

Foundation Model Selection for Remote Sensing via a Constraint-Aware Agent

Binger Chen¹^{*}, Tacettin Emre Bök^{2*}, Behnood Rasti³, Volker Markl³, and Begüm Demir³

¹ Humboldt-Universität zu Berlin, Berlin, Germany

`binger.chen@hu-berlin.de`

² 5U AI, Munich, Germany

`tacettinemre1@gmail.com`

³ Technische Universität Berlin & BIFOLD, Berlin, Germany

`behnood.rasti@gmail.com, {volker.markl,demir}@tu-berlin.de`

Abstract. Foundation Models (FMs) are increasingly integrated into remote sensing (RS) pipelines for applications such as environmental monitoring, disaster assessment, and land-use mapping. These models include unimodal vision encoders trained in a single data modality and multimodal architectures trained in multiple sensor modalities, such as synthetic aperture radar (SAR), multispectral, and hyperspectral imagery, or jointly in image-text pairs in vision-language settings. FMs are adapted to diverse perception tasks, such as semantic segmentation, image classification, change detection, and visual question answering, depending on their pretraining objectives and architectural design. However, selecting the most suitable remote sensing foundation model (RSFM) for a specific task remains challenging due to scattered documentation, heterogeneous formats, and complex deployment constraints. To address this, we first introduce the RSFM Database (RS-FMD), the first structured and schema-guided resource covering over 160 RSFMs trained on various data modalities, spanning different spatial, spectral, and temporal resolutions, considering different learning paradigms. Built upon RS-FMD, we further present **REMSA** (**R**emote-sensing **M**odel **S**election **A**gent), a constraint-aware agent that enables automated RSFM selection from natural language queries. **REMSA** combines structured FM metadata grounding with task-aware orchestration for retrieval, clarification, ranking, and explanation. In detail, it interprets user input, clarifies missing constraints, ranks models via in-context learning, and provides transparent justifications. Our system supports various RS tasks and data modalities, enabling personalized, reproducible, and efficient FM selection. To evaluate **REMSA**, we construct a benchmark of 100 expert-verified RS query scenarios. Each query is evaluated across 4 systems and 3 LLM backbones, with the top-3 selected models assessed by domain experts using a fixed suitability rubric. This results in **3,000** expert-scored task-system-model configurations under our rubric-based expert evaluation protocol. **REMSA** outperforms multiple baselines, including naive agent-based method, dense retrieval, and unstructured retrieval augmented generation approaches, showing its practical utility in real decision-making applications. **REMSA** operates entirely on publicly available metadata of open source RSFMs, without accessing private or sensitive data. Our code and data are publicly available at: <https://github.com/be-chen/REMSA>.

Keywords: Remote Sensing · Foundation Models · Model Selection · Foundation Model Database · Agents

^{*} Work done while the author was affiliated with Technische Universität Berlin & BIFOLD.

1 Introduction

With the growing availability of remote sensing (RS) missions and their onboard sensors (*e.g.*, Sentinel-2 [7], Sentinel-1 [34], EnMAP [10]), RS plays an increasingly important role in many applications such as agriculture, disaster response, urban development, and biodiversity monitoring. These applications increasingly rely on foundation models (FMs) that can generalize across various RS data modalities with different spatial, spectral, and temporal resolutions, geospatial extents and applications, while being transferable and effective even with limited labeled data in downstream RS analysis tasks. Recently, numerous FMs have emerged in the RS domain, offering powerful capabilities for modeling complex RS data. These models include vision-only encoders trained on single or multiple RS data modalities (*e.g.*, MMEarth [22], OmniSat [2], MA3E [17]) and vision-language models (VLMs) trained with paired RS data modalities and text (*e.g.*, GeoText [4], LHRs-Bot [21]). These models are pretrained on large-scale and heterogeneous RS datasets encompassing a diverse range of sensor modalities, including RGB, multispectral, hyperspectral, synthetic aperture radar (SAR), and light detection and ranging (LiDAR), across multiple spatial and temporal resolutions. Each FM exhibits its strengths in distinct downstream RS tasks, such as classification, object detection, change detection, captioning, and visual question answering (VQA). For instance, in practical RS workflows, change detection typically relies on multitemporal SAR or optical data inputs, while fine-grained land cover mapping often benefits from high-resolution optical imagery. This diversity brings new possibilities for multi-modal RS applications, but it also raises the challenge of selecting the most suitable FM for a given task with data modality and operational constraints.

Despite these advances, selecting a suitable FM for a specific RS task remains a challenge. RS practitioners must balance diverse constraints such as the available data modalities and volume, geographic coverage, computational resources, and task-specific evaluation priorities. These constraints have been shown to significantly influence RSFM generalization and robustness [25, 26]. With hundreds of remote sensing foundation models (RSFMs) now publicly available [11, 16] and no unified structured schema to organize their properties (such as model architectures, training data, or reported performance), the selection process is often manual, which is time-consuming, error-prone, and difficult to reproduce. Existing approaches rely on searching across scattered repositories and publications, manually parsing papers and model cards, and running exhaustive experiments [1, 27], all without guaranteed reproducibility or transparency. Even public RS benchmarks [11, 13, 30] mainly compare model accuracy on fixed applications, but they do not address the upstream selection problem of identifying suitable candidate FMs under user-specific modalities, data availability, compute budgets, and deployment trade-offs. This makes a unified, machine-readable database (DB) of RSFMs a necessary foundation for any systematic selection and automation.

Recent advances in large language model (LLM) agents have shown the potential of integrating language models with external tools and structured workflows to support complex tasks [8, 12, 18, 35]. However, most existing agents target general-purpose question answering or conversational assistance. To the best of our knowledge, no prior work has developed a domain-specific agent for FM selection in operational, constraint-heavy RS scenarios. In particular, RS tasks involve complex trade-offs across sensors, spatial, spectral, and temporal resolutions, as well as data availability. Existing LLMs lack the domain knowledge and structured access to model documentation to address these constraints. Hence, such an agent must provide more robust and interpretable solutions.

