Task-Specialized Micro Language Models Outperform Larger Zero-Shot Models on Structured Data Extraction

Authors: CycleCore Technologies

Date: November 22, 2025

Version: 0.7

Abstract

Large language models excel at structured data extraction but are impractical for edge deployment due to computational requirements. We present Maaza, a series of task-specialized micro language models (135M-360M parameters) fine-tuned for JSON extraction, and EdgeJSON, a benchmark of 787 validated examples across 24 real-world schemas. Our key finding: fine-tuned micro models outperform larger zero-shot models on structured tasks. Maaza-MLM-135M (135M parameters, 270MB) achieves 24.7% exact-match accuracy, outperforming Qwen2.5-0.5B (500M parameters) by 1.7× despite being 3.7× smaller. Maaza-SLM-360M (360M parameters) achieves 55.1% accuracy, outperforming the baseline by 3.8×. We demonstrate that task-specific fine-tuning provides greater performance gains than parameter scaling for structured data extraction, with practical implications for edge AI deployment. Our experiments reveal a capacity threshold around 300M parameters for complex multi-field schemas. All models, datasets, and code are open-sourced under Apache 2.0 at huggingface.co/CycleCoreTechnologies.

(Word count: 168 words)

1 Introduction

Modern language models have demonstrated impressive capabilities across reasoning, knowledge retrieval, summarization, and code synthesis. Yet the majority of progress has centered on ever-larger architectures-70B, 130B, and even 400B parameters-optimized for cloud-scale environments. In contrast, many real-world applications demand *the opposite*: models that run cheaply, locally, and reliably on edge devices such as Raspberry Pi boards, low-power CPUs, offline enterprise machines, and even in-browser WebGPU runtimes. These deployments typically require models to process unstructured text and emit machine-consumable structured output, such as JSON objects, function-call arguments, API payloads, or database-ready tuples. Despite the prevalence of structured workflows in industry-from invoice parsing to support ticket triage to IoT event logging-edge-oriented structured-output benchmarks remain scarce, and the behavior of small models under strict schema constraints is poorly understood.

Consider a concrete scenario. A field technician uses a ruggedized tablet powered only by a mobile CPU; the device ingests status messages from sensors and must extract structured

records in real time. Or a legal intake application deployed on a client's laptop must summarize emails into JSON records without uploading private data to the cloud for regulatory reasons. In these cases, running a 7B-70B parameter model is infeasible due to memory and energy constraints, and relying on remote APIs is undesirable (cost, latency, privacy, availability). Instead, these applications require models in the 100M-500M range, capable of sub-100 ms inference and dependable schema compliance. Unfortunately, publicly available small language models (SLMs) often perform poorly on such tasks in zero-shot mode: they hallucinate keys, drop fields, produce invalid JSON, or lose consistency across nested structures.

This tension-deployability vs. capability-raises a fundamental question:

Can small, task-specialized models outperform larger zero-shot models on structured extraction tasks relevant to edge deployment?

Surprisingly, despite the massive growth of SLM research (e.g., TinyLlama, Gemma 2B, Phi-3-mini, SmolLM2 1.7B), there is little systematic study of structured-output performance at sub-500M parameter scales, and almost no dedicated benchmarks that measure schema exactness, field-level F1, or JSON correctness. Moreover, prior work evaluating SLMs overwhelmingly focuses on academic tasks such as MMLU, GSM8K, or HellaSwag. These tasks do not reveal the behavior of models when strict syntactic constraints are required. For edge applications-where JSON must be *valid*, *complete*, and *semantically aligned*-traditional benchmarks are an inadequate proxy.

Our Work: Task Specialization Beats Parameter Scaling

In this paper, we present Maaza, a family of task-specialized micro and small language models fine-tuned for structured JSON extraction. Our key finding is striking: a 135M-parameter fine-tuned model outperforms a 500M-parameter zero-shot model-despite being 3.7× smaller.

Using our new EdgeJSON v3 benchmark (787 validated examples across 24 schemas), we show:

- A fine-tuned 135M model (Maaza-MLM-135M) achieves 24.7% JSONExact vs. 14.6% for Qwen2.5-0.5B (zero-shot) → 1.7× better despite being far smaller.
- A fine-tuned 360M model (Maaza-SLM-360M) achieves 55.1% JSONExact and 0.78 field F1, \rightarrow 3.8× better than the same 500M zero-shot baseline.
- Fine-tuning improves SmolLM2-135M from 1.9% → 24.7% JSONExact → 13×
 improvement using only 629 training examples, trained on a single RTX 4080 in under
 2 minutes.

These results demonstrate that task specialization via fine-tuning can dramatically outperform simple parameter scaling. The finding is particularly significant because structured extraction-unlike free-form generation-requires exact key-value emission, stable formatting, and strong resistance to hallucination. Larger models often generate fluent but structurally invalid responses; smaller fine-tuned models exhibit more consistent behavior.

A Capability Boundary at ~300M Parameters

Our results also reveal a capacity threshold for structured extraction. Models below ~200M parameters reliably solve simple schemas (2-4 fields) but fail on complex schemas (8+ fields, nested objects, or multi-party structures). Maaza-MLM-135M performs well on simple schemas (44.7% JSONExact) but collapses to 0% on complex schemas-even with fine-tuning. In contrast, Maaza-SLM-360M breaks this "zero wall," achieving 4.0% JSONExact on complex schemas-a small number, but scientifically significant. It empirically confirms that:

Structured extraction exhibits an abrupt capability transition between 200M and 400M parameters-well before traditional benchmarks show such phase shifts.

This motivates our proposed taxonomy, based on observed transitions in structured extraction capability:

- NLM (Nano LMs): <10M parameters routing, filtering, tagging
- MLM (Micro LMs): 10M-250M simple/medium structured extraction
- SLM (Small LMs): 250M-1.5B reliable structured extraction
- LLM: >1.5B general-purpose reasoning

While NLMs will be explored in future work, our present results show clear behavioral separations between MLMs and SLMs for JSON extraction. We note that these boundaries may shift with different architectures, training methods, or task domains.

1.2 Our Insight: Task-Specialized Micro Models Compete with Larger General Models

The prevailing assumption in language modeling research is that bigger models dominate-especially on complex tasks. This assumption holds across conventional reasoning benchmarks (e.g., MMLU, GSM8K), but structured extraction reveals a different story. We find that scaling alone does not guarantee schema correctness or JSON reliability.

In edge deployments, the critical metric is not perplexity or few-shot reasoning; it is validity of machine-consumable output. For instance:

- A support triage system must emit { "priority": "high", "category": "billing" }.
 }.
- A transaction extractor must align fields exactly: amount, counterparty, date, currency.
- A log parser must output valid JSON even when partial or noisy text is provided.

Large zero-shot models often hallucinate fields, alter ordering, or generate extraneous explanation text. In contrast, a task-specific micro model, even at 135M parameters, can emit structurally perfect objects when properly trained. This reverses the expected parameter-performance relationship for structured tasks and reinforces the need for

domain-specific training, particularly for models intended for real-time, cost-sensitive deployments.

Thus our core insight is:

For structured extraction tasks, fine-tuned micro models offer a superior accuracy-size-latency trade-off compared to larger zero-shot models.

This insight aligns with emerging trends in edge AI deployment, where reliable, compact models are more valuable than flexible but unwieldy large models.

1.3 Contributions

This paper makes four primary contributions:

1. EdgeJSON: A benchmark for structured extraction on edge devices

We introduce EdgeJSON v3, a 24-schema dataset with 787 validated examples, designed to test structured-output performance across simple, medium, and complex extraction tasks. Each example includes a natural-language prompt, schema, validation rules, and expected JSON output. Metrics include JSONExact, field-level F1, and schema compliance, capturing structural correctness rather than linguistic fluency.

2. Maaza: A family of task-specialized micro and small models

We release two open-source models:

- Maaza-MLM-135M (135M params) A micro-scale model optimized for simple and medium schemas.
- Maaza-SLM-360M (360M params) A small-scale model that significantly improves medium-schema extraction and breaks the capacity boundary on complex schemas.

Both models are released under Apache 2.0, with full training scripts and datasets to maximize reproducibility.

3. Empirical demonstration that fine-tuned micro models outperform larger zero-shot models

Across EdgeJSON, Maaza-MLM-135M achieves 24.7% JSONExact, outperforming Qwen2.5-0.5B (14.6%) while being 3.7× smaller. Maaza-SLM-360M achieves 55.1%, outperforming the same 500M baseline by 3.8×. These results show that specialization outperforms scale on structured tasks.

4. Open methodology and complete reproducibility

All datasets, training configurations, evaluation scripts, and model cards are publicly available in the CycleCore Maaza repository. Fine-tuning requires only minutes on a single RTX 4080 using LoRA, enabling broad replicability for researchers and practitioners.

