training objective of the CTM residual branch consists of the multistep segmentation loss and the auxiliary quality-prediction loss. The segmentation loss optimizes the actual landform segmentation, while the Quality Head independently learns the prediction quality of different thought states.

C.2 uncertain Branch Training

The uncertainty branch is trained by progressively transferring from a machine prior to expert cognition.

Because large-scale manual uncertainty annotations are unavailable at the beginning of training, we first use the predictive entropy generated by the main segmentation model as cold-start pseudo-labels to pretrain the Uncertainty Decoder.During this stage, the SegFormer backbone, CTM, and main segmentation branch remain frozen, and only the uncertainty-decoder parameters are updated, ensuring that uncertainty learning does not propagate backward into the original segmentation capability.

The goal of the cold-start stage is for the Uncertainty Decoder to learn the spatial uncertainty distribution already formed by the model. For the predicted uncertainty distribution and the fitted pseudo-label, we mainly enforce consistency with a pixel-level regression objective. Because the uncertainty map is a continuous spatial distribution in [0, 1], we use three losses: pixelwise regression, relative ranking, and spatial smoothness. First, BCEWithLogitsLoss constrains pixelwise

agreement between the uncertainty logits and the target uncertainty map. Second, a pairwise ranking loss constrains the relative uncertainty magnitude between pixels. Finally, a spatial-gradient loss constrains the local variation trend between the prediction and target maps. The overall loss is:

$$L_{\text{uct}} = L_{\text{pixel}} + 0.2L_{\text{rank}} + 0.05L_{\text{grad}}$$

At this stage, the pseudo-label is not treated as expert cognitive uncertainty in the final sense; it serves only as a machine prior that provides a reasonable initialization for the uncertainty branch when expert data are scarce.

As expert-interaction data accumulate, training enters the expert-calibration stage. Experts modify the predicted uncertainty distribution through two simple types of actions: increasing or decreasing the current uncertainty intensity within a selected region, or directly assigning a region a high, medium, or low uncertainty level. These operations produce expert-corrected labels and record the regions actually reviewed or modified by experts. During calibration, the loss follows the principle that manual supervision has priority and machine pseudo-labels provide auxiliary supervision. Regions confirmed or modified by experts use manual uncertainty labels as the primary supervision, whereas regions not yet reviewed by experts may retain a weaker pseudo-label constraint. The overall objective can be summarized as:

$$L_U = \lambda_h L_{\text{expert}} + \lambda_p L_{\text{pseudo}}.$$

Manually annotated regions have higher training priority, and the influence of pseudo-labels is gradually reduced as expert data increase.

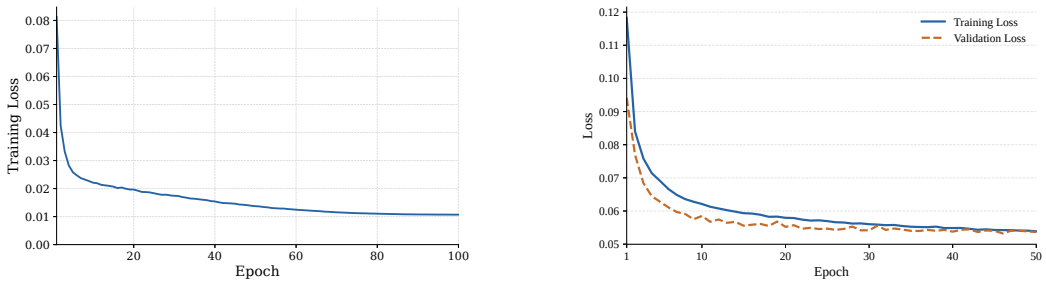
Through this strategy, the uncertainty branch first obtains a stable machine uncertainty prior from the main segmentation model and then continuously calibrates its output using expert behavior. Its predictions thereby shift from the model's own confidence distribution toward a cognitive-uncertainty representation that better matches expert interpretation. Throughout the process, the main segmentation network remains frozen, ensuring that expert feedback changes only the uncertainty branch and does not affect the CTM segmentation model itself.

C.3 Model Training Setting

Our experiments are conducted on multi-source lunar remote-sensing data. The inputs consist of the original imagery, a digital elevation model (DEM), a slope map, and a hillshade map, which are concatenated along the channel dimension before being fed into the network. The segmentation task includes 14 classes of lunar landform structures. The segmentation result for each class is represented by an independent binary channel. The dataset is divided into training and validation sets containing 6,150 and 1,085 samples, respectively. All input data are uniformly processed to a spatial size of 1024 × 1024 pixels.

During training of the CTM segmentation model, we modify the first input layer of SegFormer to accommodate four-channel inputs. In each forward pass, CTM is unrolled for a fixed number of $T = 16$ thought ticks. The training loss jointly supervises the best-performing thought tick and the final thought tick, thereby balancing the effectiveness of intermediate states with the stability of the final state. A stage-wise learning rate adjustment strategy was adopted during training. In the first two epochs, a linear warmup strategy was applied to gradually increase the learning rates of the newly introduced module parameters from a small value to the target learning rate. Subsequently, the cosine annealing strategy was employed for learning rate decay, with the minimum learning rate maintained at 0.01 times the initial learning rate. The initial learning rates of the CTM module, Residual Head, and Quality Head were all set to 1×10^{-4}. In addition, gradient clipping with a maximum gradient norm of 20 was applied to improve training stability. The model is trained using four heterogeneous accelerator cards, each with 64 GB of memory. The global batch size is set to 8, and the model is trained for 100 epochs, requiring approximately 52 hours. The loss curves during training are shown in Figure 8a.

For the Expert Uncertainty Branch, only the uncertainty decoder is updated during training. During the cold-start stage, the uncertainty maps fitted from the model's inference process are used as pseudo-labels, allowing the uncertainty branch to first learn the spatial prior of model uncertainty. This initial training uses the same accelerator configuration as the segmentation model, with a batch size of 8 and a total of 50 training epochs. The loss curves during training are shown in Figure 8b.



(a) CTM Segmentation

(b) Uncertain Branch

Figure 8: Training loss curves for the CTM segmentation model and the Uncertainty Branch during the cold-start stage.

D More Experimental Visualization Results

D.1 Comparison of CTM-based Segmentation and the SegFormer Baseline

We present several representative groups of results in Figure 9 for qualitative analysis.

Compared with the SegFormer baseline, the CTM-based segmentation model produces more refined predictions in several challenging terrain regions. In particular, the continuous thought process enables the model to adjust local segmentation boundaries, with more noticeable improvements observed for small-scale and elongated structures(figure 10h). For densely distributed terrain patterns, CTM also shows better capability in distinguishing adjacent structures and recovering some regions that are missed by the baseline prediction(figure 9l). Meanwhile, in certain cases, the CTM prediction can suppress incorrectly activated regions from the baseline model, leading to more consistent spatial delineation(figure 9p). These observations suggest that the proposed CTM mechanism can provide complementary information beyond the original segmentation representation, especially for complex terrain structures with ambiguous boundaries and fine-scale morphology.

Similarly, we provide an additional visualization of the continuous reasoning process for an impact crater and its internal structures in Figure 10, further supporting the observations presented in Section 4.2. Consistent with our previous findings, early reasoning steps primarily delineate the large-scale impact crater and prominent lobate scarps. As reasoning progresses, nested structures within the crater, such as the central peak, are gradually identified and refined. Meanwhile, the Quality Head also tends to select predictions that more closely match the ground-truth mask.

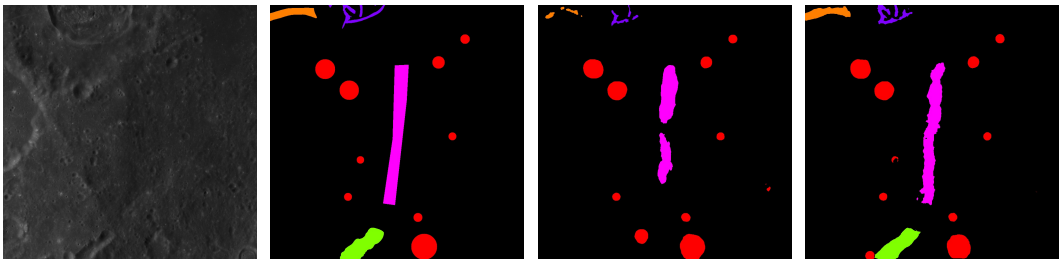
D.2 Visualization of the Uncertainty Branch during Cold-Start Training

The visualization comparison of the uncertainty prediction results during the cold-start stage of the Expert Uncertainty Branch is shown in Figure 11.