In this work, we first introduce the RSFM Database (RS-FMD), the first schema-guided resource of over 160 RSFMs, covering various data modalities, pretraining strategies, and reported




Fig. 1: Architecture of **REMSA** for remote sensing foundation model selection. **REMSA** parses an unstructured user query, grounds the extracted constraints in RS-FMD and task memory, and orchestrates retrieval, ranking, clarification, and explanation tools to produce justified model recommendations.

benchmark results. Built on RS-FMD, we propose **REMSA**, a first LLM-based agent for automated FM selection in RS. As shown in Fig. 1, **REMSA** is a modular agent that supports FM selection through structured query interpretation and task-aware tool invocation. It extracts user intent from free text input and converts it into constraints. Based on the task state, the agent orchestrates tools for retrieval from RS-FMD, ranking FMs, interactive clarification, and justified explanation. A lightweight memory mechanism further enhances accuracy and personalization. To evaluate **REMSA**, we build a new benchmark of realistic user queries and establish a rubric-based expert evaluation protocol that scores the suitability of selected FMs under user-specified constraints. We also implement a set of carefully constructed baselines, ensuring fair and meaningful comparisons with **REMSA**. **REMSA** is designed to support a broad range of users, including RS scientists, machine learning researchers, and industry practitioners who need to identify suitable RSFMs for their tasks. Because **REMSA** accepts free-text queries and incorporates structured interpretation together with multi-turn clarification, it can assist even non-experts who may not be familiar with RS modalities or FM architectures. This makes **REMSA** suitable for both exploratory use by practitioners and rigorous FM selection in research settings. Although **REMSA** adopts a modular agent design rather than proposing a new visual backbone, our contribution is methodological. We formalize RSFM selection as an under-supported stage in practical RS or computer vision pipelines: deciding how FMs should be compared, selected, and deployed under modality, data, compute, and application constraints before downstream adaptation. In summary, we make the following **contributions**:

- We introduce RS-FMD, the first structured and schema-guided DB of over **160** RSFMs. We will release it as a community resource with continuous maintenance and updates.
- We propose **REMSA**, a modular LLM agent that combines structured metadata grounding, dense retrieval, in-context ranking, clarification, explanation, memory augmentation components, and a task-aware orchestration mechanism to support complex FM selection in real RS settings.
- We construct a new benchmark dataset and rubric-based evaluation protocol for FM selection in RS, comprising 100 realistic query scenarios and **3,000** expert-scored evaluations across systems, LLM backbones, and selected models.

2 Related Work

Foundation Models and Model Selection. Due to the rapid emergence of FMs, there has been extensive research into their capabilities and benchmarks [3, 6, 36]. In RS, recent surveys and benchmarks [11, 27, 37] have systematically cataloged FMs and evaluated their performance on applications such as land cover classification, wildfire scar segmentation, urban change detection, visual question answering, etc. However, these works primarily focus on descriptive analysis or standardized evaluation, offering limited support for automated FM selection. Large-scale evaluations such as GEO-Bench-2 [30] further highlight that RSFM performance varies strongly across capability dimensions, but they mainly evaluate models under fixed benchmark protocols rather than addressing the upstream problem of selecting suitable FMs under user-specific constraints. Other work also shows that pretraining data coverage (geographic and sensor diversity) strongly affects RSFM generalization [25, 26]. While current benchmarks document these properties, they do not use them to guide model choice, further motivating automated FM selection. Additionally, recent capabilities encoding approach estimates a model’s performance on unseen downstream tasks, reducing the need for exhaustive fine-tuning [1]. Although such performance prediction methods provide valuable tools for comparative evaluation, they are complementary to our setting. They rely on standardized transfer evidence and do not address end-to-end FM selection from open-ended user queries involving modalities, data availability, compute budgets, and deployment constraints. Moreover, previous surveys and benchmarks are static and task-specific, lacking a unified schema or machine-readable representation of RSFMs. In contrast, our RS-FMD consolidates the available FMs into a structured, extensible resource that directly supports automated retrieval, comparison, and selection. Another relevant line of work is AutoML, which includes frameworks such as Auto-WEKA [33], Auto-sklearn [9], and CAML [23]. They automate the selection of parameters, algorithms, or pipelines through meta-learning and optimization techniques. Although these approaches show the feasibility of automating model choice in classical machine learning settings, they have not been applied to the selection of FMs, particularly in the RS domain. To our knowledge, there is no existing dedicated method or agent that assists scientists in selecting the most suitable FM for their specific constraints and applications. Our work fills this gap by combining RS-FMD and **REMSA**, formalizing constraint-aware RSFM selection as an upstream step in practical RS or computer vision pipelines and implementing it through structured metadata grounding, task-aware orchestration, clarification, ranking, and explanation.

Tool-Augmented Agents in RS. Recent developments in retrieval-augmented language models and tool-augmented agents such as LLaVA-Plus [18], OmniACT [12], VideoAgent [8] show the feasibility of combining LLMs with structured retrieval and external tool invocation for complex task execution and workflows. In RS, several works have explored modular agentic workflows. GeoLLM-Squad [15] introduces a multi-agent orchestration framework that decomposes geospatial tasks into specialized sub-agents, improving scalability and correctness over single-agent baselines. RS-Agent [38] integrates retrieval pipelines and tool scheduling to process spatial question answering tasks, while ThinkGeo [29] introduces a benchmark for evaluating multi-step tool-augmented agents on RS workflows. Recently, EarthDial [31] extends large vision-language models to multi-sensor earth observation data through interactive dialogue modeling, enabling conversational reasoning over complex data modalities. These agentic approaches highlight the benefits of modularity and retrieval-augmented decision process. However, they mainly target geospatial information extraction, change detection, or VQA applications rather than FM selection. **REMSA** explicitly integrates a curated FM database with structured retrieval, agentic ranking, interactive constraint resolution,

and transparent model reasoning, making it a dedicated tool-augmented agent tailored for FM selection in RS.

3 Remote Sensing Foundation Model Database

RS-FMD is a curated DB of RSFMs identified through a systematic search covering over 160 RSFMs, serving as the structured resource behind **REMSA**. It enables interpretable and constraint-aware FM selection by consolidating heterogeneous resources into a unified, machine-readable format. To build RS-FMD, we conducted a systematic search for RSFMs using multiple sources. We reviewed survey papers and popular FM lists, surveyed recent RS and ML venues, ran keyword searches on arXiv, and inspected linked GitHub repositories.

Schema Design. Each record follows a schema covering properties such as identifiers, architecture, modalities, and pretrained model weights, along with structured fields for pretraining datasets and reported benchmark evaluations when available. This schema ensures traceability, comparability, and extensibility across FMs. The full schema and an example record are in Appendix (App.) Sec. A. This comprehensive schema enables our FM selection agent to ground the selection process in documented model properties and reported capabilities, and match models to user-defined tasks and constraints. It also ensures that critical properties, such as input data modalities, spatial, spectral, and temporal characteristics, and training configurations, can be queried and filtered in a principled and automated manner.

Automated database population. Populating this database requires extracting structured information from diverse sources, such as papers, model cards, and repositories. Due to the scale and heterogeneity of available model documentation, fully manual curation is impractical. Therefore, we adopt a semi-automated knowledge extraction approach coupled with confidence-guided human verification. Our approach is a schema-guided LLM extraction pipeline inspired by a general knowledge extraction approach OneKE [19], but significantly adapted to our domain and use case. Specifically, we extend their approach by introducing our own schema definitions, adding a dedicated confidence scoring step, and optimizing prompt design for RS model descriptions. The process is entirely automated and iterative: for each FM, we collect and input a set of unstructured sources, then issue multiple LLM calls to generate independent structured outputs in each iteration. Each output is validated against the schema, parsed, and aggregated. This iterative strategy allows us to exploit both the probabilistic uncertainty of each iteration and the self-consistency across iterations. Fields for which the model produces divergent outputs or low log-probabilities are marked as uncertain and passed to the human verification stage. The resulting pipeline effectively converts complex heterogeneous text sources into machine-readable JSON records with minimal targeted manual intervention.

Confidence Score for Human Verification. Ensuring the reliability of the extracted metadata is critical for FM selection in downstream RS tasks. To this end, we define a confidence score for each field in each record, enabling targeted human verification only where the uncertainty is high. Our confidence score combines two complementary criteria: the model’s generation probability and the consistency of outputs across multiple LLM sampling rounds. For each field, we compute the confidence score as follows:

$$\text{Confidence} = w_{\log p} \cdot \text{NormalizedLogProb} + w_{\text{cons}} \cdot \text{SelfConsistency} \quad (1)$$

where **NormalizedLogProb** quantifies the LLM’s internal certainty by mapping the raw log-probability of the generated field value to a bounded range, and **SelfConsistency** measures the

fraction of LLM generations that agreed on the same value among multiple independent sampling iterations.

To ensure interpretability and numerical scaling, we normalize log-probabilities using a sigmoid function controlled by the temperature parameter. We set the temperature $\tau = 0.5$ to avoid saturation and preserve sensitivity in the moderate-confidence regime of the sigmoid function. We set $w_{\text{logp}} = 0.7$ and $w_{\text{cons}} = 0.3$ to prioritize the log-probability signal while still leveraging the stabilizing effect of self-consistency. These weights were empirically determined via a grid search on a validation set of randomly sampled 10 FM records with manually verified ground truth. We optimized for maximum agreement between the confidence score and human verification decisions, using the area under the precision-recall curve (PR-AUC) as the selection criterion. We observed that prioritizing the log-probability signal improved precision, while incorporating self-consistency helped identify low-confidence outliers. However, these weights are not necessarily fixed and can be adjusted by practitioners depending on the properties of their LLMs, model domains, or confidence calibration needs. Any field with a final confidence below a threshold $\theta = 0.75$ is automatically flagged for human review. Importantly, human annotators inspect only the flagged fields rather than full model records. Reviewing all fields for all FMs would require substantial annotation effort, as each record contains many heterogeneous metadata elements. To assess the risk of confidently incorrect extractions, we manually inspected all fields for 10 records and found high-confidence outputs to be consistently accurate, supporting the reliability of our scoring mechanism. In practice, occasional field-level errors have limited impact on FM selection, as the most decisive documented properties (modality, architecture, compute requirements, and reported performance evidence) are usually clearly stated and rarely mis-extracted.

Diversity of Coverage. The current release of RS-FMD spans a broad range of RSFMs pretrained on various data modalities (multispectral, hyperspectral, SAR, LiDAR, and text) and employing diverse model architectures (transformer-based encoders, CNN–transformer hybrids, vision–language models). Pretraining data sources range from small curated datasets to million-scale image collections, and spatial resolutions span from sub-meter imagery to coarse multispectral composites. By consolidating these heterogeneous information sources into a schema-guided resource, RS-FMD supports reproducible comparison, metadata-driven analysis, and automated retrieval. We will maintain RS-FMD in a public repository under a permissive license through a community-in-the-loop update process. In addition to monthly rolling checks by the maintainers for newly released RSFMs, model authors can submit documentation for new models directly. **REMSA** then automatically extracts the corresponding schema fields and presents them for inspection, while maintainers review the submitted records before inclusion to ensure consistency and reliability. This design reduces maintenance latency compared with fully manual curation, while retaining quality control over the structured metadata.

4 REMSA Agent Architecture

The goal of **REMSA** is to automate the upstream selection of FMs for RS tasks through a decision-centered, modular agentic workflow. **REMSA** integrates structured knowledge grounding, LLM-assisted ranking with in-context learning, and iterative clarification to produce transparent and reproducible selections. Selecting an appropriate RSFM is challenging, as the models differ in data modalities, pretraining strategies, reported benchmark evidence, and resource requirements. In addition, users often provide incomplete or ambiguous task descriptions, requiring the agent to infer user intent and reconcile trade-offs among candidate models. To address these challenges, **REMSA**

provides an integrated pipeline combining different agent components and external tools. This pipeline decomposes RSFM selection into structured retrieval, constraint filtering, ranking, clarification, explanation, and memory augmentation under a customized orchestration mechanism. This section describes the agent workflow and the details of each component and tool.