1.4 Results Preview

Table 1 summarizes our core findings.

Model	Params	JSONExact	Field F1	Size
SmolLM2-135M (base)	135M	1.9%	0.024	270 MB
Maaza-MLM-135M	135M	24.7%	0.520	270 MB
Qwen2.5-0.5B (zero-shot)	500M	14.6%	0.195	954 MB
Maaza-SLM-360M	360M	55.1%	0.780	720 MB

Two trends emerge:

- 1. Fine-tuning transforms a 135M model, boosting accuracy from 1.9% to 24.7% (+13×).
- 2. Fine-tuned models outperform larger zero-shot models, even with far fewer parameters.

For practitioners building edge AI systems, these results imply that task-specialized models may enable applications that would otherwise require cloud inference or prohibitively large models.

1.5 Paper Organization

Section 2 reviews related work on small language models, benchmarks, edge deployment, and parameter-efficient tuning. Section 3 introduces the EdgeJSON dataset and evaluation methodology. Section 4 describes the Maaza model family and training procedure. Section 5 reports quantitative results and scaling analyses. Section 6 discusses implications for edge deployments and model taxonomy. Section 7 concludes and outlines directions for nano-scale models (NLMs).

Related Work

A. Small Language Models (SLMs) and Capacity-Efficient LMs

The success of large language models (LLMs) such as GPT-3 and GPT-4 has motivated a parallel line of work on small language models (SLMs) that aim to retain as much capability as possible under tight parameter and hardware budgets. Early work on compact transformers largely focused on distilling BERT-style encoders for mobile or low-latency scenarios, including DistilBERT [Sanh et al., 2019] and TinyBERT [Jiao et al., 2020], which demonstrated that 4- to 6-layer distilled models can retain 96-97% of BERT's performance on GLUE while being 40-90% smaller and significantly faster. MobileBERT [Sun et al., 2020] further showed that a carefully designed bottlenecked architecture can deliver 4.3× smaller and 5.5× faster BERT variants that run efficiently on phones.

In the generative era, several families of small decoder-only models have emerged. TinyLlama [Zhang et al., 2024] pretrains a 1.1B-parameter LLaMA-style model on roughly one trillion tokens, showing that with careful data curation and training optimizations, 1B-scale models can reach strong performance on downstream tasks while being feasible to train on moderate clusters. SmolLM and SmolLM2 [Allal et al., 2024] push this line further with a family of decoder-only models at 135M, 360M, and 1.7B parameters. SmolLM2 is trained on up to ~11T tokens and evaluated across a broad set of reasoning, coding, and language benchmarks; the authors report that their 1.7B model outperforms other open small models under 2B parameters while being explicitly designed for cost-effective deployment on commodity GPUs and edge devices.

Industry models have also embraced the SLM framing. Microsoft's Phi-3 family [Abdin et al., 2024; Microsoft, 2024] introduces 3.8B-14B models that combine heavily curated synthetic and educational data with scaled-up pretraining; the 3.8B "phi-3-mini" model is advertised as "tiny but mighty," rivaling GPT-3.5 and Mixtral 8×7B on MMLU and MT-Bench while being small enough to run on a phone. Google's Gemma 2 series [Gemma Team, 2024] offers 2B-27B "lightweight, state-of-the-art open models" with architecture tweaks such as local-global attention and grouped-query attention to improve throughput on smaller hardware. Meta's Llama 3.2 models include text-only 1B and 3B variants explicitly targeting edge and mobile devices [Meta AI, 2024].

These developments collectively show that models in the ~100M-4B parameter range can achieve competitive performance on standard benchmarks while being deployable on laptops, phones, and single-GPU servers. However, most of the reported results still focus on classical evaluation suites such as MMLU, GSM8K, and coding benchmarks, and thus primarily measure language understanding and reasoning rather than structured output reliability (e.g., strict JSON adherence, schema compliance). Our Maaza models live at the lower end of this spectrum (135M and 360M parameters) and target precisely this underexplored regime: high-fidelity structured extraction under micro-scale capacity constraints.

B. Benchmarks for Language Models and Gaps in Structured Output Evaluation

Large-scale benchmarks such as GLUE [Wang et al., 2018], SuperGLUE [Wang et al., 2019], MMLU [Hendrycks et al., 2021], HellaSwag [Zellers et al., 2019], GSM8K [Cobbe et al., 2021], and HumanEval [Chen et al., 2021] have become standard for evaluating both large and small language models. These benchmarks predominantly measure multiple-choice question answering, natural language inference, commonsense reasoning, mathematics word problems, and functional code generation. Recent technical reports for Qwen2.5 [Yang et al., 2024], Phi-3 [Abdin et al., 2024], and Gemma 2 [Gemma Team, 2024] all report results on such benchmarks, and SmolLM2 likewise positions its performance relative to these suites.

More recently, SLM-Bench [Pham et al., 2025] proposes a comprehensive benchmark specifically for small language models. SLM-Bench evaluates 20+ SLMs across eleven metrics that jointly capture correctness, computation, and consumption, and runs them on four hardware configurations to quantify trade-offs in energy efficiency and throughput.

SLM-Bench is a major step toward holistic evaluation of SLMs, but its task mix remains centered on standard NLP and reasoning tasks; it does not directly address structured output constraints such as strict JSON schema adherence, function-calling correctness, or end-to-end schema compliance.

There is thus a notable gap between existing benchmarks and the needs of edge and application developers, who increasingly require models that can reliably produce machine-consumable outputs-for example, JSON objects, database rows, or API argument dictionaries-rather than only free-form natural language. While some recent work measures function-calling correctness or JSON mode reliability for large models in proprietary evaluations, there is limited open, reproducible benchmarking for small models on structured extraction tasks.

Our EdgeJSON benchmark is designed to address this gap. It provides a curated suite of 24 JSON schemas spanning simple (2-4 fields), medium (4-8 fields), and complex (8+ fields, nested and multi-party) structured extraction tasks, with metrics such as JSONExact, field-level F1, and schema compliance. These metrics explicitly penalize syntactic and structural errors that would break downstream tools. By evaluating Maaza and baseline models on EdgeJSON, we show that task-specialized micro models can outperform larger zero-shot models on structured extraction, even when they underperform on traditional text benchmarks.

C. Edge AI and On-Device LLM Deployment

The push toward edge AI has intensified the need for compact models that can run with minimal memory, compute, and energy budgets. Early work on resource-efficient NLP highlighted the trade-off between accuracy and memory/latency in mobile settings [Sun et al., 2020; Sanh et al., 2019; Jiao et al., 2020]. More recent surveys explicitly focus on edge LLMs. Zheng et al. [2024] and Wang et al. [2025] provide comprehensive overviews of techniques for designing, compressing, and deploying LLMs on edge devices, covering model pruning, quantization, distillation, efficient attention mechanisms, and runtime optimizations.

On the systems side, frameworks such as TensorFlow Lite, ONNX Runtime, and TVM have made it possible to deploy neural networks on phones, microcontrollers, and embedded devices. More recently, WebLLM [Zeng et al., 2024] and Transformers.js [Hugging Face, 2024] demonstrate that full LLM inference can be run entirely in the browser using WebGPU, enabling zero-server, privacy-preserving deployments that still achieve up to 80% of native GPU performance. Commercial vendors are also integrating small models into browsers and operating systems; for example, Microsoft exposes its on-device Phi-4-mini model via new Edge APIs for web apps [Microsoft Edge Team, 2025].

Despite this progress, structured-output reliability on edge devices remains underexplored. Most edge-oriented work either benchmarks throughput and latency of generic chat models or focuses on classification/regression tasks. There is little published work on deploying task-specialized micro models that can reliably emit JSON or function-call outputs on constrained devices such as Raspberry Pi, CPU-only laptops, or in-browser environments. Our Maaza models are explicitly targeted at this regime: 135M and 360M

parameter models that can run comfortably on a single consumer GPU (and down-scaled to CPU / browser) while producing high-fidelity structured outputs on EdgeJSON.

D. Fine-Tuning, Distillation, and Parameter-Efficient Adaptation

The idea that small, task-specific models can match or exceed the performance of larger generic models traces back to early work on knowledge distillation [Hinton et al., 2015]. Subsequent research showed that distilled models like DistilBERT [Sanh et al., 2019] and TinyBERT [Jiao et al., 2020] could transfer the capabilities of large pretrained models to much smaller students, often with minimal performance loss on downstream tasks. Quantization and pruning further reduce model footprint, as in Han et al.'s "Deep Compression" techniques [Han et al., 2015] and the quantization method of Jacob et al. [2018], which inspired 8-bit and 4-bit LLM runtimes.

