# WaterPark: A Robustness Assessment of Language Model Watermarking

**Anonymous ACL submission** 

### Abstract

Various watermarking methods ("watermarkers") have been proposed to identify LLMgenerated texts; yet, due to the lack of unified evaluation platforms, many critical questions remain under-explored: i) What are the strengths/limitations of various watermarkers, especially their attack robustness? ii) How do various design choices impact their robustness? iii) How to optimally operate watermarkers in adversarial environments? To fill this gap, we systematize existing LLM watermarkers and watermark removal attacks, mapping out their design spaces. We then develop WATERPARK, a unified platform that integrates 10 state-ofthe-art watermarkers and 12 representative attacks. More importantly, by leveraging WA-018 TERPARK, we conduct a comprehensive assessment of existing watermarkers, unveiling the impact of various design choices on their attack robustness. We further explore the best practices to operate watermarkers in adversarial environments. We believe our study sheds light on current LLM watermarking techniques while WATERPARK serves as a valuable testbed to facilitate future research.<sup>1</sup>

#### Introduction 1

011

019

The recent advances in large language models (LLMs), including GPT (Openai) and Llama (Touvron et al., 2023), have significantly enhanced our capabilities of general-purpose text generation and complex problem-solving, but also raised concerns about misuse through disinformation (Izacard et al., 2022), phishing (Bender et al., 2021), and academic dishonesty (Compilatio). There is thus a pressing need for the capability of identifying LLMgenerated content.

A simple approach is to train classifiers to distinguish between LLM- and human-generated texts (Lütkebohle). However, as LLMs improve,



Figure 1: Illustration of LLM watermarking and watermark removal attacks.

041

042

043

044

047

048

050

051

054

057

058

060

061

062

064

065

066

067

068

069

071

072

this distinction becomes less clear. Watermarking has emerged as an alternative solution, embedding statistical signals ("watermarks") during generation to verify LLM-produced texts. Various watermarking methods ("watermarkers") have been developed (Aaronson and Kirchner; Liu et al., 2023a; Liu et al., 2024; Kirchenbauer et al., 2023a; Hu et al., 2024; Zhao et al., 2023; Kuditipudi et al., 2023), each with unique design choices and desirable properties, raising a set of intriguing questions:

RQ1 – What are the strengths and limitations of various watermarkers, especially their robustness against manipulations?

RQ2 – How do different design choices impact the attack robustness of watermarkers?

RQ3 – What are the best practices for operating watermarkers in adversarial environments?

Despite recent efforts to benchmark LLM watermarkers, existing research is limited in addressing these questions. WaterBench (Tu et al., 2023) primarily focuses on watermarking effectiveness; MarkMyWords (Piet et al., 2023) mainly evaluates the robustness of a specific watermarker (Kirchenbauer et al., 2023a); MarkLLM (Pan et al., 2024) focuses on providing a platform to compare watermark detectability, basic robustness, and text quality. Zhao et al. (2024) provides a comprehensive overview of watermarking techniques. Moreover, these studies lack an in-depth analysis of how a watermarker's design choices impact its robustness. Consequently, the aforementioned questions remain largely unexplored.

<sup>&</sup>lt;sup>1</sup>All the source code and data are publicly available: http s://anonymous.4open.science/r/WaterPark

Previous Conclusion	Refined Conclusion	Explanation	Consistency
UG (Zhao et al., 2023) is not robust due to its context-free design (Piet et al., 2023).	UG shows higher resilience than other watermarkers to paraphrasing attacks.	UG's context-free design ensures consistency between detection and generation, avoiding issues in text- dependent designs.	0
UPV (Liu et al., 2023a) has a fairly low false positive rate.	UPV is more prone to false positive cases compared to TGRL.	While model-based detection incurs	Ð
UPV shows strong robustness to rewriting and outperforms TGRL under paraphrasing attacks.	Both UPV and TGRL struggle against GPT-based paraphrasing attacks.	score-based detection, it fails to offer higher flexibility in countering paraphrasing attacks.	O
RDF (Kuditipudi et al., 2023) signif- icantly outperforms TGRL against substitution attacks.	RDF shows strong robustness against synonym substitution and other lexical editing attacks.	RDF's distribution-transform strat- egy is more robust than the distribution-shift strategy against lexical editing.	•
RDF's (Kuditipudi et al., 2023) edit score-based detection is insensitive to local misalignment caused by to- ken insertion (Piet et al., 2023).	RDF (edit score) shows higher resilience against lexical editing attacks compared to GO (Aaronson and Kirchner) (plain score) when token length is fixed. How- ever, this advantage diminishes as token length varies.	The edit score-based detection is ro- bust to lexical editing but sensitive to varying token length.	D
SIR (Liu et al., 2024) shows strong resilience against paraphrasing at- tacks.	SIR's effectiveness decreases as the intensity of paraphrasing attacks increases.	SIR's distribution-reweight strategy introduces higher uncertainty, mak- ing it sensitive to the intensity of paraphrasing attacks.	D
UB (Hu et al., 2024) is more robust than TGRL to substitution attacks.	UB and TGRL are comparably robust to synonym substitution attacks; yet, UB is more vulnerable to paraphrasing attacks.	The distribution-reweight strategy is not superior to the distribution-shift strategy.	0

Table 1: Conclusions in prior work and WATERPARK (○ – inconsistent; ● – partially inconsistent; ● – consistent).