4.1 Agent Workflow

Fig. 1 illustrates the architecture of **REMSA** that is composed of two main layers: the **LLM agent core** and a set of **external tools**. The LLM agent core consists of two key components: the *Interpreter*, which parses user inputs and extracts user intent, and the *Task Orchestrator*, which dynamically decides which external tool to invoke at each step based on the current task state. When a user submits a free-text query, the *Interpreter* transforms it into a structured representation of constraints. We prompt the LLM with a carefully designed schema that covers both mandatory and optional fields relevant to RSFM selection (see App. Sec. B for complete schema). Specifically, the parser extracts the target **application** (*e.g.*, land cover classification, surface water segmentation) and the required **modality** (*e.g.*, multispectral, SAR) as mandatory fields to narrow the FM search space. Then **REMSA** integrates broader practical constraints through optional fields and clarification steps, including data availability, compute budget, fine-tuning requirements, and output quality priorities. Once constraints are available, the *Task Orchestrator* initiates a control loop that manages the entire selection process. At each step, it first evaluates the current task state, including constraint completeness, candidate set size, hard constraint violations, and ranking confidence. Then it invokes the appropriate tool accordingly. If no mandatory constraints are missing, the orchestrator calls the *Retrieval Tool* to generate an initial candidate set. If the candidate set is small and all constraints are satisfied, the *Ranking Tool* is applied directly. If there are too many candidates or if ranking results yield low confidence scores, the orchestrator calls the *Clarification Generator Tool* to ask the user for additional input. The updated query is then passed back through the same loop. Once the top- k result is obtained, the *Explanation Generator Tool* is invoked to produce the final report. This decision process is governed by predefined thresholds for ranking confidence, constraint coverage, candidate set size, and clarification rounds. The orchestration makes the selection process adaptive and transparent, rather than a single-step retrieval or unconstrained LLM recommendation. To support personalization and long-term refinement, **REMSA** also integrates *Task Memory*, which stores past user interactions in a lightweight vector database. Relevant memory entries are retrieved via cosine similarity to improve future interactions. More details on the implemented workflow are in App. Sec. C. To enhance **REMSA**’s reliability, we have several built-in mechanisms to mitigate failures. The orchestrator monitors confidence signals and triggers clarification rounds when ranking is uncertain. Rule-based constraints filter candidates that violate hard requirements. A fallback “closest-match” mode returns the most compatible alternative when no candidate fully satisfies the constraints. Our modular design also allows for integrating explicit feedback mechanisms (*e.g.*, an optional LLM-as-a-Judge component that re-evaluates low-quality selections), making **REMSA** extensible to more robust error mitigation strategies.

4.2 Agent Tools

The following tools operate as callable interfaces outside of the agent core. Each tool is invoked independently by the orchestrator, depending on the task state, supporting retrieval, ranking, clarification, and explanation within the RSFM selection workflow. Our design supports extensibility for future tool integration.

Retrieval Tool. To generate an initial candidate set, the retrieval tool encodes both the structured constraints and the FM entries in the RS-FMD using Sentence-BERT embeddings [28]. To preserve the metadata structure in the embedding, each metadata field is prefixed with a type-indicator token (*e.g.*, [APPLICATION], [MODALITY]) before encoding. **REMSA** uses Facebook AI Similarity Search (FAISS) [20] for an efficient search based on cosine similarity. The tool returns a list of the most relevant FMs determined by a configurable similarity threshold. Users can adjust it based on their domain requirements. In our experiments, we set this threshold empirically based on preliminary experiments to ensure broad coverage while minimizing irrelevant matches. This tool is optimized for high recall: it includes soft matches and does not enforce strict constraints, allowing the downstream pipeline to handle finer filtering.

Ranking Tool. While the retrieval tool provides a broad list of relevant FMs, it cannot fully capture user-specific needs and deployment trade-offs. The ranking tool therefore combines deterministic constraint filtering with LLM-based preference ordering over the remaining candidates. The ranking tool refines the candidate FM list using a hybrid strategy to balance efficiency, flexibility, and interpretability:

- *Rule-Based Filtering:* Candidates that violate hard constraints, such as required modality, sensor support, or mandatory deployment requirements, are eliminated using deterministic logic. These hard constraints are defined based on fields extracted by the interpreter.
- *In-Context LLM Ranking:* The remaining candidates are re-ranked by an LLM prompted with the structured query and FM metadata, using expert-crafted few-shot examples that demonstrate how to balance application fit, modality compatibility, reported evidence, and deployment constraints. The LLM returns an ordered list with brief justifications, leveraging in-context ranking rationale without any model training (details in App. Sec. D). We also compute a field-aware confidence score for each selection following Sec. 3.

Clarification Generator Tool. If the orchestrator detects insufficient constraints or a low overall confidence score of selected FMs, it invokes the clarification tool. This tool inspects the parsed schema to determine missing or under-specified fields (*e.g.*, modality, region, compute budget, or output-quality priority) and formulates clarification questions. The tool generates each question based on the interpreter schema. We limit the clarification to three rounds to avoid user fatigue. **REMSA** will integrate the responses with initial user input, parse and merge them into the evolving task specification, in order to iteratively refine the selection process.

Explanation Generator Tool. Once a ranking is available, this tool generates structured, human-readable explanations. It uses a prompt-driven LLM to synthesize the rationale for each selected FM, including key reasons about suitability, constraint satisfaction, and trade-offs among candidates. Each JSON format output includes the model name, bullet points for explanation, and links to the corresponding paper and code repository. This tool enhances transparency and user trust by exposing the selection process (prompt is in App. Sec. E).

5 Evaluation Protocol and Benchmark

Evaluating FM selection in RS is challenging due to the lack of dedicated benchmarks. Previous works mainly focus on evaluating model accuracy after adaptation on fixed tasks or datasets, rather than assessing the upstream ability to recommend suitable FM under diverse real-world deployment constraints. In this work, we leverage RS-FMD to construct an agent-based benchmark for upstream FM selection, systematically covering diverse models, modalities, and deployment constraints.

Benchmark Construction. Our evaluation protocol relies on structured expert scoring rather than agreement with human-selected answers, ensuring methodological rigor without imposing excessive annotation overhead. We curate a benchmark of 100 diverse natural language queries to keep expert evaluation feasible while still ensuring diversity and realism. All queries can be found in our source code. Each selected model-query pair is anonymized with respect to the generating system and evaluated by two experts with complementary RS and ML expertise. Experts do not provide a separate ground-truth top-1 FM. Instead, they score the suitability of each selected model under the user-specified constraints using a structured rubric. Full details of the expert procedure are in App. Sec. G. In total, the evaluation yields 3,000 expert ratings, as experts assess the top-3 FMs selected by **REMSA** and 3 baseline systems across the query set. Each instance is carefully rated across seven criteria, resulting in a substantial expert evaluation workload despite the modest number of unique queries. To maximize representativeness, we create the queries using structured templates of various scenarios and instantiate them (Full templates are in App. Sec. H.) The queries diversify over data availability, computational resources, application complexity, and evaluation priorities. The resulting queries cover a wide range of tasks, including flood mapping with SAR data, crop type classification using multispectral or hyperspectral imagery, urban expansion monitoring with optical time series, and environmental monitoring or hazard-related applications, such as sea ice and wildfire detection. These tasks cover both single-date and multi-temporal analysis, single- and multi-modal inputs, and varying resource environments. All queries were reviewed by a domain expert for factual accuracy and corrected for consistency.

Baselines. There is limited prior work on automated FM selection for RS deployment, and existing AutoML or agent systems cannot directly perform this task. We therefore design baselines that serve as both meaningful comparisons and component analysis of **REMSA**, with each baseline removing or modifying specific components to assess their contributions:

1. **REMSA-NAIVE:** Same toolset and DB as **REMSA**, but employs only basic sequential orchestration without **REMSA**'s adaptive, task-aware control logic. It relies on LangChain's default single-step execution, where the LLM independently chooses tools without structured workflow or multi-turn coordination [14]. This baseline tests the effectiveness of our orchestration mechanism.
2. **DB-RETRIEVAL:** Returns the top- k models from the FAISS-based dense retrieval over RS-FMD, with ranking, clarification, memory, and orchestration removed. This serves as a retrieval-only baseline and isolates the contribution of LLM-based ranking and constraint reasoning.
3. **UNSTRUCTURED-RAG:** A generic RAG setup where LLM receives the query and unstructured FM descriptions and outputs top- k FMs with brief justifications (prompt in App. Sec. F). This baseline tests whether a general-purpose LLM can perform FM selection without our structured, modular agent.

Evaluation Protocol and Criteria. For each query, **REMSA** and all baselines output their top-3 selections. These model-query pairs were evaluated independently and blindly by the two experts using the 7 criteria in Tab. 1. After independent scoring, substantial disagreements were resolved according to the predefined rubric, with the corresponding model documentation re-examined when necessary. The evaluation protocol was fixed before scoring, and no candidate sets, system outputs, rubric criteria, or model records were modified after scoring to avoid post-hoc bias. Each selected FM was rated on a 1-5 scale (0.5 precision) across the 7 criteria, covering task relevance and deployment suitability under real-world constraints. Several criteria use explicit rules. For example, generalizability combines geographic, modality, and dataset-scale factors, popularity relies on citations or GitHub activity, and recency is based on publication year (More details are in App. Sec. G.).

Table 1: Expert rubric for scoring the suitability of selected query–model pairs.

Criterion	Description
Application Compatibility	Whether the model fits the user requested application
Modality Match	Whether the model supports the required input data modality
Reported Performance	Performance reported on similar datasets or applications
Efficiency	Suitability for the user’s computational resources
Popularity	Based on GitHub repository stars and citations
Generalizability	Diversity and scale of pretraining data
Recency	Whether the model reflects recent developments

The criteria are designed to be transparent, reproducible, and grounded in practical needs rather than ad-hoc user preferences. The final score is a weighted sum of all criteria ratings (Our weight settings are in App. Sec. I.). The score is linearly mapped to a 1-100 scale to better show the differences.

Exhaustive empirical benchmarking of all candidate models under all query-specific settings is infeasible, because such downstream accuracy–cost benchmarking would require adapting and evaluating many heterogeneous RSFMs for each user-specific task, modality, region, and resource setting, which is beyond the scope of this work. Instead, we develop a protocol that offers a reproducible and practical proxy for evaluating the quality of real-world FM selection workflows. To support broader community adoption, we publicly release the full set of evaluation queries, expert guidelines, scoring criteria, and model metadata used in the evaluation. This enables reproducibility and provides a standardized foundation for future research on FM selection in RS and beyond. Our evaluation does not assume a single ground-truth “best” FM. Instead, experts score each selected model-query pair according to a fixed rubric, and systems are compared based on the suitability scores of their returned candidates. Thus, the evaluation measures model selection quality under user constraints, rather than agreement with independently selected human answers or downstream task accuracy after fine-tuning. **REMSA** returns top- k FMs with explanations, enabling users to choose based on their own preferences.

6 Experiments

We conduct experiments to comprehensively evaluate the quality of **REMSA**’s RSFM recommendations under user-specified constraints. Since limited prior work directly targets real FM selection under diverse deployment constraints, we develop our own baselines. This section presents our experiment setup, quantitative results, and sensitivity analysis, followed by a discussion of limitations and representative examples.

Experimental Setup. We use GPT-4.1 [24] as the primary LLM core for **REMSA** and all baselines to ensure fairness. To evaluate robustness across different backbones, we additionally conduct experiments with DeepSeek3.2 [5] and LLaMA-3.3-70B [32]. **REMSA** is designed to be LLM-agnostic and can readily support different LLMs without architectural changes. Our benchmark consists of 100 diverse natural language user queries. For each input, **REMSA** and all baselines (all described in Sec. 5) select the top-3 candidate FMs for comparison. Domain experts rate each candidate using the criteria in Tab. 1, and we report both single-model and set-level suitability scores to evaluate overall selection quality. During evaluation, all clarification rounds in **REMSA** were executed automatically, with the system interacting with a separate LLM simulating user responses.

Table 2: Model selection quality comparison to the baselines (GPT-4.1)

System	Avg Top-1	Avg Set	Top-1 Hit	HQ Hit	MRR
REMSA (Ours)	75.76	75.03	21.33%	40.00%	0.34
REMSA -NAIVE	72.67	72.00	20.00%	37.33%	0.29
DB-RETRIEVAL	67.37	68.87	12.00%	17.33%	0.23
UNSTR.-RAG	71.23	68.39	13.33%	30.67%	0.24

No human was involved in these interactions, ensuring consistency and preventing evaluator bias across systems and LLM backbones.

Evaluation Metrics. We use complementary metrics to evaluate both the top-ranked model and the overall returned set quality: (1) *Average Top-1 Score* (expert-assigned suitability score of the top-ranked model), (2) *Average Set Score* (average expert-assigned suitability score of the top-3 models), (3) *Top-1 Hit Rate* (fraction where the system’s top-ranked model receives the highest expert score among pooled candidates), (4) *High-Quality Hit Rate* (fraction where the top model scores ≥ 80), and (5) *Mean Reciprocal Rank - MRR* (rank of the highest-scored pooled candidate within the system’s top-3).

6.1 Comparison to baselines

As shown in Tab. 2, **REMSA** consistently outperforms all baselines across the evaluated metrics, demonstrating its effectiveness for FM selection under various real constraints. Illustrative examples of expert-scored model-query pairs are provided in App. Sec. J. Under GPT-4.1, **REMSA** achieves the highest Average Top-1 Score (75.76) and Average Set Score (75.03), indicating not only that the top-ranked models receive stronger suitability scores, but also that the top-3 selections offer competitive and diverse alternatives. It also achieves the highest Top-1 Hit Rate (21.33%), High-Quality (HQ) Hit Rate (40.00%), and MRR (0.34), indicating both precision at rank 1 and stable ranking quality. Compared to DB-RETRIEVAL, which relies on similarity-based retrieval over structured metadata, **REMSA** improves Top-1 Hit Rate from 12.00% to 21.33%, HQ Hit Rate from 17.33% to 40.00%, and MRR from 0.23 to 0.34. This underscores the value of structured decision logic beyond retrieval, especially when user queries involve constraints (*e.g.*, modality, resolution, compute budget) that require the composition of multiple constraints rather than direct metadata matching. Although UNSTR.-RAG has access to full model descriptions, its performance remains lower due to the lack of structured guidance and modular reasoning. While it achieves moderate Average Top-1 scores (71.23) compared to **REMSA**, its Top-1 Hit Rate (13.33%) and MRR (0.24) remain substantially below **REMSA**. This result shows that **REMSA**’s ability to combine structured schema grounding with task-aware tool orchestration enables more precise alignment with user needs. Both **REMSA** and **REMSA**-NAIVE perform notably better than retrieval-only or unstructured RAG baselines, showing the effectiveness of our modular architecture: grounding the selection process in a structured schema and enabling tool-based reasoning provides a substantial advantage. However, **REMSA** improves further across all major evaluation metrics compared to **REMSA**-NAIVE. For example, Average Top-1 Score improves from 72.67 to 75.76, HQ Hit Rate from 37.33% to 40.00%, and MRR from 0.29 to 0.34. These gains suggest that our orchestration logic, including multi-turn clarification and decision heuristics, contributes meaningfully to selection quality, especially when model choices are ambiguous or task formulations are complex.

Table 3: Model selection quality comparison to the baselines (DeepSeek3.2)

System	Avg Top-1	Avg Set	Top-1 Hit	HQ Hit	MRR
REMSA (Ours)	75.35	73.81	18.67%	40.00%	0.30
REMSA -NAIVE	72.03	71.83	16.51%	36.89%	0.26
DB-RETRIEVAL	67.37	68.87	12.00%	17.33%	0.23
UNSTR.-RAG	69.19	70.94	10.67%	24.00%	0.24

Table 4: Model selection quality comparison to the baselines (LLaMA-3.3-70B)

System	Avg Top-1	Avg Set	Top-1 Hit	HQ Hit	MRR
REMSA (Ours)	73.39	70.34	14.67%	32.00%	0.26
REMSA -NAIVE	69.02	69.00	14.23%	29.47%	0.24
DB-RETRIEVAL	67.37	68.87	12.00%	17.33%	0.23
UNSTR.-RAG	69.87	68.04	10.00%	26.67%	0.22

Latency Trade-off. To assess the latency-performance trade-off, we measure the average end-to-end runtime per query. As expected, single-step methods are faster: DB-Retrieval takes 0.77s, Unstr.-RAG 11.9s, and **REMSA**-Naive 22.7s, whereas **REMSA** requires **31.7s** due to multi-stage decision processing and optional clarification. Despite this moderate overhead, **REMSA** delivers the highest expert-scored selection quality across major metrics, indicating that its additional reasoning steps yield meaningful and consistent gains.

6.2 Effect of different LLM backbones

To evaluate the robustness of **REMSA** under different LLM backbones, we further compare results using GPT-4.1, DeepSeek3.2, and LLaMA3.3-70B. As shown in Tab. 2, Tab. 3, and Tab. 4, across all three cores, **REMSA** consistently outperforms the corresponding baselines, indicating that the improvements primarily stem from the overall architecture rather than a specific LLM. Selection quality is strongest with GPT-4.1 (Avg Top-1: 75.76; MRR: 0.34), followed closely by DeepSeek3.2 (75.35; 0.30), while LLaMA3.3-70B shows slightly lower but still competitive results (73.39; 0.26). Notably, HQ Hit Rate remains relatively stable across GPT-4.1 and DeepSeek3.2 (both 40.00%), suggesting that high-quality candidate identification is robust to backbone choice. The consistent margin over **REMSA**-NAIVE and retrieval-based baselines across all LLM backbones demonstrates that structured grounding and tool-based orchestration provide consistent gains independent of the underlying language model.

6.3 Sensitivity Analysis on Evaluation Criteria

To understand which rubric dimensions drive the observed selection quality, we perform a sensitivity analysis by removing each scoring criterion individually from the expert evaluation protocol. As shown in Tab. 5, the performance is generally robust in most dimensions, but some removals reveal important insights into which criteria contribute the most to effective model selection. Both the removal of Application Compatibility and Modality Match lead to notable performance drops, confirming that **REMSA** actively prioritizes functionally appropriate models aligned with the user’s

Table 5: Sensitivity of expert-scored suitability metrics to individual rubric criteria.

Criteria Setting	Avg Set	Top-1 Hit	MRR	Note
Full Scoring (All Criteria)	75.03	22.67%	0.38	
w/o Application Compatibility	73.32	21.33%	0.36	Green: Increase Red: Drop
w/o Modality Match	70.88	22.67%	0.36	
w/o Reported Performance	75.05	22.67%	0.38	
w/o Efficiency	80.23	25.33%	0.38	
w/o Popularity+Recency	75.13	25.33%	0.39	
w/o Generalizability	75.10	22.67%	0.38	

objective. Notably, removing Reported Performance and Generalizability yields minimal change in overall results, implying that these dimensions are either captured implicitly through other criteria or are less decisive in the current benchmark setup. In contrast, removing Efficiency or Popularity+Recency actually leads to a modest performance gain. This suggests that while these criteria add practical relevance for deployment, they may occasionally favor well-known or resource-efficient models over candidates that better satisfy the core application and modality constraints. The sensitivity results further suggest that **REMSA** does not primarily rely on superficial indicators such as citations or recency, but instead emphasizes core compatibility and constraint satisfaction in its final decisions.

7 Conclusion and Discussion

We proposed **REMSA**, an LLM agent that combines a structured FM database with task-aware orchestration for real RSFM selection problems. By orchestrating modular tools for retrieval, ranking, clarification, explanation, and memory augmentation, **REMSA** enables transparent and consistent RSFM selections under user-specified constraints. Our structured RSFM database RS-FMD consolidates heterogeneous descriptions for transparent retrieval and comparison. Across our rubric-based expert benchmark, **REMSA** outperforms retrieval-only, unstructured RAG, and naive agent baselines in selection quality. This work opens several directions for future research, including expanding the benchmark to more complex scenarios, developing unified execution interfaces and lightweight accuracy-cost benchmarks for well-standardized RSFM categories, incorporating learned performance-prediction modules when sufficient standardized transfer data becomes available, and exploring adaptive FM selection strategies.

Although **REMSA** performs well in RSFM selection, several *limitations* remain. Our benchmark includes 100 expert-scored queries, which may miss rare or emerging use cases, although the evaluation requires substantial effort with 3,000 expert ratings. In addition, the ranking relies on in-context learning rather than supervised training, which may limit selection quality on complex queries.

Acknowledgements

This work is supported by the European Research Council Project Agent-BigEarth (No. 101292498) and the European Union Horizon Project PROTEUS (No. 101296397).

References

1. Adorni, P., Pham, M., et al.: Towards efficient benchmarking of foundation models in remote sensing: A capabilities encoding approach. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR (2025), https://openaccess.thecvf.com/content/CVPR2025W/MORSE/html/Adorni_Towards_Efficient_Benchmarking_of_Foundation_Models_in_Remote_Sensing_A_CVPRW_2025_paper.html, online; accessed 20-February-2026
2. Astruc, G., Gonthier, N., et al.: Omnisat: Self-supervised modality fusion for earth observation. In: Computer Vision - ECCV (2024). https://doi.org/10.1007/978-3-031-73390-1_24
3. Cekmeceli, K., Himmetoglu, M., et al.: Do vision foundation models enhance domain generalization in medical image segmentation? In: Computer Vision - ECCV Workshops (2024). https://doi.org/10.1007/978-3-031-91672-4_12, online; accessed 20-February-2026
4. Chu, M., Zheng, Z., et al.: Towards natural language-guided drones: Geotext-1652 benchmark with spatial relation matching. In: Computer Vision - ECCV (2024). https://doi.org/10.1007/978-3-031-73247-8_13, online; accessed 20-February-2026
5. DeepSeek-AI: Deepseek-v3.2: Pushing the frontier of open large language models. CoRR **abs/2512.02556** (2025). <https://doi.org/10.48550/ARXIV.2512.02556>, online; accessed 20-February-2026
6. Don, M., Pinson, S., et al.: Foundation model or finetune? evaluation of few-shot semantic segmentation for river pollution. In: Computer Vision - ECCV Workshops (2024). https://doi.org/10.1007/978-3-031-92089-9_21, online; accessed 20-February-2026
7. Drusch, M., Bello, U.D., et al.: Sentinel-2: Esa’s optical high-resolution mission for gmes operational services. *Remote sensing of Environment* **120**, 25–36 (2012)
8. Fan, Y., Ma, X., et al.: Videoagent: A memory-augmented multimodal agent for video understanding. In: Computer Vision - ECCV (2024). https://doi.org/10.1007/978-3-031-72670-5_5, online; accessed 20-February-2026
9. Feurer, M., Klein, A., et al.: Efficient and robust automated machine learning. In: Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems (2015), <https://proceedings.neurips.cc/paper/2015/hash/11d0e6287202fced83f79975ec59a3a6-Abstract.html>, online; accessed 20-February-2026
10. Guanter, L., Kaufmann, H., et al.: The enmap spaceborne imaging spectroscopy mission for earth observation. *Remote. Sens.* **7**(7), 8830–8857 (2015). <https://doi.org/10.3390/RS70708830>, online; accessed 20-February-2026
11. Guo, X., Lao, J., et al.: Skysense: A multi-modal remote sensing foundation model towards universal interpretation for earth observation imagery. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR (2024). <https://doi.org/10.1109/CVPR52733.2024.02613>, online; accessed 20-February-2026
12. Kapoor, R., Butala, Y.P., et al.: Omniact: A dataset and benchmark for enabling multimodal generalist autonomous agents for desktop and web. In: Computer Vision - ECCV (2024). https://doi.org/10.1007/978-3-031-73113-6_10, online; accessed 20-February-2026
13. Lacoste, A., Lehmann, N., et al.: Geo-bench: Toward foundation models for earth monitoring. In: Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems, NeurIPS (2023), http://papers.nips.cc/paper_files/paper/2023/hash/a0644215d9cff6646fa334dfa5d29c5a-Abstract-Datasets_and_Benchmarks.html, online; accessed 20-February-2026
14. LangChain: Langchain. <https://python.langchain.com/docs/introduction/> (2025), online; accessed 20-February-2026
15. Lee, C., Paramanayakam, V., et al.: Multi-agent geospatial copilots for remote sensing workflows. CoRR **abs/2501.16254** (2025). <https://doi.org/10.48550/ARXIV.2501.16254>, online; accessed 20-February-2026