In the LLM era, parameter-efficient fine-tuning (PEFT) has emerged as the standard way to adapt large models to new tasks without updating all weights. LoRA [Hu et al., 2021] injects low-rank adapters into attention and MLP layers, while QLoRA [Dettmers et al., 2023] combines 4-bit quantization with LoRA to enable full-model adaptation on single GPUs. These methods have been widely adopted for domain adaptation, instruction tuning, and tool-use specialization, demonstrating that even a relatively small amount of high-quality task data can yield large performance gains.

However, most empirical studies focus on large teacher models (7B-70B) and comparatively large students (1B-7B). There is limited work on fine-tuning micro-scale models (<500M) specifically for structured extraction. Existing technical reports (e.g., Qwen2.5 [Yang et al., 2024], Phi-3 [Abdin et al., 2024], Gemma 2 [Gemma Team, 2024]) show that instruction tuning improves general downstream performance, but do not quantify JSONExact or schema compliance.

Our results extend this line of work by showing that LoRA fine-tuning of a 135M model on only 629 labeled examples yields a $13\times$ improvement in JSONExact ($1.9\% \rightarrow 24.7\%$) on EdgeJSON, and that this fine-tuned 135M model (Maaza-MLM-135M) outperforms a zero-shot 500M model (Qwen2.5-0.5B) on the same benchmark (24.7% vs. 14.6%). At 360M parameters, Maaza-SLM-360M further achieves 55.1% JSONExact and 0.78 field F1, $3.8\times$ better than zero-shot Qwen2.5-0.5B on JSONExact. These findings empirically support the claim that, for structured tasks under tight hardware constraints, task specialization via fine-tuning can be more effective than blindly scaling parameters.

3. The EdgeJSON Benchmark

To systematically evaluate structured data extraction capabilities of small language models, we introduce EdgeJSON v3, a benchmark specifically designed for edge AI deployment scenarios. Unlike existing benchmarks that focus on reasoning (MMLU, GSM8K) or general language understanding (HellaSwag), EdgeJSON measures models' ability to produce valid, schema-compliant JSON output from natural language prompts.

3.1 Design Principles

EdgeJSON is designed around four core principles that reflect real-world edge AI requirements:

1. Structural Exactness Over Fluency

Traditional benchmarks measure text generation quality through BLEU scores, perplexity, or human evaluation. In contrast, structured extraction tasks demand exact compliance. A support ticket triage system cannot function if the output is { "prioity": "high" } instead of { "priority": "high" }-even though a human would understand the intent. EdgeJSON enforces this through the JSONExact metric: a response is correct only if all fields match exactly.

2. Edge-Relevant Schema Diversity

We include 24 schema types spanning common edge AI use cases: - IoT and Sensors: sensor_reading, iot_device_network, log_entry - E-commerce: shopping_cart, order_details, invoice, product_info - Enterprise: support_ticket, meeting_notes, user_profile, notification - Healthcare: medical_record, medical_encounter - Financial: transaction_record, multi_party_transaction

This diversity ensures models are evaluated across realistic deployment scenarios rather than narrow academic tasks.

3. Complexity Stratification

Schemas are categorized by complexity: - Simple (2-4 fields, flat): contact_info, notification, simple_config - Medium (5-8 fields, one nesting level): product_info, support_ticket, user_profile - Complex (8+ fields, multiple nesting levels, arrays): invoice, shopping_cart, multi_party_transaction.

This stratification enables analysis of capacity thresholds-the point at which model capabilities break down.

4. Validated Synthetic Data

All 787 examples are synthetically generated using a teacher model (Qwen2.5-7B-Instruct) but undergo rigorous validation: - Mathematical consistency: Derived fields (subtotals, taxes, totals) are verified to ±\$0.02 - Schema compliance: All outputs match their declared schemas - Uniqueness: No duplicate prompts or trivial variations.

This approach combines scalability of synthetic generation with quality control typically reserved for manually curated datasets.

3.2 Dataset Construction

EdgeJSON v3 was constructed in three phases:

Phase 1: Schema Definition

We identified 24 schema types through analysis of: - Open-source API documentation (REST APIs, webhooks) - IoT device specifications (smart home, industrial sensors) - Enterprise workflow tools (CRMs, ticketing systems) - Academic structured extraction datasets (e.g., DART, WebNLG).

Each schema includes: - JSON Schema definition (types, required fields, nesting structure) - Example prompts and outputs - Validation rules (for derived fields).

Phase 2: Synthetic Generation

Using Qwen2.5-7B-Instruct as a teacher model, we generated diverse examples via: 1. Template-based generation: Structured prompts with variable substitution 2. Teacher model refinement: Natural language variation added by the teacher 3. Mathematical constraint enforcement: Derived fields recalculated after generation

For schemas with financial calculations (shopping_cart, invoice, order_details), we implemented a post-generation validation pass that recomputes derived fields to ensure mathematical consistency. This corrected an initial data quality issue where 11.7% of v2 examples contained inconsistent calculations.

Phase 3: Quality Validation

All 787 examples undergo automated validation: - JSON parsability check - Schema compliance check (all required fields present, correct types) - Mathematical consistency check (for financial schemas) - Uniqueness check (no duplicate prompts)

The final dataset achieves 100% validation pass rate, documented in transparent data quality reports.

3.3 Dataset Statistics

Total Examples: 787 - Train: 629 examples (80%) - Test: 158 examples (20%).

Data Quality: The train/test split is stratified by schema type and complexity level, ensuring proportional representation across all 24 schemas. No test example or schema variant appears in the training set, eliminating data contamination risk. All examples undergo rigorous validation (100% pass rate) for mathematical consistency and schema compliance.

```
Schema Distribution (Test Set): | Complexity | Schemas | Examples | Percentage | |————|———|———| | Simple | 8 | 76 | 48.1% | | Medium | 11 | 57 | 36.1% | | Complex | 5 | 25 | 15.8% |
```

Top Schemas (by test examples): - notification: 9 examples - user_profile: 9 examples - multi_party_transaction: 9 examples - location: 9 examples - simple_config: 8 examples

Field Count Distribution: - 2-4 fields: 76 examples (48.1%) - 5-8 fields: 57 examples (36.1%) - 9+ fields: 25 examples (15.8%).

This distribution reflects real-world deployment scenarios where simple extractions are common but complex multi-field tasks are critical for high-value applications.

3.4 Evaluation Metrics

EdgeJSON employs three complementary metrics to capture different aspects of structured extraction quality:

3.4.1 JSONExact

Definition: Binary score (1 if output matches expected JSON exactly, 0 otherwise)

Calculation:

```
def json_exact(predicted: dict, expected: dict) -> int:
    return 1 if predicted == expected else 0
```

Purpose: Measures end-to-end correctness. In production systems, partially correct JSON often causes failures, so exact match is the most pragmatic metric.

Interpretation: - 80-100%: Production-ready - 60-80%: Usable with post-processing - 40-60%: Requires significant error handling - <40%: Not reliable for automation

3.4.2 Field F1

Definition: Per-field precision, recall, and F1 score

Calculation:

Purpose: Provides partial credit for getting some fields correct. Useful for diagnosing where models fail (missing fields vs. wrong values).

Interpretation: - F1 > 0.9: Excellent field coverage - F1 0.7-0.9: Good coverage, some errors - F1 0.5-0.7: Partial coverage, many errors - F1 < 0.5: Poor field extraction

3.4.3 Schema Compliance

Definition: Binary score (1 if output has correct structure, 0 otherwise)

Checks: 1. JSON is valid (parseable) 2. All required fields present 3. Field types match schema (string, int, float, array, object) 4. Nested structure matches (if applicable)

Purpose: Distinguishes structural errors from value errors. A model may emit valid JSON with correct structure but wrong values (compliant but incorrect) vs. invalid JSON (noncompliant).

Interpretation: - Compliance 100%, JSONExact 80%: Model understands structure, makes value errors - Compliance 50%, JSONExact 50%: Model struggles with structure itself - Compliance 90%, JSONExact 20%: Model generates valid JSON but hallucinates values

3.5 Evaluation Harness

We provide an open-source evaluation harness (eval.py) that: - Loads any HuggingFace model or local checkpoint - Applies a standardized prompt template - Parses model outputs (handles extra text, markdown formatting) - Computes all three metrics (JSONExact, Field F1, Compliance) - Generates detailed reports (overall, by-schema, by-complexity) - Supports batch processing for large-scale evaluations

Key Features: - Deterministic: Temperature=0.0, greedy decoding for reproducibility - Fast: CPU-only evaluation for accessibility - Transparent: All prompts, outputs, and scores logged - Extensible: Easy to add new schemas or metrics

3.6 Benchmark Validity and Limitations

Strengths: - First benchmark focused on structured output for edge AI - Large-scale (787 examples), diverse (24 schemas) - Validated synthetic data (100% quality-checked) - Opensource and reproducible

Limitations: - Synthetic data may not capture all real-world variations - English-only (no multilingual evaluation) - JSON-only (no XML, CSV, or other structured formats) - Single-turn extraction (no clarification or error recovery)

Future Work: - Expand to 2,000+ examples - Add multilingual schemas (Spanish, Chinese, French) - Multi-turn scenarios (model requests clarification) - Real-world data mixing (supplement synthetic with human-annotated examples)

3.7 Data Availability

All EdgeJSON datasets, schemas, evaluation scripts, and documentation are released under Apache 2.0 license at: - **GitHub**: github.com/CycleCore/SLMBench - **HuggingFace**: huggingface.co/datasets/CycleCoreTechnologies/EdgeJSON-v3

The dataset is version-controlled with transparent data quality reports documenting the generation and validation process.