To bridge this gap, this work conducts a systematic study of state-of-the-art LLM watermarkers, focusing on their attack robustness. We aim to understand how various design choices affect attack resilience and identify best practices for operating watermarkers in adversarial environments.

We develop WATERPARK, the first open-source platform dedicated to evaluating the attack robustness of LLM watermarkers in a unified and comprehensive manner. As of 12/15/2024, WATERPARK integrates 10 state-of-the-art watermarkers, 12 representative watermark removal attacks, and 8 key metrics. Moreover, WATERPARK offers a comprehensive suite of tools for in-depth robustness assessment, including next-token distribution comparisons, attack combination analyses, and what-if scenario evaluations. Leveraging WATERPARK, we empirically evaluate the attack resilience of representative LLM watermarkers, leading to many interesting findings, which challenge the conclusions in prior work, as summarized in Table 1. We also explore how a watermarker's design choices impact its attack robustness, unveiling critical trade-offs between different types of robustness.

# 2 LLM Watermarking

A large language model (LLM) is typically an autoregressive model that generates the next token  $x_t$ based on previous tokens  $x_{<t} \triangleq x_1, \ldots, x_{t-1}$  (including its prompt), modeled as sampling from a conditional distribution  $p(x_t|x_{< t})$ .

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

Conceptually, a watermark is a pattern embedded in a given signal (e.g., text) for identifying the signal's source. In the context of LLMs, watermarks can be used to prove that a given text is LLM-generated (or even generated by specific LLMs). An LLM watermarking method ("watermarker") comprises three components: the LLM, watermarking procedure (generator), and detection procedure (detector), as shown in Figure 1.

Typically, the generator produces watermarked texts iteratively. At each iteration, with access to a secret key k, the previous tokens  $x_{<t}$ , and the LLM's next-token distribution  $p(x_t|x_{<t})$ , the generator generates a perturbed distribution  $\tilde{p}(x_t|x_{<t})$ , from which the next token  $x_t$  is sampled. Meanwhile, with access to a secret key k', the detector determines whether M generates a given text. Note that in symmetric schemes, k = k', while in asymmetric schemes (Fairoze et al., 2023),  $k \neq k'$ . This study mainly focuses on symmetric watermarking schemes due to their widespread adoption.

The current watermarkers can be categorized based on the information carried by the watermarks (e.g., one-bit versus multi-bit) and their key design choices, including context dependency, generation strategy, and detection method. The key design factors of each category are deferred to §B.

100

#### **Platform** 3

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147

148

149

150

151

152

153

155

156

157

158

159

160

161

162

164

165

171

173

#### 3.1 **Threat Model**

One critical property of LLM watermarkers is their robustness against potential attacks. We assume a threat model similar to prior work (Piet et al., 2023; Zhao et al., 2023; Sankar Sadasivan et al., 2023; Kirchenbauer et al., 2023b), as shown in Figure 1. We assume the adversary has access to sample watermarked and non-watermarked texts, but cannot reproduce the watermarking procedure or interact with the detection procedure. The adversary modifies the watermarked text  $\tilde{T}$  as the altered text T' such that the following objectives are met: effectiveness – T' evades the watermark detector (i.e., detected as non-watermarked) and quality -T' preserves  $\tilde{T}$ 's original semantics.

# 3.2 Metrics

Effectiveness. At a high level, WATERPARK evaluates the watermark detector's accuracy in detecting watermarked texts mainly using two metrics: true positive rate (TPR) and false positive rate (FPR). TPR measures the fraction of watermarked texts detected as watermarked, while FPR measures the fraction of non-watermarked samples wrongly detected as watermarked. Formally, let  $S_+$  and  $S'_+$  respectively be the sets of ground-truth and detected watermarked texts ( $S_{-}$  and  $S'_{-}$  correspondingly).

$$\text{TPR} = \frac{|S_+ \cap S'_+|}{|S_+|} \qquad \text{FPR} = \frac{|S_- \cap S'_+|}{|S_-|} \quad (1)$$

In WATERPARK, we plot the receiver operating characteristic (ROC) curve that measures TPR against FPR across varying settings of detection thresholds. In particular, the area under the ROC curve (AUC) evaluates the overall effectiveness of each watermarker. Moreover, to compare two watermarkers under specific settings, we may also measure their TPRs under a fixed FPR (e.g., 1%).

Fidelity. To evaluate the impact of watermarking 166 on text quality, WATERPARK employs the follow-167 ing metrics to measure the difference between the original text T and the watermarked text T. We 169 employ WER (word error rate), BLEU (Papineni 170 et al., 2002), BERTScore (Zhang et al., 2020) and P-SP (Wieting et al., 2022) as metrics to assess fidelity. