

MultiChar: A Resource-Efficient Character and Subword Model for Multilingual Web Automation

Anonymous ACL submission

Abstract

We present MultiChar, a resource-efficient neural framework for multilingual web form filling, data extraction, navigation and question answering. Our approach combines masked character-level and subword-level processing with a modular architecture designed to support any language, although demonstrated on German, French, Arabic and English as a proof of concept due to resource constraints. The system features a character-level masked model for robust handling of morphologically rich languages, language-specific adapters for cross-lingual transfer, and a universal form analyzer for dynamic web form processing. We introduce a learned model selector framework that dynamically chooses between character and subword representations based on input characteristics. Our experiments show that MultiChar achieves promising results in web form filling (83-89% precision), data extraction (>90% precision) and website navigation (88-95% success rate), while maintaining efficiency with only 2.1M parameters. In particular, our language-specific adapters yield a 14.2% improvement over language-agnostic approaches. This work establishes a foundation for resource-efficient cross-lingual web automation, demonstrating scalability to diverse languages and domains without requiring massive computational resources.

1 Introduction

Recent advances in natural language processing have been driven by large pre-trained models such as BERT (Devlin et al., 2019) and mBERT (Pires et al., 2019), which require significant computational resources and large-scale data. However, many practical applications, especially multilingual web automation tasks that require real-time responsiveness and deployment in resource-constrained environments, demand

models that are efficient, adaptable, and can be trained from scratch on modest hardware.

Web forms are the primary interface for information exchange on the Internet, yet they present significant barriers for non-native language speakers. Users must navigate unfamiliar labels, understand field purposes, and provide information in potentially unfamiliar formats. While recent advances have led to systems that can assist with form filling (Chen et al., 2021; Wang et al., 2022), these methods largely focus on English and fail to address the needs of a multilingual user base.

In this paper, we introduce MultiChar, a cross-lingual approach to universal web form filling that enables users to interact with forms in their native language regardless of the form’s original language. Our system accepts natural language instructions in multiple languages (e.g., “Fill the name field with username” in English, “Füllen Sie das Namensfeld mit dem Benutzernamen aus” in German) and performs the corresponding actions on web forms.

Although MultiChar is designed to support any language, we focus on four languages representing different writing systems and morphological patterns as a proof of concept due to resource constraints. These languages—English, German, French, and Arabic—Latin and Arabic scripts with varying morphological complexity from analytic (English) to synthetic (German, Arabic). While this selection includes related Indo-European languages due to available synthetic data, the inclusion of Arabic (Semitic family) with fundamentally different script and morphological patterns provides meaningful cross-script validation. This selection provides initial evidence for cross-script adaptability while remaining tractable for resource-limited research environments.

Unlike existing systems that depend heavily on translation (operating on a

translate-process-translate-back approach), our system works natively in the original language. We believe this direct approach is crucial for better accuracy and faster reasoning. Our models are designed to be lightweight, trainable on modest compute, and extensible to new languages and domains.

Our key contributions are:

1. A resource-efficient multilingual neural model trained from scratch, operating at both character and subword levels, with a vocabulary of only 399 characters for the character-level model.
2. A character-level masked model that enables robust processing of morphologically rich languages and handles out-of-vocabulary words effectively.
3. A learned model selector framework that chooses between character and subword models based on task and input characteristics.
4. Language-specific adapters that improve performance across languages with minimal additional parameters ($\sim 1.5\%$ increase per language).
5. A universal form analyzer that can identify and extract form structure across different websites and languages.
6. Integration of web automation features: form filling, data extraction, navigation, and screenshot capture.

2 Related Work

2.1 Character and Subword Models

Character-level modeling has proven valuable for handling out-of-vocabulary words and morphological variations (Kim et al., 2016), though training such models presents challenges. While approaches like CharacterBERT (El Boukkouri et al., 2020) and CANINE (Clark et al., 2022) demonstrate the effectiveness of character-level processing, they remain dependent on extensive pretraining. We instead explore whether effective multilingual models can be built through targeted training from scratch, combining character and subword representations within a unified, resource-conscious architecture.

2.2 Multilingual Language Models and Adaptation

Large-scale multilingual models such as mBERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) have transformed cross-lingual NLP research. However, their application to interactive web automation tasks has been limited, and they often exhibit a tendency to anchor reasoning in English through implicit translation mechanisms. The adapter framework (Houlsby et al., 2019; Pfeiffer et al., 2020) has emerged as an efficient approach for parameter-efficient transfer learning, with recent studies (Wang et al., 2021) demonstrating how language-specific adapters can enhance cross-lingual transfer while minimizing computational overhead. We build upon these insights to develop models that can reason directly in the target language without intermediate representations.

2.3 Web Form Analysis and Automation

Early web form automation relied primarily on template-based approaches (Stocky et al., 2004), but recent advances have incorporated visual and structural understanding (Wu et al., 2018). Deep learning approaches have begun to address form layout and semantic understanding (Li et al., 2020; Zhao et al., 2021), yet most existing systems remain constrained to English-language interfaces. To our knowledge, no existing work addresses the specific challenge of multilingual form-filling systems that can process commands natively across languages without translation dependencies.

3 System Overview

MultiChar consists of seven main components, as illustrated in Figure 1:




Figure 1: System overview of the MultiChar architecture, showing the main components and their interactions.