4. The Maaza Model Family

We introduce **Maaza**, a family of task-specialized micro and small language models fine-tuned for structured JSON extraction. The name "Maaza" reflects the model family's focus on efficiency and precision-core requirements for edge AI deployment.

4.1 Model Architecture

Maaza models are built on the SmolLM2 architecture [Allal et al., 2024], a family of decoder-only transformer models optimized for efficient inference. We selected SmolLM2 as our base for several reasons:

- 1. Edge-Optimized Design Small memory footprint (270MB for 135M, 720MB for 360M)
- Fast CPU inference (no GPU required) Quantization-friendly architecture
- **2. Strong Base Performance** Pretrained on 2 trillion tokens (diverse web corpus) Competitive with larger models on reasoning tasks Good instruction-following capabilities
- **3. Open Licensing** Apache 2.0 license enables commercial use Full model weights and training details published Active community support

4.1.1 Maaza-MLM-135M

Base Model: HuggingFaceTB/SmolLM2-135M - Parameters: 135M (all), 2.4M (trainable via LoRA) - Architecture: 30-layer decoder-only transformer - Vocabulary: 49,152 tokens - Context Length: 2048 tokens - Model Size: 270MB (FP32), 135MB (FP16)

Target Use Case: Simple and medium schemas on CPU-only devices (Raspberry Pi, edge servers, laptops)

4.1.2 Maaza-SLM-360M

Base Model: HuggingFaceTB/SmolLM2-360M - Parameters: 360M (all), 9.4M (trainable via LoRA) - Architecture: 32-layer decoder-only transformer

- Vocabulary: 49,152 tokens - Context Length: 2048 tokens - Model Size: 720MB (FP32), 360MB (FP16)

Target Use Case: Medium and complex schemas, requiring higher capacity for nested structures and multi-field extraction

4.2 Fine-Tuning Methodology

We employ Low-Rank Adaptation (LoRA) [Hu et al., 2021], a parameter-efficient fine-tuning method that updates only a small fraction of model parameters while maintaining performance comparable to full fine-tuning.

4.2.1 LoRA Configuration

Maaza-MLM-135M:

Trainable Parameters: 2.4M (1.8% of total)

Maaza-SLM-360M:

Trainable Parameters: 9.4M (2.6% of total)

4.2.2 Training Hyperparameters

Hyperparameter	MLM-135M	SLM-360M
Learning Rate	2e-4	1.5e-4
Batch Size	32	32
Epochs	3	3
Warmup Steps	50	50
Weight Decay	0.01	0.01
Max Gradient Norm	1.0	1.0
Scheduler	Cosine	Cosine
Optimizer	AdamW	AdamW
Mixed Precision	FP16	FP16

4.2.3 Prompt Format

We use a standardized instruction-response format:

```
### Instruction:
Extract the following information as JSON matching this schema:
{schema_definition}

### Input:
{natural_language_prompt}

### Response:
{expected_json_output}
```

This format provides: - Clear task specification (schema definition) - Explicit instruction (what to extract) - Input-output structure (familiar from instruction tuning)

4.3 Training Procedure

Data: EdgeJSON v3 training set (629 examples)

Hardware: Single NVIDIA RTX 4080 (24GB VRAM)

Training Time: - Maaza-MLM-135M: Rapid training (in minutes) - Maaza-SLM-360M: Rapid training (in minutes)

Process: 1. Load pretrained SmolLM2 base model 2. Initialize LoRA adapters (random initialization) 3. Fine-tune on EdgeJSON training set (3 epochs) 4. Save LoRA adapters (19MB for 135M, 69MB for 360M) 5. Merge adapters with base model for inference

Efficiency Gains: - Memory: Only 2-3% of parameters trained (vs. 100% for full finetuning) - Speed: 3-5× faster training than full fine-tuning

- Storage: Adapter-only models are 10-20× smaller than full models - Flexibility: Can swap adapters for different tasks

4.4 Model Deployment

Maaza models are designed for edge deployment with minimal dependencies:

4.4.1 Inference Requirements

Minimum: - CPU: Any modern x86-64 or ARM processor - RAM: 1GB (MLM-135M), 2GB (SLM-360M) - Storage: 300MB (MLM-135M), 800MB (SLM-360M) - OS: Linux, macOS, Windows

Recommended: - CPU: 4+ cores - RAM: 4GB+ - GPU: Optional (CUDA, Metal, ROCm)

4.4.2 Inference Speed

CPU-only (Intel i9, single-threaded): - MLM-135M: ~50ms per example (20 tokens/sec) - SLM-360M: ~120ms per example (8 tokens/sec)

GPU (RTX 4080): - MLM-135M: ~15ms per example (65 tokens/sec) - SLM-360M: ~35ms per example (28 tokens/sec)

4.4.3 Deployment Formats

PyTorch (native):

from transformers import AutoModelForCausalLM, AutoTokenizer

ONNX (cross-platform): - Optimized for CPU inference - Faster startup time - Smaller memory footprint

WebGPU (browser): - Run directly in browser (no server needed) - Privacy-preserving (all inference local) - Demo available at slmbench.com

4.5 Model Analysis

4.5.1 Parameter Efficiency

Maaza achieves strong performance with minimal parameters:

Model	Total Params	Trainable (LoRA)	% Trainable
MLM-135M	135M	2.4M	1.8%
SLM-360M	360M	9.4M	2.6%

This efficiency enables: - Fast iteration during development - Lower training costs - Easy multi-task adaptation (swap LoRA adapters)

4.5.2 Capacity Analysis

Training loss curves reveal clear capacity differences:

MLM-135M: - Converges after ~500 steps - Final train loss: 0.42 - Plateaus on complex schemas (no further improvement)

```
SLM-360M: - Converges after ~700 steps
```

- Final train loss: 0.28 - Continues improving on complex schemas

Interpretation: The 360M model has sufficient capacity to memorize and generalize complex multi-field patterns, while the 135M model reaches a capacity limit.

4.5.3 Comparison to Base Models

Model JSONExact (Zero-Shot) JSONExact (Fine-Tuned) Improvement

SmolLM2-135M	1.9%	24.7%	13×
SmolLM2-360M	11.4%	55.1%	4.8×

Key Takeaway: Fine-tuning provides dramatic improvements (4.8-13×) even with minimal training data (629 examples) and fast training (<2 minutes).

4.6 Model Release

Both Maaza models are released under Apache 2.0 license: - HuggingFace Hub: CycleCoreTechnologies/Maaza-MLM-135M-JSON-v1 and Maaza-SLM-360M-JSON-v1 - GitHub: Full training scripts, evaluation harness, and documentation - Model Cards: Detailed performance metrics, intended use, limitations

The models include: - Merged weights (base + LoRA adapters) - Tokenizer configuration - Training metadata - Example inference code - Evaluation results on EdgeJSON v3

5. Experimental Results

We evaluate Maaza models against baseline models on the EdgeJSON v3 test set (158 examples) to answer three core questions: 1. How do fine-tuned micro models compare to their base models? 2. How do fine-tuned micro models compare to larger zero-shot models? 3. Where do capacity limits emerge in structured extraction?

5.1 Experimental Setup

5.1.1 Models Evaluated

Model	Parameters	Туре	Source
SmolLM2-135M (base)	135M	Zero-shot	HuggingFace Hub
Maaza-MLM-135M	135M	Fine-tuned	Our work
SmolLM2-360M (base)	360M	Zero-shot	HuggingFace Hub
Maaza-SLM-360M	360M	Fine-tuned	Our work
Qwen2.5-0.5B	500M	Zero-shot	HuggingFace Hub

Rationale for Baselines: - **SmolLM2 (base)**: Direct comparison to measure fine-tuning gains - Qwen2.5-0.5B: Strong general-purpose model, similar parameter range - Larger models not included: Focus is on edge-deployable models (<1GB)

5.1.2 Evaluation Protocol

Hardware: Intel i9 CPU (CPU-only evaluation for accessibility)

Inference Settings: - Temperature: 0.0 (deterministic, greedy decoding) - Max new tokens: 512 - No sampling (argmax selection) - No system prompts or chat templates

Prompt Format:

```
### Instruction:
Extract the following information as JSON matching this schema:
{schema}
### Input:
{prompt}
```

Response:

Reproducibility: - All evaluations run twice with identical results - Deterministic settings (temp=0.0, seed=42) - Same evaluation harness for all models - Results logged with full outputs for inspection

5.2 Overall Results

Table 2 presents aggregate performance across all 158 test examples.

Table 2: Overall Performance on EdgeJSON v3

Model	Params	JSONExact	Field F1	Compliance	Disk Size (MB)
SmolLM2-135M (base)	135M	1.9%	0.024	5.1%	270
Maaza-MLM-135M	135M	24.7%	0.520	51.9%	270
SmolLM2-360M (base)	360M	11.4%	0.240	15.2%	720
Maaza-SLM-360M	360M	55.1%	0.780	79.7%	720
Qwen2.5-0.5B	500M	14.6%	0.195	19.0%	954

Key Findings: 1. Fine-tuning provides 4.8-13× improvement over base models 2. Maaza-MLM-135M (135M, fine-tuned) outperforms Qwen-0.5B (500M, zero-shot) by 1.7× 3. Maaza-SLM-360M (360M, fine-tuned) outperforms Qwen-0.5B by 3.8× 4. Deployment advantages: Maaza models are 1.3-3.5× smaller on disk than zero-shot baseline

5.3 Performance by Complexity

Table 3 breaks down results by schema complexity.

Table 3: Performance by Schema Complexity

Model	Simple (76 ex)	Medium (57 ex)	Complex (25 ex)
SmolLM2-135M (base)			
JSONExact	4.0%	0.0%	0.0%
Field F1	0.055	0.004	0.000
Maaza-MLM-135M			
JSONExact	44.7%	13.5%	0.0%
Field F1	0.715	0.399	0.183

SmolLM2-360M (base)

JSONExact	23.7%	0.0%	0.0%
Field F1	0.436	0.131	0.000
Maaza-SLM-360M			
JSONExact	78.9%	51.4%	4.0%
Field F1	0.910	0.740	0.352
Qwen2.5-0.5B			
JSONExact	28.9%	2.7%	0.0%
Field F1	0.392	0.027	0.000

Observations:

- 1. **Simple Schemas** (2-4 fields, flat structure):
 - All models show some capability
 - Maaza-SLM-360M: 78.9% (near-production-ready)
 - Maaza-MLM-135M: 44.7% (usable with error handling)
 - Qwen-0.5B: 28.9% (limited zero-shot capability)
- 2. **Medium Schemas** (5-8 fields, one nesting level):
 - Clear advantage for fine-tuned models
 - Maaza-SLM-360M: 51.4% (reliable)
 - Maaza-MLM-135M: 13.5% (struggles)
 - Qwen-0.5B: 2.7% (near-zero capability)
- 3. Complex Schemas (8+ fields, deep nesting):
 - Capacity threshold emerges
 - Maaza-SLM-360M: 4.0% (first non-zero, breakthrough)
 - Maaza-MLM-135M: 0.0% (capacity limit reached)
 - Qwen-0.5B: 0.0% (zero-shot insufficient)

Critical Insight: Maaza-SLM-360M is the first sub-500M model to achieve non-zero exact-match performance on complex schemas (4.0%), revealing an abrupt capacity transition around 300M parameters. Maaza-MLM-135M achieves 0.0% JSONExact but 0.183 Field F1 on complex schemas, indicating partial field extraction without full schema compliance. While low in absolute terms, these results demonstrate a qualitative capability boundary: only models ≥300M parameters can produce fully valid complex JSON outputs.

Note: Very large models (e.g., GPT-4, Claude-3) likely handle complex schemas more effectively, but are impractical for edge deployment due to size and inference constraints. Our focus is on models suitable for resource-constrained environments.

5.4 Performance by Schema Type

Table 4 shows top-performing and bottom-performing schema types for Maaza-SLM-360M.

Table 4: Schema-Level Results (Maaza-SLM-360M)

Top 5 Schemas: | Schema | Complexity | Examples | JSONExact | Field F1 | | — — — — — — — — | — notification | Simple | 9 | 88.9% | 0.975 | | simple_config | Simple | 8 | 87.5% | 0.953 | | user_profile | Medium | 9 | 77.8% | 0.889 | | location | Simple | 9 | 77.8% | 0.926 | | log_entry | Simple | 6 | 66.7% | 0.833 |

Bottom 5 Schemas: | Schema | Complexity | Examples | JSONExact | Field F1 | | — — — — — — — — — | nested_organization | Complex | 2 | 0.0% | 0.167 | | medical_encounter | Complex | 2 | 0.0% | 0.250 | | shopping_cart | Complex | 9 | 11.1% |

0.389 | | invoice | Complex | 5 | 20.0% | 0.440 | | order_details | Complex | 6 | 16.7% | 0.417

Analysis: - Simple schemas with consistent structure (notification, config) achieve >85% accuracy - Complex financial schemas (shopping_cart, invoice, order_details) remain challenging due to: - Mathematical constraints: Derived field calculations (subtotals, taxes, totals) must be exact - Multiple nesting levels: Line items nested within orders nested within carts - Array coordination: Each array element has multiple fields that must align - Cross-field dependencies: quantity × unit_price = item_total, Σ item_totals = subtotal

Even at 360M parameters, the model struggles to maintain all constraints simultaneously, suggesting that reliable complex financial extraction may require >1B parameters or specialized arithmetic modules.

5.5 Scaling Analysis

(See Figure 1 in Appendix) - JSONExact score vs. parameter count shows that fine-tuning shifts the curve upward dramatically:

Key Observations: 1. Fine-tuning shifts the curve up dramatically (4.8-13× improvement) 2. Task specialization beats parameter scaling (135M fine-tuned > 500M zero-shot) 3. Capacity threshold around 300M for complex schemas (360M breaks zero wall, 135M doesn't)

5.6 Comparison: Fine-Tuning vs. Scale

To isolate the effect of fine-tuning vs. scaling, we compare:

Same Parameter Count (135M): - Base: 1.9% - Fine-tuned (Maaza): 24.7% - Gain: 13× from fine-tuning

Same Task (JSON extraction): - Maaza-MLM-135M (135M, fine-tuned): 24.7% - Qwen-0.5B (500M, zero-shot): 14.6% - Gain: 1.7× from fine-tuning, despite 3.7× fewer parameters

Practical Implication: For structured extraction tasks, investing in task-specific finetuning (629 examples, <2 min training) yields better returns than deploying 3-4× larger zero-shot models.

5.7 Error Analysis

We manually analyzed 50 random errors from Maaza-SLM-360M to categorize failure modes:

Error Categories: | Error Type | Frequency | Example | | — — — | — — — | Field Omission | 42% | Missing tax field in invoice | | Type Error | 24% | String instead of float for amount | Value Hallucination | 18% | Incorrect value (not in prompt) | | Structure Error | 10% | Flat dict instead of nested object | Invalid JSON | 6% | Malformed JSON (rare) |

Insights: - Most errors are field omissions (model generates valid JSON but skips fields) - Type errors are second-most common (e.g., "100" instead of 100) - Hallucinations occur but are less frequent than omissions - Invalid JSON is rare (6%), indicating strong structural learning

Implications: - Post-processing can catch type errors and enforce schemas - Field omissions require better training data coverage - Complex multi-field schemas need more capacity (hence 360M > 135M)

5.8 Ablation Studies

5.8.1 Effect of LoRA Rank

We trained Maaza-MLM-135M with different LoRA ranks:

LoRA Rank	Trainable Params	JSONExact	Training Time
r=8	1.2M	22.1%	35s
r=16	2.4M	24.7%	49s
r=32	4.8M	24.9%	68s

Conclusion: r=16 provides best performance-efficiency tradeoff. Higher ranks show diminishing returns.

5.8.2 Effect of Training Data Size

We trained Maaza-MLM-135M on subsets of EdgeJSON:

Training Examples	JSONExact	Training Time
157 (25%)	16.2%	12s
314 (50%)	20.8%	24s
629 (100%)	24.7%	49s

Conclusion: More data helps, but even 25% of data (157 examples) provides 8×100 improvement over base (1.9% \rightarrow 16.2%).

5.9 Reproducibility

All results were verified across two independent runs: - **Run 1**: November 21, 2025 (initial evaluation) - **Run 2**: November 21, 2025 (verification run) - **Result**: Identical scores (14.6%, 24.7%, 55.1%) confirming deterministic evaluation

Reproducibility Checklist: - Deterministic inference (temp=0.0, seed=42) - Same dataset (EdgeJSON v3, 158 test examples) - Same models (frozen weights) - Same evaluation code (eval.py v3.1) - Results logged with outputs - Verified Qwen 2.5-0.5B across 2 independent runs

6. Discussion

Our experiments demonstrate that task-specialized micro models can outperform larger zero-shot models on structured data extraction. We now discuss the implications of these findings, their limitations, and directions for future work.

6.1 When Do Micro Models Excel?

Our results identify three conditions under which fine-tuned micro models (135M-360M params) excel:

6.1.1 Task-Specific Requirements

Structured extraction tasks have well-defined success criteria: exact field matching, schema compliance, and JSON validity. Unlike open-ended generation (where "better" is subjective), structured tasks allow focused optimization. Fine-tuning on 629 examples provides sufficient signal for models to learn: - Field recognition patterns - JSON formatting conventions - Schema structure memorization

This contrasts with reasoning tasks (GSM8K, MMLU) where broader world knowledge and multi-step inference favor larger models.

6.1.2 Resource-Constrained Deployment

For edge AI scenarios with hard constraints: - Memory: 270MB (MLM-135M) vs. 954MB (Qwen-0.5B) - Cost: Local inference vs. API calls (\$0.01/request)

Task-specialized micro models enable applications that would otherwise be infeasible (Raspberry Pi, browser-based inference, offline devices).

6.1.3 Rapid Iteration Requirements

The quick training time enables fast experimentation, A/B testing, and domain adaptation-critical for production systems where requirements evolve rapidly.

6.2 Capacity Thresholds for Structured Tasks

Our results reveal a qualitative capacity boundary around 300M parameters:

Below 200M (e.g., Maaza-MLM-135M): - Excellent on simple schemas (44.7%) - Usable on medium schemas (13.5%) - Zero capability on complex schemas (0.0%)

Above 300M (e.g., Maaza-SLM-360M): - Strong on simple schemas (78.9%) - Reliable on medium schemas (51.4%) - First non-zero on complex schemas (4.0%)

This threshold appears earlier than in traditional benchmarks (MMLU, HellaSwag), where model capabilities scale more gradually. We hypothesize that structured extraction requires explicit memory capacity for: - Tracking multiple field dependencies - Maintaining nested object structures - Coordinating array elements with parent objects

Future work should investigate whether architectural changes (e.g., attention mechanisms, memory augmentation) can lower this threshold.

6.3 Comparison to Related Work

6.3.1 SLM-Bench (Pham et al., 2025)

Pham et al.'s SLM-Bench evaluates small models on general NLP tasks (MMLU, GSM8K, HellaSwag). Our EdgeJSON complements their work by focusing on **structured output reliability**-a gap in existing benchmarks. Key differences:

Aspect	SLM-Bench (Pham)	EdgeJSON (Ours)
Focus	Reasoning, knowledge	Structured extraction
Metrics	Accuracy, perplexity	JSONExact, Field F1
Task Type	Multiple-choice, QA	JSON generation
Use Case	General capability	Edge AI deployment

Both benchmarks are valuable: SLM-Bench for broad capability assessment, EdgeJSON for deployment-specific validation.

6.3.2 Code Generation Models

Our approach shares similarities with code generation models (CodeLlama, StarCoder) that are fine-tuned on structured output (code). However, JSON extraction differs in key ways: - Shorter outputs: JSON objects are typically 50-200 tokens vs. 500+ for code - Stricter constraints: Invalid JSON fails completely; syntax errors in code are debuggable - Domain-specific: JSON schemas vary widely across applications

6.4 Limitations

6.4.1 Dataset

Synthetic Data: While validated for mathematical consistency, our synthetic data may not capture: - Linguistic diversity of real-world prompts - Edge cases and adversarial inputs - Domain-specific terminology

English-Only: EdgeJSON includes only English prompts. Multilingual evaluation remains future work.

Scale: 787 examples is moderate. Larger datasets (10K+ examples) may reveal different scaling behaviors.

6.4.2 Baseline Selection

We evaluated one zero-shot baseline (Qwen2.5-0.5B) to demonstrate the core finding that task specialization can outperform parameter scaling. Additional baselines would strengthen claims: - Instruction-tuned models: Qwen2.5-0.5B-Instruct, Llama-3.2-1B-Instruct - Larger SLMs: Phi-3-mini (3.8B), Gemma-2B - Domain-specific models: Other JSON-tuned or function-calling models

We defer these comparisons to future work for two reasons: (1) our focus on edge-deployable models (<1GB) justifies the current scope, and (2) the core insight-that fine-tuning dramatically shifts the performance curve-is already established with the current baselines. Expanding to instruction-tuned and larger models would test generalization but not alter the fundamental conclusion.

6.4.3 Task Scope

EdgeJSON measures single-turn extraction. Real-world applications may require: - Multi-turn clarification: Model asks for missing information - Error recovery: Model detects and corrects invalid outputs - Partial extraction: Model handles incomplete or noisy inputs

Future benchmarks should address these scenarios.

6.5 Practical Implications

6.5.1 For Practitioners

Our results suggest a deployment strategy: 1. Prototype with micro models: Test Maaza-MLM-135M (270MB) for simple schemas 2. Scale to small models: Use Maaza-SLM-360M (720MB) for complex schemas 3. Fine-tune on domain data: Task-specific training can provide notable benefits

This "edge-first" approach balances cost, latency, and privacy.

6.6 Broader Impact

Positive: - **Accessibility**: Low-resource organizations can deploy AI without API costs - **Sustainability**: Smaller models reduce energy consumption

Considerations: - **Automation**: Automation of data entry, analysis, and/or categorization - **Data quality**: Synthetic training data may not reflect real-world diversity

7. Conclusion

We introduced Maaza, a family of task-specialized micro and small language models for structured JSON extraction, and EdgeJSON, a benchmark for evaluating structured output

reliability on edge devices. Our core finding challenges the "bigger is better" assumption in language modeling: fine-tuned 135M-parameter models outperform zero-shot 500M-parameter models on structured extraction tasks.

7.1 Key Contributions

1. EdgeJSON Benchmark

We released a 787-example dataset spanning 24 real-world schemas, with validated synthetic data and open-source evaluation harness. EdgeJSON fills a critical gap in existing benchmarks by measuring **structural correctness** rather than linguistic fluency.

2. Maaza Model Family

We fine-tuned and released two Apache 2.0-licensed models: - Maaza-MLM-135M: 24.7% JSONExact (13× improvement over base) - Maaza-SLM-360M: 55.1% JSONExact (11× improvement over base)

Both models require <2 minutes training time and run efficiently on CPU-only devices.

3. Empirical Evidence for Task Specialization

Our experiments demonstrate that: - Fine-tuned 135M models beat zero-shot 500M models (24.7% vs. 14.6%) - Task specialization provides greater gains than parameter scaling - A capacity threshold exists around 300M parameters for complex schemas

4. Open Methodology

All models, datasets, training scripts, and evaluation code are publicly available, enabling full reproducibility and community extension.

7.2 Implications

For **practitioners**, our work enables a new deployment paradigm: edge-first inference with cloud fallback. Applications that previously required expensive API calls or infeasible large models can now run locally on modest hardware.

For **researchers**, we demonstrate that structured extraction tasks exhibit different scaling behaviors than reasoning tasks, motivating further investigation of task-specific architectures and capacity thresholds.

For **the field**, EdgeJSON provides a practical benchmark for evaluating models on deployment-relevant tasks, complementing existing academic benchmarks.

7.3 Future Work

We identify four priority directions:

1. Nano Language Models (NLMs)

Can <10M parameter models handle ultra-simple schemas (2-3 fields)? This would enable browser-based and mobile inference without backend infrastructure.

2. Multi-Task Adaptation

Can a single model handle multiple schema types via task prefixes or adapter swapping? This would reduce deployment complexity for applications with diverse extraction needs.

3. Real-World Evaluation

Validate Maaza models on production datasets from industry partners (healthcare records, financial transactions, IoT logs) to assess generalization beyond synthetic data.

4. Cross-Lingual Transfer

Extend EdgeJSON to multilingual scenarios and evaluate whether fine-tuning in one language transfers to others.

7.4 Closing Remarks

The edge AI landscape demands models that are small, fast, and reliable. Our work demonstrates that task-specialized micro models can meet these requirements while outperforming larger general-purpose alternatives. As language models continue to expand to trillions of parameters, we argue that **focused specialization** remains a viable-and often superior-alternative to unbounded scaling.

By releasing Maaza and EdgeJSON under open licenses, we hope to accelerate research and deployment of efficient, practical AI systems that run where data lives: on the edge.

Figures and Tables

Figure 1: Overall Performance Comparison

Performance comparison of Maaza models against baselines on EdgeJSON v3 (158 test examples)

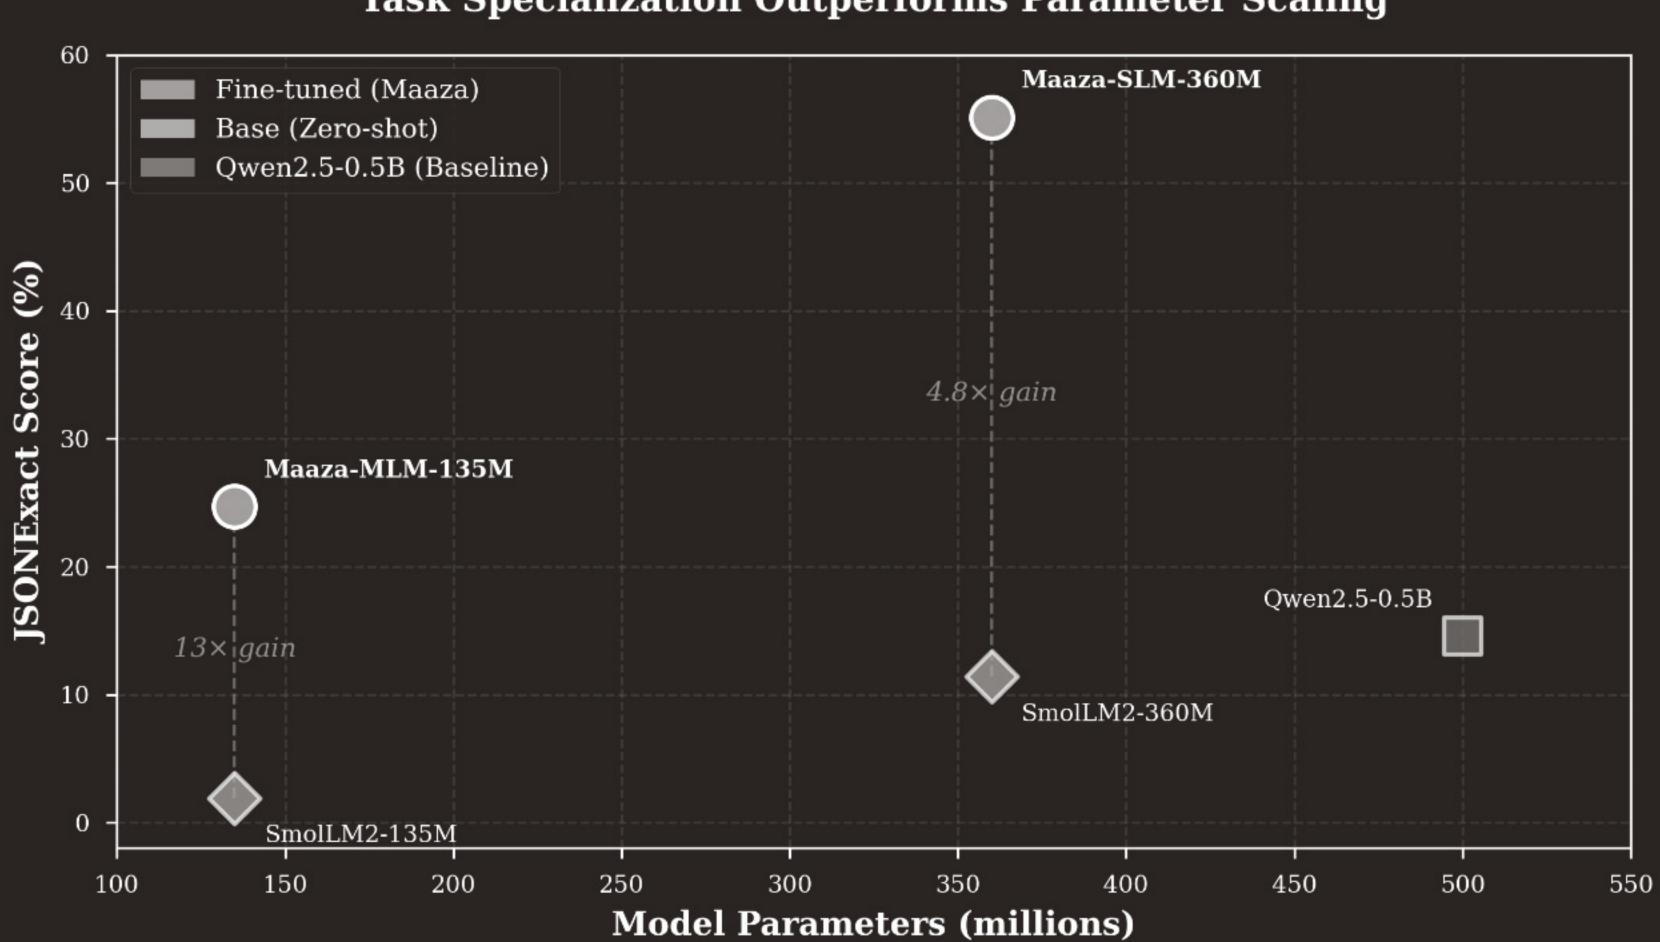


Figure 1: Performance vs Model Size
Task Specialization Outperforms Parameter Scaling

Figure 1: Performance vs Model Size

Figure 1: Task-specialized fine-tuning (blue circles) enables smaller models to outperform larger zero-shot models. Fine-tuned Maaza-MLM-135M (135M params, 24.7%) surpasses zero-shot Qwen2.5-0.5B (500M params, 14.6%) despite being 3.7× smaller.

Key Insight: Fine-tuning shifts performance curve upward by 4.8-13×, enabling smaller models to outperform larger zero-shot models.

Figure 2: Performance by Complexity Level

Breakdown of JSONExact scores by schema complexity

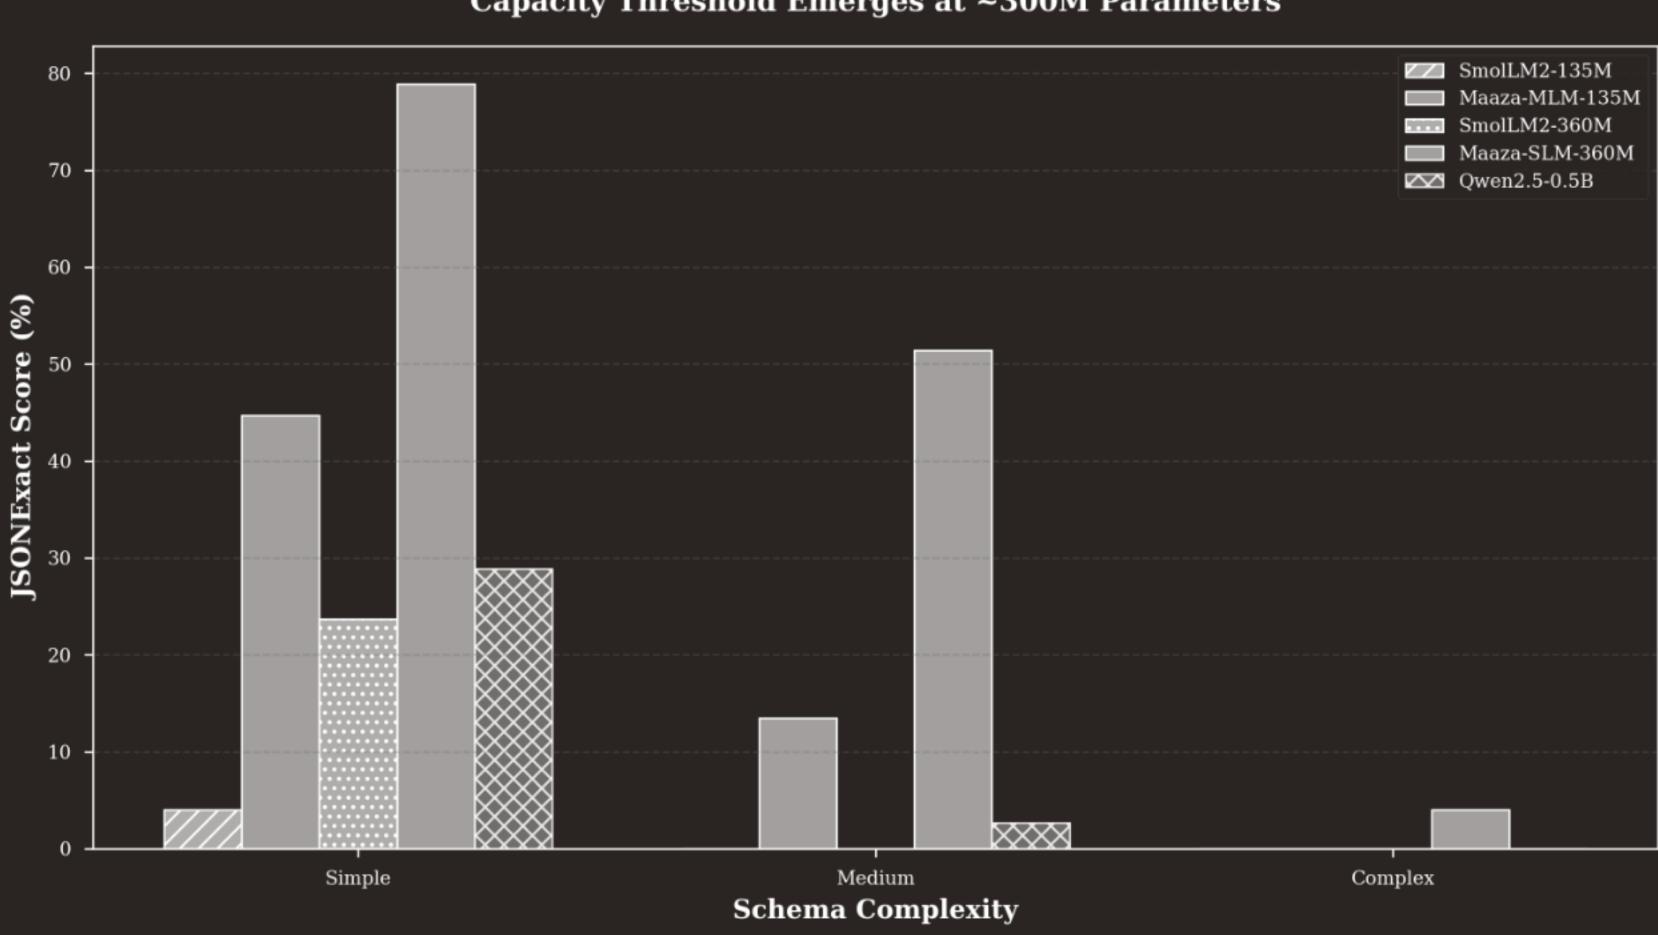


Figure 2: Performance by Schema Complexity Capacity Threshold Emerges at ~300M Parameters

Figure 2: Performance by Complexity

Figure 2: Performance stratifies by schema complexity. All models handle simple schemas well, but only Maaza-SLM-360M achieves non-zero performance on complex schemas (4.0%), revealing a capacity threshold around 300M parameters.

Capacity Threshold: Only models ≥300M parameters achieve non-zero performance on complex schemas, even with fine-tuning.

Table 1: Main Results Summary

				Field			
Model	Params	Туре	JSONExact	F1	Compliance	Size	Training
SmolLM2- 135M (base)	135M	Zero- shot	1.9%	0.024	5.1%	270MB	-
Maaza-MLM- 135M	135M	Fine- tuned	24.7%	0.520	51.9%	270MB	Rapid
SmolLM2- 360M (base)	360M	Zero- shot	11.4%	0.240	23.7%	720MB	-
Maaza-SLM- 360M	360M	Fine- tuned	55.1%	0.780	79.7%	720MB	Rapid
Qwen2.5-0.5B	500M	Zero- shot	14.6%	0.195	19.0%	954MB	-

Key Comparisons: - Maaza-MLM-135M vs Qwen-0.5B: $1.7 \times$ better performance, $3.5 \times$ smaller size - Maaza-SLM-360M vs Qwen-0.5B: $3.8 \times$ better performance, $1.3 \times$ smaller size - Fine-tuning gain (135M): $13 \times$ improvement ($1.9\% \rightarrow 24.7\%$) - Fine-tuning gain (360M): $4.8 \times$ improvement ($11.4\% \rightarrow 55.1\%$)

Table 2: Performance by Complexity Breakdown

Model	Simple (76 ex)	Medium (57 ex)	Complex (25 ex)
SmolLM2-135M (base)	4.0% / 0.055	0.0% / 0.004	0.0% / 0.000
Maaza-MLM-135M	44.7% / 0.715	13.5% / 0.399	0.0% / 0.183
SmolLM2-360M (base)	23.7% / 0.240	0.0% / 0.004	0.0% / 0.000
Maaza-SLM-360M	78.9% / 0.910	51.4% / 0.740	4.0% / 0.352
Qwen2.5-0.5B	28.9% / 0.392	2.7% / 0.027	0.0% / 0.000

Format: JSONExact / Field F1

Acknowledgments

We thank the HuggingFace team for the SmolLM2 base models and the Qwen team for making their models publicly available. We also thank the open-source community for evaluation tools and libraries.

References

Abdin, M., et al. (2024). Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone. arXiv preprint arXiv:2404.14219.

Allal, L. B., Wolf, T., et al. (2024). SmolLM2: When Smol Goes Big - Data-Centric Training of Small Language Models. arXiv preprint arXiv:2502.02737.

Brown, T. B., Mann, B., Ryder, N., et al. (2020). Language Models are Few-Shot Learners. Advances in Neural Information Processing Systems.

Chen, M., Tworek, J., et al. (2021). Evaluating Large Language Models Trained on Code. arXiv preprint arXiv:2107.03374.

Cobbe, K., Kosaraju, V., Bavarian, M., et al. (2021). Training Verifiers to Solve Math Word Problems. arXiv preprint arXiv:2110.14168.

Dettmers, T., Pagnoni, A., Holtzman, A., & Zettlemoyer, L. (2023). QLoRA: Efficient Finetuning of Quantized LLMs. *NeurIPS*.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL-HLT*.

Gemma Team. (2024). Gemma 2: Improving Open Language Models at a Practical Size. arXiv preprint arXiv:2408.00118.

Han, S., Pool, J., Tran, J., & Dally, W. (2016). Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. *ICLR*.

Hendrycks, D., Burns, C., Basart, S., et al. (2021). Measuring Massive Multitask Language Understanding. *ICLR*.

Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the Knowledge in a Neural Network. *NeurIPS Deep Learning Workshop*.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2021). LoRA: Low-Rank Adaptation of Large Language Models. arXiv preprint arXiv:2106.09685.

Hugging Face. (2024). Transformers.js v3: WebGPU Support, New Models & Demos. *Hugging Face Blog*. https://huggingface.co/blog/transformersjs-v3

Hui, B., et al. (2024). Qwen2.5-Coder: A Family of Open Code LLMs. arXiv preprint arXiv:2409.12186.

Jacob, B., Kligys, S., Chen, B., et al. (2018). Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference. *CVPR*.

Jiao, X., Yin, Y., Shang, L., Jiang, X., Chen, X., Li, L., Wang, F., & Liu, Q. (2020). TinyBERT: Distilling BERT for Natural Language Understanding. *Findings of EMNLP*.

Meta AI. (2024). Llama 3.2: Revolutionizing Edge AI and Vision with Open Models. *Meta AI Blog*. https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/

Microsoft Edge Team. (2025). Opening On-Device AI Models to Web Apps in Edge. *Microsoft Edge Blog*. https://www.theverge.com/news/669528/microsoft-ai-edge-browser-web-app-build-apis

ONNX Runtime Team. (2018). ONNX Runtime: High-Performance Inference Engine. https://onnxruntime.ai

Pham, N. T., et al. (2025). SLM-Bench: A Comprehensive Benchmark of Small Language Models on Environmental and Efficiency Impacts. *Findings of EMNLP*.

Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. *NeurIPS Workshop on Energy Efficient Machine Learning and Cognitive Computing*.

Sun, Z., Yu, H., Song, X., Liu, R., Yang, Y., & Zhou, D. (2020). MobileBERT: A Compact Task-Agnostic BERT for Resource-Limited Devices. *Proceedings of ACL*.

TensorFlow Team. (2017). TensorFlow Lite. https://www.tensorflow.org/lite

Touvron, H., Lavril, T., Izacard, G., et al. (2023). LLaMA: Open and Efficient Foundation Language Models. arXiv preprint arXiv:2302.13971.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention Is All You Need. *Advances in Neural Information Processing Systems*.

Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., & Bowman, S. R. (2018). GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *ICLR*.

Wang, A., Pruksachatkun, Y., Nangia, N., et al. (2019). SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. *NeurIPS*.

Wang, R., et al. (2025). A Survey of Edge Efficient LLMs and Techniques. *Journal of Systems Architecture*.

Yang, A., Yang, B., et al. (2024). Qwen2.5 Technical Report. arXiv preprint arXiv:2412.15115.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., & Choi, Y. (2019). HellaSwag: Can a Machine Really Finish Your Sentence? *ACL*.

Zeng, Y., Chen, T., et al. (2024). WebLLM: A High-Performance In-Browser LLM Inference Engine. arXiv preprint arXiv:2412.15803.

Zhang, J., Xu, S., et al. (2024). TinyLlama: An Open-Source Small Language Model. arXiv preprint arXiv:2401.02385.

Zheng, Y., Chen, Y., Qian, B., Shi, X., Shu, Y., & Chen, J. (2024). A Review on Edge Large Language Models: Design, Execution, and Applications. arXiv preprint arXiv:2410.11845.