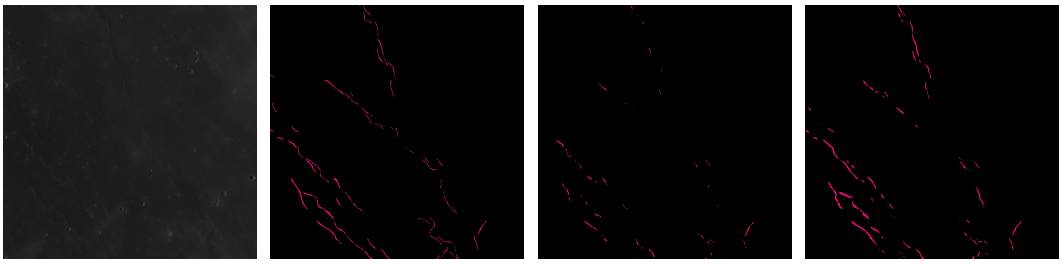
The comparison results demonstrate that the uncertainty branch can effectively approximate the initial uncertainty distribution during the cold-start stage. Nevertheless, this stage only serves as an initialization process rather than the final goal of uncertainty modeling. It enables the uncertainty prediction head to learn the mapping from backbone multi-scale features to spatial uncertainty representations based on model-derived priors. With subsequent expert feedback and the collection of sufficient expert-adjusted uncertainty annotations, the uncertainty branch can be further calibrated to gradually bridge the gap between machine uncertainty and expert-oriented cognitive uncertainty.

E Uncertainty Metrics

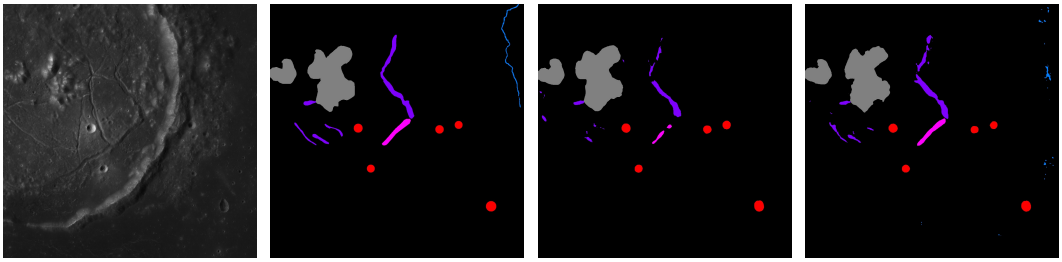
To quantitatively evaluate the uncertainty prediction results, we adopt three complementary metrics: Mean Absolute Error (MAE), Spearman correlation coefficient, and Pearson correlation coefficient. Given a predicted uncertainty map $\hat{U}$ and the target uncertainty map U, both maps are flattened into $N = H \times W$ pixel-wise values for calculation.



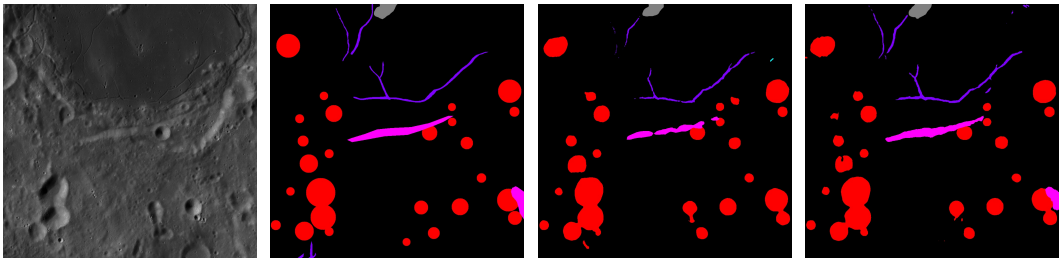
(a) Image (b) GT mask (c) SegFormer results (d) CTM results



(e) Image (f) GT mask (g) SegFormer results (h) CTM results



(i) Image (j) GT mask (k) SegFormer results (l) CTM results



(m) Image (n) GT mask (o) SegFormer results (p) CTM results

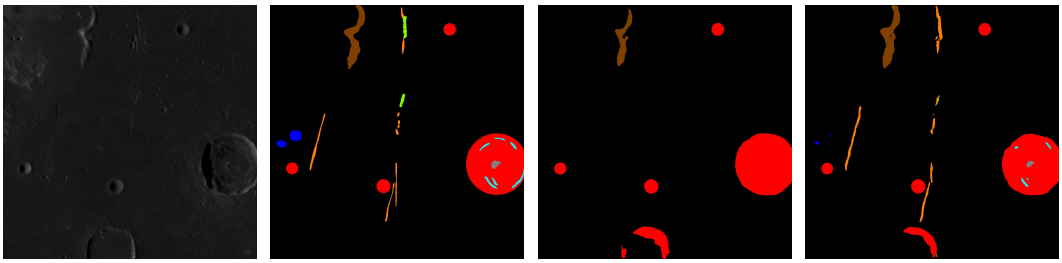
Figure 9: Supplementary results. Each row shows one comparison group. From left to right, the four columns represent the lunar image, Ground Truth Mask, SegFormer baseline prediction, and the final prediction of our CTM-based segmentation model, respectively.

MAE measures the average absolute difference between the predicted uncertainty intensity and the target uncertainty value:

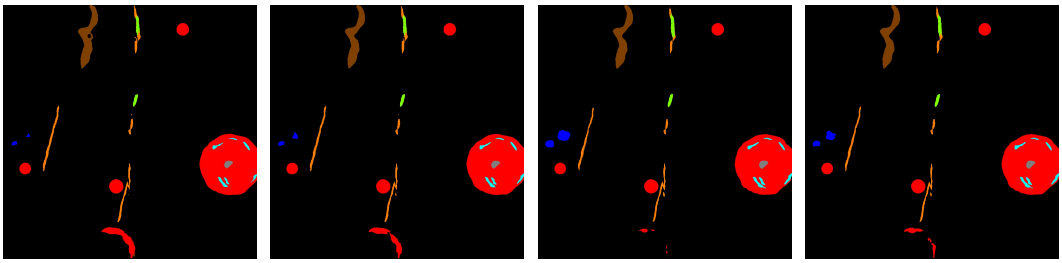
$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N} |\hat{U}_i - U_i|.$$

A lower MAE indicates that the predicted uncertainty values are closer to the target uncertainty distribution.

Spearman correlation evaluates whether the predicted uncertainty preserves the relative ordering of uncertainty among different spatial locations. Instead of directly comparing numerical values, it

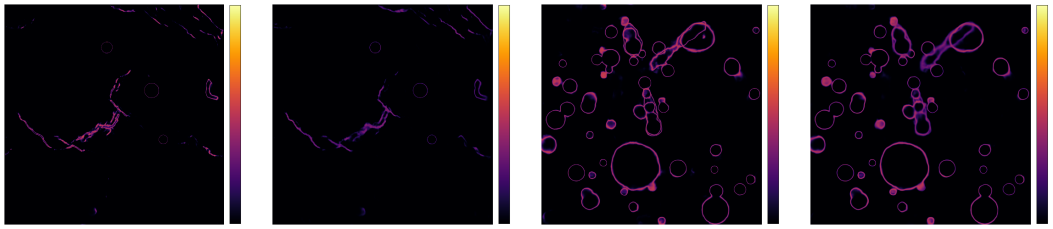


(a) Image (b) GT mask (c) t1 (d) t



(e) t4 (f) t6 (g) t8-selected (h) t16-final

Figure 10: Additional visualization of a representative continuous reasoning process, illustrating the progressive refinement and increasingly clear delineation of geological structures across thought steps.



(a) GT-1 (b) Predict-1 (c) GT-2 (d) Predict-2

Figure 11: Visualization evaluation of the uncertainty branch during the cold-start stage. (a) and (b) show the Ground Truth and predicted uncertainty maps of the first example, while (c) and (d) show those of the second example.

first converts the predicted and target uncertainty values into ranked sequences and then calculates the correlation between these rankings. Therefore, it reflects the consistency of spatial uncertainty trends.

Pearson correlation measures the linear correlation between predicted and target uncertainty values:

$$r = \frac{\sum_i (\hat{U}_i - \overline{\hat{U}})(U_i - \overline{U})}{\sqrt{\sum_i (\hat{U}_i - \overline{\hat{U}})^2} \sqrt{\sum_i (U_i - \overline{U})^2}}.$$

where $\hat{U}_i$ and U_i represent the predicted and target uncertainty values at pixel i, respectively. $\overline{\hat{U}}$ and $\overline{U}$ denote the corresponding mean uncertainty values. The Pearson coefficient ranges from -1 to 1, with larger positive values indicating stronger linear consistency between the predicted and target uncertainty distributions.

Together, these three metrics provide complementary evaluations of uncertainty prediction: MAE measures absolute numerical accuracy, Spearman evaluates preservation of relative uncertainty ordering, and Pearson measures global linear consistency.