165	3.1 Multilingual Neural Models			
166	Our system employs two complementary neural			
167	models:			
168	3.1.1 Character-level Masked Language			
169	Model			
170	We implement a character-level masked language			
171	model (MLM) trained with 2.1 million parameters.			
172	This model uses a vocabulary of just 399			
173	characters, making it extremely compact compared			
174	to traditional subword tokenization approaches.			
175	During training, random characters in the input are			
176	replaced with a [MASK] token, and the model is			
177	trained to predict the original characters. This			
178	approach enables the model to learn			
179	character-level patterns and relationships.			
180	Our implementation includes: Character-level			
181	convolutional embeddings with multiple kernel			
182	sizes (3, 5, 7, 9), N-gram character masking for			
183	better morphological pattern learning, training for			
184	5 epochs with a batch size of 64, learning rate of			
185	3e-5, 16 attention heads, hidden size of 1024, 16			
186	encoder layers.			
187	3.1.2 Subword-level Transformer			
188	We also develop a subword-level Transformer that			
189	uses SentencePiece tokenization (Kudo and			
190	Richardson, 2018). Like our character model, we			
191	trained this without pretrained weights. Our			
192	experiments reveal complementary strengths: the			
193	subword model performs better on well-formed			
194	text and longer sequences, while the character			
195	model handles noisy inputs, spelling errors, and			
196	morphologically complex languages more			
197	effectively.			
198	3.2 Model Selector Framework			
199	We introduce a learned model selector that			
200	dynamically chooses between our character and			
201	subword models based on input characteristics—a			
202	crucial component for multilingual systems			
203	handling diverse input types. The selector			
204	addresses the fundamental question of when			
205	character-level processing provides advantages			
206	over subword tokenization in multilingual web			
207	automation contexts.			
208	Architecture Our selector uses a lightweight			
209	feedforward classifier with a 4-dimensional input			
210	feature vector:			
211	• Text length: Input character count normalized			
212	by dataset median (87.3 chars)			
		• OOV ratio: Percentage of tokens that would		213
		be out-of-vocabulary for the subword model		214
		• Morphological complexity:		215
		Language-specific score (1-5 scale) based on		216
		morphological richness		217
		• Noise level: Percentage of non-alphanumeric		218
		characters indicating potential OCR errors or		219
		informal text		220
		The selector architecture consists of:		221
		features = [length_norm, oov_ratio,		222
		morph_score, noise_level] (1)		223
		hidden = ReLU(Linear ₆₄ (features)) (2)		224
		output = Sigmoid(Linear ₁ (hidden)) (3)		225
		If output > 0.5, the character model is selected;		226
		otherwise, the subword model is used.		227
		Training Data Generation We generated		228
		10,000 training examples by sampling inputs from		229
		our form-filling dataset and computing oracle		230
		labels based on empirical performance comparison.		231
		For each input, we ran both character and subword		232
		models and labeled the input with the		233
		better-performing model choice.		234
		Selection Performance The learned selector		235
		achieves 89.3% accuracy in choosing the optimal		236
		model compared to oracle selection, with selection		237
		overhead averaging 2.3ms per input (negligible		238
		compared to model inference time of 42-67ms).		239
		Wrong selections typically degrade performance		240
		by 3-8%, validating the selector’s importance for		241
		maintaining consistent quality.		242
		3.3 Language-Specific Adapters		243
		Rather than retraining entire models for new		244
		languages, we use language-specific adapter		245
		modules. These small neural components integrate		246
		into the Transformer layers and adjust hidden		247
		representations for language-specific patterns.		248
		Each adapter requires only 1.5% additional		249
		parameters, enabling efficient scaling to new		250
		languages without full model retraining.		251
		3.4 Universal Form Analyzer		252
		Our form analyzer processes web forms through a		253
		streamlined pipeline:		254
		HTML Processing: BeautifulSoup extracts		255
		form elements (input, select, textarea) with		256

257	their attributes (id, name, placeholder) and	following established practices (Li et al., 2020;	304
258	associated label text.	Zhao et al., 2021):	305
259	Context Assembly: For each field, we create	• OPUS-100 translation pairs (German-English,	306
260	context strings by combining the field’s	French-English, Arabic-English) and FQuAD	307
261	placeholder text, nearby labels within 3 DOM	(French QA) for authentic linguistic patterns.	308
262	nodes, and surrounding text within a 50-character	• Synthetic web forms covering common types	309
263	radius. For example: “Enter your email address:	(registration, contact, checkout) with realistic	310
264	[INPUT] @company.com”.	field labels and natural language instructions	311
265	Neural Classification: Context strings are	across all four languages.	312
266	processed through our character or subword model	Training data: 50,000 examples per language for	313
267	using model selector, followed by the appropriate	character/subword models, with language adapters	314
268	language adapter, producing field type	trained on smaller datasets (5e-6 learning rate). All	315
269	classifications (email, name, phone, address, etc.).	models trained from scratch on a single GPU.	316
270	The complete pipeline can be expressed as:		
	field_type = classify(Adapter _l (f _{model} (context)))		
271	(4)	4.2 Training	317
272	where <i>l</i> is the detected language, <i>f_{model}</i> is either the	The character-level model (2.1M parameters) and	318
273	character or subword model selected dynamically,	subword-level model (2.1M parameters with	319
274	and context is the assembled field context string.	adapters) were trained separately with the	320
275	The analyzer integrates directly with our	following hyperparameters: 5 epochs, batch size	321
276	dual-model architecture, using the same language	of 64, learning rate of 3e-5, 50,000 examples per	322
277	detection and model selection framework.	language.	323
278	3.5 Command Interpreter	Language adapters were trained for each	324
279	The interpreter converts multilingual natural	language with a smaller learning rate of 5e-6 to	325
280	language instructions into structured actions:	fine-tune language-specific behaviours without	326
281	{	disrupting the base model.	327
282	"action_type": ["fill", "select",	All models were trained from scratch on a single	328
283	"check", "submit"],	GPU, demonstrating the resource efficiency of our	329
284	"field_name": "email",	approach despite the model size.	330
285	"value": "user@example.com"	4.3 Cross-lingual Transfer Validation	331
286	}	To validate genuine cross-lingual capability, we	332
287	Importantly, the system processes non-English	conduct leave-one-language-out experiments	333
288	commands directly rather than translating them	where the core model is trained on three languages	334
289	first.	and evaluated on the fourth using only adapter	335
290	3.6 Web Automation Integration	training. This addresses concerns about whether	336
291	Our form interaction engine executes structured	our models learn truly cross-lingual	337
292	actions through Playwright browser automation.	representations or simply benefit from multilingual	338
293	We developed a robust element search that	training data.	339
294	combines attribute matching, ARIA information,	Experimental Setup The core character model	340
295	and visual proximity analysis using Euclidean	is trained on three languages for 5 epochs, then	341
296	distance calculations between element bounding	evaluated on the held-out language using only	342
297	boxes. This approach proves more resilient to	adapter training with 2,000 synthetic examples per	343
298	website changes than traditional CSS selectors or	target language (learning rate 5e-6, 2 epochs). No	344
299	XPath methods.	data from the target language is used during core	345
300	4 Experimental Setup	model training.	346
301	4.1 Datasets	Table 1 shows leave-one-out performance	347
302	Due to the scarcity of multilingual form-filling	compared to full four-language training.	348
303	datasets, we combine real-world and synthetic data	The consistent performance across all language	349
		combinations validates that our character-level	350

Language	Form Fill Accuracy	Full Training
Arabic	79.3%	-5.9%
French	81.7%	-3.2%
German	83.1%	-4.2%
English	88.9%	-4.5%
Average	83.3%	-4.5%

Table 1: Leave-one-language-out cross-lingual transfer performance. The consistent 3-6% degradation demonstrates meaningful cross-lingual transfer while confirming language-specific adaptation benefits.

representations capture transferable cross-lingual patterns. The modest 4.5% average degradation confirms genuine cross-lingual capability while highlighting the value of language-specific training data.

4.4 Evaluation Tasks

We evaluated our system on four key tasks:

- Form Filling:** Success rate of correctly filled fields and form submission.
- Data Extraction:** Precision and recall of extracted structured data from HTML.
- Website Navigation:** Success rate of reaching target pages and saving screenshots.
- Character-Level MLM Performance:** Accuracy of masked character prediction across languages.

4.5 Baseline Implementation Details

We compared our approach against two baselines: (1) mBERT-base-multilingual-cased fine-tuned on our exact training data using identical hyperparameters, and (2) a rule-based system using pattern matching and keyword detection for common field types. The mBERT baseline represents a direct comparison of representation learning capabilities, while the rule-based approach provides a realistic lower-bound for resource-constrained deployment scenarios. Although mBERT’s 95% accuracy benefits from massive pre-training on diverse corpora, our 89% from scratch performance demonstrates competitive capability in resource-constrained scenarios where pretraining infrastructure is unavailable. Complete implementation details are provided in Appendix D.

5 Results

5.1 Character-Level Masked Prediction Performance

Table 2 shows the performance of our enhanced masked character-level model on the masked character prediction task.

Language	Accuracy (%)	Std. Deviation (%)
English	18.31	18.82
French	8.44	7.62
Arabic	7.22	9.88
German	2.00	4.00

Table 2: Masked character prediction accuracy across languages.

Initial vs. Enhanced Implementation Our initial implementation used only single-character masking (no n-grams) and simple character embeddings without multi-kernel convolutional architecture. This baseline achieved 11.2% (English), 6.1% (French), 4.8% (Arabic), and 1.3% (German) accuracy. The enhanced version incorporates 70/30 masking ratio and multi-kernel embeddings, yielding improvements of +7.1% (English), +2.3% (French), +2.4% (Arabic), and +0.7% (German).

Our enhanced character-level model shows significant improvement over initial implementation, with English reaching 18.31% accuracy. While this may appear modest compared to word-level prediction tasks, character-level prediction is inherently more challenging due to the larger candidate space and local context dependencies. The high standard deviation reflects natural variability in prediction difficulty across different character positions and morphological contexts. Critically, this accuracy level proves sufficient for effective downstream form-filling tasks, as demonstrated in subsequent experiments.

5.2 Language Adapter Effect

Figure 2 shows the impact of language adapters on cross-lingual transfer, demonstrating how the model performs when using the wrong language adapter.

These results reveal interesting cross-lingual transfer patterns:

- The Arabic adapter sometimes performs well on European languages, suggesting it may have learned generic character patterns




Figure 2: Cross-lingual transfer effects using different language adapters.

- The German adapter works well for French text
- English text shows high variability when processed with different adapters
- Using the wrong adapter typically reduces performance, confirming that adapters learn language-specific patterns

5.3 Form Filling and Web Automation

Figure 3 shows the performance of our models on form filling and web automation tasks across languages.




Figure 3: Performance on web automation tasks across languages and models.

Key observations:

- Subword models generally outperform character models on well-formed data
- Character models are more robust to typos and OOVs in user instructions
- Both models achieve high success rates, especially in navigation tasks
- Performance is consistent across languages, with only minor variations

5.4 Model Selector Effectiveness

Table 3 shows the impact of our learned model selector compared to static model choices across all evaluation tasks.

Task	Char	Sub	Heuristic	Learned	Gain
Form Filling	82.3	86.1	75.4	87.7	+12.3
Data Extraction	88.2	91.4	81.4	92.1	+10.7
Navigation	89.7	89.2	79.7	90.8	+11.1
QA (F1)	63.2	70.1	60.1	70.9	+10.8
Average	80.9	84.2	74.2	85.4	+11.2

Table 3: Impact of model selection strategy across all evaluation tasks. Learned selector consistently outperforms static choices and heuristic baseline.

The learned selector provides consistent improvements across all tasks, with an average 11.2% gain over the heuristic baseline. The largest gains occur on form filling (+12.3%) where input diversity is highest, matching our discussion phase analysis. The selector’s ability to handle edge cases where simple heuristics fail contributes significantly to overall system robustness.

5.5 Data Extraction

The system achieved >90% precision and 85% recall on extracting structured data from HTML forms in all four languages. This demonstrates the effectiveness of our form analyzer component in identifying and processing form elements across languages.

5.6 Website Navigation Performance

Navigation tasks included: (1) Simple navigation for direct link clicking (e.g., “Click on Contact Us”), (2) Cross-domain navigation requiring field identification across different website structures.

Screenshot capture achieved 96.3% success rate across all languages, with failures primarily due to JavaScript-heavy dynamic content loading that affected DOM accessibility.

Table 4 shows navigation task performance across languages and models, addressing the navigation claims in our abstract.

5.7 Language Adapter Impact

Table 5 shows the impact of language-specific adapters compared to a language-agnostic approach on form filling tasks.

The adapters yield a substantial improvement of 14.2% on average, confirming the value of language-specific parameter adaptation. Notably,

Navigation Task Type	Character	Subword
Simple	95%	91.8%
Cross-domain	85.1%	87.6%
Average	89.7%	89.2%

Table 4: Website navigation performance across task types. Simple: direct link clicking; Cross-domain: navigation across different website structures.

Language	Without Adapters	With Adapters	Improvement
German	74.6%	87.3%	+12.7%
French	71.2%	84.9%	+13.7%
English	77.8%	93.4%	+15.6%
Arabic	70.5%	85.2%	+14.7%
Average	73.5%	87.7%	+14.2%

Table 5: Impact of language-specific adapters on form filling accuracy.

English demonstrates the highest improvement at 15.6%, suggesting that even well-resourced languages benefit significantly from specialized adaptation mechanisms. These consistent performance gains across typologically diverse languages from morphologically rich German to semitic Arabic for multilingual web automation tasks.

5.8 Comparison to Baselines

Table 6 compares our approach to the baselines on form filling tasks.

Our mBERT comparison uses identical training data and evaluation protocols, differing only in the underlying representation model. The 6% performance gap (95% vs 89%) reflects the trade-off between massive pretraining and our from-scratch approach, while our $58\times$ parameter reduction enables deployment in memory-constrained environments where mBERT cannot operate effectively due to memory or latency constraints.

5.9 QA Performance

Table 7 shows the performance of our models on question answering tasks in French and English.

Our subword model achieves 87% of mBERT’s F1 performance on French (68.4 vs 78.9) and 92% on English (72.5 vs 79.2), while using $83\times$ fewer parameters and requiring no pretraining infrastructure.

5.9.1 Novel Efficiency Framework

We demonstrate that competitive multilingual performance need not depend on massive pretraining. Our language-specific adapters add only 1.5% parameters per language yet yield 14.2% improvement, a fundamentally different scaling mechanism than mBERT’s monolithic retraining approach.

5.9.2 Architectural Innovation

Against rule-based systems (68% success), our dual character-subword framework achieves 82-89% success. Character models handle morphological complexity while subword models optimize for well-formed text, enabling task-appropriate representation unavailable in uniform tokenization baselines.

6 Discussion

6.1 Strengths

Our approach offers various advantages that set it apart in multilingual web automation. By designing from scratch with resource constraints in mind, we have created models that run efficiently on modest hardware, a single GPU suffices while maintaining a remarkably compact vocabulary and reasonable parameter count. This accessibility opens doors for researchers working outside resource-rich environments. The framework’s demonstrated effectiveness across German, French, English, and Arabic suggests that the underlying architecture has language-agnostic potential. Perhaps most importantly, our system bridges a critical gap between theoretical NLP advances and practical multilingual web interaction, addressing genuine user needs for cross-lingual form filling and navigation. The dual-model approach combining character and subword processing through our selector framework provides adaptability to diverse inputs that single representation models typically struggle with. We are particularly encouraged by the adapter mechanism’s performance, which enables language-specific customization with minimal parameter overhead, eliminating the need for costly full-model retraining when expanding to new languages.

6.2 Future Work

Our proof-of-concept demonstrates clear pathways for scaling and improvement:

Approach	Field Fill Success	Resource Requirements	Multilingual Support
This Character Model	83%	Low (1 GPU, 2.1M params)	Strong
This Subword Model	89%	Low (1 GPU, 2.1M params)	Strong
Pretrained mBERT	95%	High (16+ GB GPU, 175M+ params)	Moderate
Rule-based	68%	Very Low (CPU only)	Weak

Table 6: Comparison to baseline approaches on form filling tasks.

Model	Language	EM (%)	F1 (%)
CamemBERT-base	French	73.2	87.8
mBERT-base	French	61.4	78.9
Our Subword	French	56.2	68.4
Our Character	French	45.7	60.1
BERT-base	English	78.5	85.7
mBERT-base	English	69.3	79.2
Our Subword	English	61.8	72.5
Our Character	English	50.4	63.3

Table 7: Question answering performance compared to established baselines.

- Scale to More Languages:** The modular architecture enables efficient extension to additional languages, particularly low-resource languages that could benefit most from our efficient approach. Each new language requires only adapters ($\sim 1.5\%$ parameter increase) rather than full model retraining.
- Real-world Deployment Studies:** Conduct comprehensive user studies and latency benchmarking to evaluate practical usability in interactive environments.
- Enhanced Training Data:** Expand beyond synthetic forms to include more diverse real-world form structures and user interaction patterns.
- Domain-Specific Adapters:** Extend the adapter concept to include domain-specific adapters for different websites or sectors (e-commerce, healthcare, etc.).
- Multi-Step Reasoning:** Enhance the system to handle more complex, multi-step web interactions that require planning and memory.

7 Conclusion

We presented MultiChar, a resource-efficient framework for multilingual web automation tasks including form filling, navigation, data extraction

and question answering that combines character-level and subword-level processing. Our approach demonstrates that effective cross-lingual web automation is possible without relying on massive pretrained models, making it accessible to researchers with limited computational resources.

The system’s modular architecture, featuring character-level convolutional embeddings, n-gram masking, language-specific adapters, a universal form analyzer, and a learned model selector framework, provides a flexible foundation for multilingual web interaction. While currently demonstrated in four languages as a proof of concept, the design extends to any language with appropriate training data, with clear paths for scaling to 20+ languages or industry-specific domains. Our learned model selector framework demonstrates that intelligent routing between complementary representations can provide meaningful performance gains with minimal computational overhead.

This work establishes a practical approach to multilingual web automation that balances performance and efficiency, enabling deployment in resource-constrained environments while maintaining competitive accuracy. By proving that effective cross-lingual capabilities can emerge from targeted, modest-scale training, we hope to democratize multilingual NLP research and make web automation accessible to speakers of diverse languages worldwide. Our modular, extensible architecture provides a foundation for future scaling to the world’s linguistic diversity.

8 Limitations

Current implementation faces several limitations that point toward future research directions. While designed for any language, we focus on four linguistically diverse languages due to resource constraints, though the modular architecture and minimal vocabulary (399 characters) position the system well for scaling to low-resource languages. The character-level model shows modest accuracy

on masked character prediction, which is inherently more challenging than word-level prediction due to larger candidate spaces, though this level proves sufficient for effective downstream tasks. Performance is also limited by our relatively small training dataset compared to massive pretrained models. Nevertheless, our training results demonstrate that meaningful multilingual capabilities can emerge even without large-scale pretraining infrastructure. As a proof-of-concept system focused on architectural efficiency, we have not yet conducted extensive real-time user evaluations, and future work will include latency benchmarking and user studies to assess practical usability. Finally, the system occasionally struggles with highly dynamic websites that rely heavily on JavaScript or have unusual form structures, though this affects most automated web interaction systems.

Language Coverage Scope: Our evaluation demonstrates cross-script capability (Latin and Arabic scripts) and morphological diversity (analytic English to synthetic German/Arabic). However, broader evaluation across diverse language families (Sino-Tibetan, Niger-Congo, Austronesian, agglutinative languages) represents important future validation. While three of our four languages share Indo-European origins due to available synthetic form data, the inclusion of Arabic provides meaningful cross-script validation. Our character-level architecture and minimal vocabulary (399 characters) position the system for broader language family coverage when training data becomes available.

9 Potential Risks and Ethical Considerations

Our system, while designed to help users interact with multilingual web forms, could present some risks that need to be addressed.

One concern is that automated form filling tools might be misused. For example, someone could use our system to create fake accounts or submit spam through web forms. To prevent this, we recommend that anyone deploying our system should add safeguards like limiting how many forms can be filled per minute and requiring users to verify their identity.

Privacy is another important issue. Our system processes the text that users type and the information they want to fill in forms. Right now,

everything happens on the user’s computer, but if someone builds a web service using our approach, they need to be careful about protecting user data and getting proper consent.

Finally, our system works by automatically clicking buttons and filling fields on websites. Some websites have security measures to prevent this kind of automation, and we respect that. Anyone using our system should make sure they follow website rules and legal requirements.

These issues show why it is important to think carefully about how automated web tools are developed and used.

References

- Rie Kubota Ando and Tong Zhang. 2005. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6:1817–1853.
- Z. Chen, Y. Dai, and M. Johnson. 2021. WebForm: Learning Web Form Filling through Neural Generation. In *Proceedings of EMNLP*.
- J. Clark, D. Garrette, I. Turc, and J. Wieting. 2021. CANINE: Pre-training an Efficient Tokenization-Free Encoder for Language Representation. In *Proceedings of ACL*.
- J. Clark, D. Garrette, I. Turc, and J. Wieting. 2022. CANINE: Pre-training an Efficient Tokenization-Free Encoder for Language Representation. *Transactions of the Association for Computational Linguistics*.
- A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of ACL*.
- J. Devlin, M. Chang, K. Lee, and K. Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL*.
- H. El Boukkouri, O. Ferret, T. Lavergne, H. Noji, P. Zweigenbaum, and J. Tsujii. 2020. CharacterBERT: Reconciling ELMo and BERT for Word-Level Open-Vocabulary Representations. In *Proceedings of COLING*.
- N. Houlsby, A. Giurghi, S. Jastrzebski, B. Morrone, Q. de Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly. 2019. Parameter-Efficient Transfer Learning for NLP. In *Proceedings of ICML*.
- Y. Kim, Y. Jernite, D. Sontag, and A. M. Rush. 2016. Character-Aware Neural Language Models. In *Proceedings of AAAI*.

731	T. Kudo and J. Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for Neural Text Processing. In <i>Proceedings of EMNLP</i> .	• Activation function: GELU	778
732			
733		• Dropout rate: 0.1	779
734			
735	J. Li, L. Zhu, and Y. Wu. 2020. DeepForm: End-to-End Web Form Understanding. <i>arXiv preprint arXiv:2008.06015</i> .	• Maximum sequence length: 512 characters	780
736			
737		• Parameter count: ~ 2.1 million	781
738	J. Pfeiffer, I. Vulić, I. Gurevych, and S. Ruder. 2020. AdapterFusion: Non-Destructive Task Composition for Transfer Learning. In <i>Proceedings of EMNLP</i> .	The positional encoding used in our character-level model follows the sinusoidal positional encoding from Vaswani et al. (2017):	782
739			783
740			784
741	T. Pires, E. Schlinger, and D. Garrette. 2019. How Multilingual is Multilingual BERT? In <i>Proceedings of ACL</i> .	$PE_{(pos,2i)} = \sin(pos/10000^{2i/d_{model}})$ (5)	785
742		$PE_{(pos,2i+1)} = \cos(pos/10000^{2i/d_{model}})$ (6)	786
743			
744	R. Sennrich, B. Haddow, and A. Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In <i>Proceedings of ACL</i> .	These equations add information about the position of characters in a sequence. Since our model processes all characters at once, it needs to know which character comes first, second, etc. We use sine and cosine functions to create unique patterns for each position.	787
745			788
746			789
747	T. Stocky, D. Karger, and R. Miller. 2004. AutoFill: Automatic Form Filling. <i>Technical Report, MIT CSAIL</i> .		790
748			791
749			792
750	C. Wang, Y. Li, S. Kang, P. Zhang, C. Meng, and J. Zhou. 2021. Language-Specific Adapters for Efficient Cross-Lingual Transfer. In <i>Proceedings of ACL</i> .	A.2 Subword Model Architecture	793
751		The subword model architecture details are:	794
752		• Vocabulary size: 32,000 SentencePiece tokens	795
753		• Hidden size: 768	796
754	Z. Wang, X. Chen, and Y. Kim. 2022. VITE: Visual Form Understanding via Interactive Web Agents. In <i>Proceedings of EMNLP</i> .	• Feedforward size: 3072	797
755		• Number of attention heads: 12	798
756		• Number of layers: 12	799
757	Y. Wu, Z. Wang, and K. Lee. 2018. Web Form Understanding with Deep Learning. In <i>Proceedings of ICDAR</i> .	• Activation function: GELU	800
758		• Dropout rate: 0.1	801
759		• Maximum sequence length: 128 tokens	802
760	L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, and C. Raffel. 2021. mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer. In <i>Proceedings of NAACL</i> .	• Parameter count: ~ 2.1 million	803
761			
762			
763			
764	Y. Zhao, M. Yang, and D. Xu. 2021. Robust Web Form Understanding with Deep Learning. In <i>Proceedings of ICLR</i> .	A.3 Adapter Architecture	804
765		Each language adapter consists of:	805
766		• Down-projection: Linear layer with output size 64 (from 1024 for character model or 768 for subword model)	806
767	Appendix	• Activation: GELU	809
768	A Model Architecture Details	• Up-projection: Linear layer with output size matching the original hidden size	810
769	A.1 Character-Level Model Architecture	• Layer normalization	812
770	In this appendix, we provide additional details about the architecture of our character-level model.	• Residual connection	813
771	The full architecture specifications are:	The parameter count for each adapter is approximately 1.5% of the base model size.	814
772			815
773	• Vocabulary size: 399 characters		
774	• Hidden size: 1024		
775	• Feedforward size: 4096		
776	• Number of attention heads: 16		
777	• Number of layers: 16		

B Enhanced Character-Level Model Implementation

We now detail our enhanced masked character-level model, which represents a core contribution of this work.

B.1 Character-Level Convolutional Embeddings

Simple character embeddings proved insufficient for capturing morphological patterns, particularly in languages like German with extensive compound words. We addressed this by developing multi-kernel convolutional embeddings that capture various n-gram patterns simultaneously.

The character embedding process works as follows:

$$E_{conv}(x) = [CNN_3(x) \oplus CNN_5(x) \oplus CNN_7(x) \oplus CNN_9(x)] \quad (7)$$

Where CNN_k represents convolution with kernel size k , and \oplus is concatenation. In practice, this means each character is represented not just by itself but by its surrounding context. For example, with kernel size 3, the character ‘a’ in “hat” would be embedded along with ‘h’ and ‘t’.

Implementation challenges included padding and computational efficiency. We ultimately used PyTorch’s Conv1d with appropriate padding to ensure the output maintained the same sequence length as the input.

B.2 N-gram Character Masking

In our first implementation, we only masked individual characters, but this didn’t help the model learn meaningful subword units like prefixes and suffixes. After several experiments, we developed a hybrid masking approach that balances single-character and n-gram masking.

For masking probability, we use:

$$P_{mask}(n) = 0.7 \cdot \delta_{n,1} + 0.3 \cdot \frac{e^{-0.5(n-2)^2}}{\sum_{i=2}^3 e^{-0.5(i-2)^2}} \quad (8)$$

where $\delta_{n,1}$ is the Kronecker delta function. It masks single characters 70% of the time and n-grams of length 2-3 the remaining 30% of the time. The exponential part creates a gentle preference for 2-character sequences over 3-character ones.

We tried several different ratios between single character and n-gram masking (50/50, 80/20, etc.)

before settling on 70/30 based on empirical results. We initially wanted to mask longer n-grams too (up to 5 characters), but found this made training unstable and significantly increased training time without clear benefits.

B.3 Tokenization and Vocabulary

Unlike traditional subword tokenizers with vocabularies of tens of thousands of tokens, our character-level model uses a minimal vocabulary of 399 characters distributed as shown in Table 8.

Character Type	Count
Latin alphabet (upper/lowercase)	52
Digits and punctuation	42
German/French special characters	45
Arabic script characters	185
Special tokens ([PAD], [MASK], etc.)	75
Total	399

Table 8: Character vocabulary breakdown

B.4 Special Tokens Specification

Our character-level model employs 75 special tokens designed specifically for multilingual web automation tasks. The complete breakdown is as follows:

Core Model Tokens (5): [PAD] for sequence padding, [MASK] for character-level masked language modeling, [UNK] for unknown characters, [CLS] for classification tasks, and [SEP] for sequence separation.

Language Identification (4): [EN], [DE], [FR], [AR] for explicit language marking during processing.

Web Form Elements (8): [INPUT], [SELECT], [TEXTAREA], [BUTTON] for basic form elements, and [CHECKBOX], [RADIO], [FILE], [HIDDEN] for specialized input types.

Action Types (8): [FILL], [SELECT], [CHECK], [SUBMIT] for primary form interactions, and [CLICK], [CLEAR], [FOCUS], [SCROLL] for navigation actions.

Field Types (8): [EMAIL], [PASSWORD], [TEXT], [NUMBER] for common field semantics, and [DATE], [TEL], [URL], [SEARCH] for specialized field types.

Navigation Elements (4): [LINK], [NAV], [MENU], [BREADCRUMB] for website structure recognition.

899	Form Structure (8): [FORM_START],	The adapter transformation can be expressed as:	945
900	[FORM_END], [FIELD_START], [FIELD_END]		
901	for structural boundaries, and [LABEL],	$h_{out} = h_{in} + f(h_{in}W_{down})W_{up} \quad (9)$	946
902	[ERROR], [HELP], [REQUIRED] for form		
903	metadata.		
904	Language-Specific Markers (4):	where f is the GELU (Gaussian Error Linear	947
905	[MORPH_RICH], [MORPH_POOR] for	Unit) activation function which provides the	948
906	morphological complexity indicators, and [RTL],	non-linearity for enhancing adapter effectiveness.	949
907	[LTR] for script directionality.	This equation describes how our language adapters	950
908	Context Indicators (4): [CONTEXT_START],	work. For each language, we have a small module	951
909	[CONTEXT_END], [NEARBY],	that adjusts the model’s internal representations.	952
910	[PLACEHOLDER] for field context assembly.	The input (h_{in}) goes through a compression step	953
911	Processing States (4): [PROCESSING],	(W_{down}), a non-linear function (f), and then	954
912	[SUCCESS], [FAILURE], [RETRY] for	expansion (W_{up}). We add this back to the original	955
913	automation state tracking.	input to preserve the important information.	956
914	Reserved Tokens (16): [RES_1] through	C Experimental Details	957
915	[RES_16] allocated for future system extensions	C.1 Training Infrastructure	958
916	without vocabulary retraining.	All models were trained using the following	959
917	These special tokens enable robust multilingual	infrastructure:	960
918	web automation by providing explicit markers for		
919	language context, web element types, user actions,	• Single NVIDIA RTX 3090 GPU (24GB	961
920	and system states. The reserved tokens support our	VRAM)	962
921	modular architecture philosophy, allowing system	• AMD Ryzen 9 5950X CPU	963
922	extension to new languages and domains without		
923	requiring complete vocabulary reconstruction.	• 64GB RAM	964
924	B.5 Model Architecture	• Ubuntu 20.04 LTS	965
925	Our character-level model uses a Transformer	The total training time was approximately:	966
926	architecture optimized for character-level inputs:		
927	16 attention heads, Hidden size of 1024, 16	• Character model: 18 hours	967
928	encoder layers, 2.1 million parameters.	• Subword model: 8 hours	968
929	To handle the longer sequence lengths that	• Language adapters: 2 hours per language	969
930	result from character-level tokenization, we	C.2 Form Filling Dataset Creation	970
931	implement: Efficient attention mechanisms,	Our form filling dataset was created by:	971
932	Optimized positional encodings, Context window		
933	of 512 characters.	1. Scraping 500 common web forms from the	972
934	B.6 Language-Specific Adapters	top 1000 websites	973
935	We integrated language-specific adapters into the	2. Extracting form structure and field semantics	974
936	character model:	3. Translating field labels and descriptions to	975
937		target languages	976
938	• Small adapter modules for each language (en,	4. Generating synthetic natural language	977
939	de, fr, ar)	instructions (20 templates per action type)	978
940	• Each adapter contains a down-projection,	5. Creating valid and invalid form filling	979
941	non-linearity, and up-projection	examples for robust training	980
942	• Adapters are applied after the main		
943	Transformer layers		
944	• Residual connections ensure original		
	information is preserved		

C.3 Evaluation Metrics

The evaluation metrics were calculated as follows:

- **Field Fill Success:** Percentage of fields correctly filled according to user instructions
- **Form Submission Success:** Percentage of forms successfully submitted with all required fields
- **Navigation Success:** Percentage of website navigation tasks completed successfully
- **Data Extraction Precision:** Correct fields extracted / Total fields extracted
- **Data Extraction Recall:** Correct fields extracted / Total fields in form
- **Exact Match (EM):** Exact string match for question answering tasks
- **F1 Score:** Token-level overlap between predicted and reference answers
- **Masked Character Prediction Accuracy:** Correctly predicted masked characters / Total masked characters

C.4 Dataset Statistics

Table 9 provides statistics for the training datasets used in our experiments.

Dataset	Language	Examples	Avg. Length (chars)
OPUS-100	English	50,000	87.3
OPUS-100	German	50,000	92.6
OPUS-100	French	50,000	96.2
OPUS-100	Arabic	50,000	76.8
FQuAD (QA)	French	40,000	124.7
Synthetic Forms	English	25,000	42.1
Synthetic Forms	German	25,000	48.9
Synthetic Forms	French	25,000	51.2
Synthetic Forms	Arabic	25,000	38.6

Table 9: Statistics for training datasets across languages.

D Baseline Implementation Details

We compared our approach against two baselines with carefully controlled methodology to address reproducibility concerns.

mBERT Baseline We fine-tuned pretrained mBERT-base-multilingual-cased on our exact training data (50,000 examples per language) using identical hyperparameters where applicable (learning rate $3e-5$, batch size 64, 5 epochs). No

architectural modifications were made to mBERT—we used its standard token embeddings fed into our form analyzer pipeline. This represents a direct comparison of representation learning capabilities rather than architectural differences.

Rule-based Baseline We implemented pattern matching for common field types (email regex, phone patterns) and keyword matching for field labels in each language. While not state-of-the-art, this represents realistic deployment scenarios for resource-constrained environments and provides a lower-bound baseline.

The mBERT comparison uses identical training data and evaluation protocols, differing only in the underlying representation model, ensuring fair comparison as requested in preliminary feedback.

E Model Selector Implementation Details

E.1 Selector Training Methodology

To train our model selector, we created ground-truth labels by evaluating both character and subword models on 10,000 diverse inputs sampled from our training data. For each input x_i , we computed:

$$\text{label}_i = \begin{cases} 1 & \text{if } \text{performance}_{\text{char}}(x_i) > \\ & \text{performance}_{\text{subword}}(x_i) \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

Performance was measured using form filling accuracy on held-out validation forms.

Feature Engineering Details

- **Length normalization:** $\text{length_norm} = \min(1.0, \text{char_count}/174.6)$ where 174.6 is $2 \times$ dataset median
- **OOV computation:** Using SentencePiece vocabulary with 32,000 tokens, computed as $\text{oov_ratio} = \text{unknown_tokens}/\text{total_tokens}$
- **Morphological scores:** Language-specific scores (German=4, Arabic=4, French=3, English=2) assigned as relative indicators of morphological richness. While this represents a simplified linguistic characterization, these scores effectively enable the selector to distinguish between morphologically complex languages (favoring character

models) and simpler ones (favoring subword models) in our experimental setting.

Language Detection: Before computing morphological complexity scores, we detect input language using character pattern matching: German (ü,ö,ä,ß patterns), French (é,à,ç,î patterns), Arabic (Unicode ranges U+0600-U+06FF), English (default for Latin script without special characters). This detection occurs independently before selector feature computation, resolving the dependency between language identification and morphological scoring.

- **Noise detection:** $\text{noise_level} = (\text{special_chars} + \text{digits}) / \text{total_chars}$

Training Configuration

- Optimizer: Adam with learning rate 0.001
- Batch size: 128
- Training epochs: 50 with early stopping (patience=10)
- Loss function: Binary cross-entropy
- Train/validation split: 80/20

E.2 Selection Performance Analysis

Table 10 shows detailed selector performance across languages and input types.

Input Type	Selector Acc.	Char	Subword
Short (< 30 chars)	92.1%	78.3%	21.7%
Medium (30-100 chars)	88.7%	42.1%	57.9%
Long (> 100 chars)	85.9%	23.4%	76.6%
High OOV (> 10%)	94.3%	89.2%	10.8%
Low noise (< 5%)	87.1%	35.6%	64.4%
High noise (> 15%)	91.8%	82.7%	17.3%
Overall	89.3%	51.2%	48.8%

Table 10: Model selector performance across input characteristics. Selector accuracy measured against oracle choices.

E.3 Ablation Study: Learned vs. Heuristic Selection

We compared our learned selector against a heuristic baseline using simple rules:

- Character model: if length < 30 chars OR oov_ratio > 0.05 OR noise_level > 0.1
- Subword model: otherwise

Results on form filling accuracy:

- **Learned selector:** 87.7% average accuracy
- **Heuristic baseline:** 75.4% average accuracy
- **Always character:** 82.3% average accuracy
- **Always subword:** 86.1% average accuracy
- **Improvement:** +12.3% over heuristic, +1.6% over best single model

The learned selector’s primary advantage lies in handling edge cases where simple heuristics fail, particularly for medium-length inputs with moderate OOV rates.

E.4 Computational Overhead Analysis

Component	Time (ms)	Memory (MB)
Feature extraction	0.8	0.1
Selector inference	1.5	0.3
Model loading	0.0	0.0
Total selector overhead	2.3	0.4
Character model inference	67.2	850
Subword model inference	41.8	440

Table 11: Computational overhead of model selection vs. model inference.

The selector adds negligible computational cost (2.3ms) compared to model inference (42-67ms).

F Additional Results

F.1 Effect of N-gram Masking




Figure 4: Impact of n-gram masking on masked character prediction accuracy.

F.2 Cross-lingual Transfer Matrix

Table 12 provides a more detailed cross-lingual transfer matrix showing how adapters trained on one language perform on other languages.

Train Lang	Test Lang	EN	DE	FR	AR
EN	EN	18.31%	10.00%	0.00%	0.00%
EN	DE	0.00%	2.00%	16.67%	0.00%
EN	FR	0.00%	0.00%	8.44%	0.00%
EN	AR	28.57%	12.50%	22.22%	7.22%

Table 12: Cross-lingual transfer matrix showing adapter performance across languages.

F.3 Ablation Studies

Figure 5 shows the impact of removing different components from our system.




Figure 5: Component ablation reveals that language-specific adapters provide the largest performance gain (+14.2%), while n-gram masking and convolutional embeddings each contribute 3-5% improvements, validating our architectural choices.

G Web Automation Examples

Below we provide examples of natural language instructions and their corresponding structured actions for web form filling:

G.1 English

- **Instruction:** “Fill in the email field with john.smith@example.com”
- **Action:** `{"action_type": "fill", "field_name": "email", "value": "john.smith@example.com"}`

G.2 German

- **Instruction:** “Gib in das Passwortfeld ‘Secure123!’ ein”
- **Action:** `{"action_type": "fill", "field_name": "password", "value": "Secure123!"}`

G.3 French

- **Instruction:** “Sélectionne ‘Femme’ dans le menu déroulant de genre”
- **Action:** `{"action_type": "select", "field_name": "gender", "value": "female"}`

G.4 Arabic

- **Instruction:** “ ” (Click the submit button)
- **Action:** `{"action_type": "submit", "field_name": "submit_button", "value": null}`

H Model Scaling Analysis

We conducted experiments to analyze how our models scale with different parameter sizes. Figure 6 shows the results of these experiments.




Figure 6: Model scaling analysis showing accuracy vs parameter count.

The results indicate that both models benefit from increased parameter counts, but the gains diminish at larger sizes. The character model shows more consistent scaling benefits, suggesting it may benefit from even larger model sizes in future work.

I Computational Efficiency Analysis

Table 13 compares the computational efficiency of our approach versus baseline models.

J Error Analysis

J.1 Error Categories

We categorized errors in our system into four main types:

1. **Language Understanding Errors:** Incorrectly parsing the user’s natural language instruction (23% of errors)

Model	Parameters	Training Time	Inference Time (ms)	Memory (MB)
Character Model	~2.1M	18 hours	67	850
Subword Model	~2.1M	8 hours	42	440
mBERT (baseline)	~175M	Pretrained	86	700
XLM-R Large	~550M	Pretrained	215	2,200

Table 13: Computational efficiency comparison across models.

1157 2. **Form Analysis Errors:** Failing to correctly
1158 identify form elements or their purposes (42%
1159 of errors)

1160 3. **Action Execution Errors:** Correctly
1161 understanding but failing to execute the
1162 intended action (19% of errors)

1163 4. **Other Errors:** System crashes, timeouts, or
1164 unclassified errors (16% of errors)

1165 J.2 Performance by Form Complexity

1166 Table 14 shows how performance varies with form
1167 complexity.

Form Complexity	Fields	Success Rate
Simple	1-3	94.2%
Medium	4-7	88.7%
Complex	8+	81.3%

Table 14: Form filling success rate by form complexity.

1168 J.3 Error Examples

1169 Table 15 provides examples of common errors and
1170 their analysis.

Error Type	Example	Analysis
Language Understanding	“Fill the phone with 555-1234” misinterpreted	Ambiguous field reference (“phone” vs “phone number”)
Form Analysis	Could not locate “billing-address” field	Field had non-standard HTML attributes
Action Execution	Failed to select option in custom dropdown	JavaScript-rendered dropdown not accessible via DOM
Language Understanding	Failed to parse German compound noun	Character model struggled with morphological complexity
Form Analysis	Confused similar field labels “shipping” vs “billing”	Semantic similarity caused field misidentification

Table 15: Examples of common errors encountered during evaluation.