16. Li, Y., Tan, J., et al.: Unleashing the potential of remote sensing foundation models via bridging data and computility islands. *The Innovation* (2025)
17. Li, Z., Hou, B., et al.: Masked angle-aware autoencoder for remote sensing images. In: *Computer Vision - ECCV* (2024). https://doi.org/10.1007/978-3-031-73242-3_15, online; accessed 20-February-2026
18. Liu, S., Cheng, H., et al.: Llava-plus: Learning to use tools for creating multimodal agents. In: *Computer Vision - ECCV* (2024). https://doi.org/10.1007/978-3-031-72970-6_8, online; accessed 20-February-2026
19. Luo, Y., Ru, X., et al.: Oneke: A dockerized schema-guided LLM agent-based knowledge extraction system. In: *Companion Proceedings of the ACM on Web Conference 2025, WWW* (2025). <https://doi.org/10.1145/3701716.3715189>, online; accessed 20-February-2026
20. Meta: Faiss. <https://ai.meta.com/tools/faiss/> (2025), online; accessed 20-February-2026
21. Muhtar, D., Li, Z., et al.: Lhrs-bot: Empowering remote sensing with vgi-enhanced large multimodal language model. In: *Computer Vision - ECCV* (2024). https://doi.org/10.1007/978-3-031-72904-1_26
22. Nedungadi, V., Kariyaa, A., et al.: Mmearth: Exploring multi-modal pretext tasks for geospatial representation learning. In: *Computer Vision - ECCV* (2024). https://doi.org/10.1007/978-3-031-73039-9_10, online; accessed 20-February-2026
23. Neutatz, F., Lindauer, M., et al.: Automl in heavily constrained applications. *VLDB J.* **33**(4), 957–979 (2024). <https://doi.org/10.1007/S00778-023-00820-1>, online; accessed 20-February-2026
24. OpenAI: Gpt-4.1. <https://openai.com/index/gpt-4-1/> (2025), online; accessed 20-February-2026
25. Plekhanova, E., Robert, D., et al.: Ssl4eco: A global seasonal dataset for geospatial foundation models in ecology. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR* (2025), https://openaccess.thecvf.com/content/CVPR2025W/EarthVision/html/Plekhanova_SSL4Eco_A_Global_Seasonal_Dataset_for_Geospatial_Foundation_Models_in_CVPRW_2025_paper.html, online; accessed 20-February-2026
26. Purohit, M., Muhawenayo, G., et al.: How does the spatial distribution of pre-training data affect geospatial foundation models? *CoRR* **abs/2501.12535** (2025). <https://doi.org/10.48550/ARXIV.2501.12535>, online; accessed 20-February-2026
27. Ramachandran, R., Roy, e.a.: A primer for assessing foundation models for earth observation. *Cornell University* (2025)
28. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP* (2019). <https://doi.org/10.18653/V1/D19-1410>, online; accessed 20-February-2026
29. Shabbir, A., Munir, M.A., et al.: Thinkgeo: Evaluating tool-augmented agents for remote sensing tasks. *CoRR* **abs/2505.23752** (2025). <https://doi.org/10.48550/ARXIV.2505.23752>, online; accessed 20-February-2026
30. Simumba, N., Lehmann, N., et al.: Geo-bench-2: From performance to capability, rethinking evaluation in geospatial AI. *CoRR* **abs/2511.15658** (2025). <https://doi.org/10.48550/ARXIV.2511.15658>, online; accessed 20-February-2026
31. Soni, S., Dudhane, A., et al.: Earthdial: Turning multi-sensory earth observations to interactive dialogues. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR* (2025). <https://doi.org/10.1109/CVPR52734.2025.01334>, online; accessed 20-February-2026
32. Team, L.: The llama 3 herd of models. *CoRR* **abs/2407.21783** (2024). <https://doi.org/10.48550/ARXIV.2407.21783>, online; accessed 20-February-2026
33. Thornton, C., Hutter, F., et al.: Auto-weka: combined selection and hyperparameter optimization of classification algorithms. In: *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD* (2013). <https://doi.org/10.1145/2487575.2487629>, online; accessed 20-February-2026
34. Torres, R., Snoeij, P., et al.: Gmes sentinel-1 mission. *Remote sensing of Environment* **120**, 9–24 (2012)

- 35. Wang, X., Zhang, Y., et al.: Videoagent: Long-form video understanding with large language model as agent. In: Computer Vision - ECCV (2024). https://doi.org/10.1007/978-3-031-72989-8_4, online; accessed 20-February-2026
- 36. Wang, Y., Li, K., et al.: Internvideo2: Scaling foundation models for multimodal video understanding. In: Computer Vision - ECCV (2024). https://doi.org/10.1007/978-3-031-73013-9_23, online; accessed 20-February-2026
- 37. Xiao, A., Xuan, W., et al.: Foundation models for remote sensing and earth observation: A survey. CoRR **abs/2410.16602** (2024). <https://doi.org/10.48550/ARXIV.2410.16602>, online; accessed 20-February-2026
- 38. Xu, W., Yu, Z., et al.: Rs-agent: Automating remote sensing tasks through intelligent agents. CoRR **abs/2406.07089** (2024). <https://doi.org/10.48550/ARXIV.2406.07089>, online; accessed 20-February-2026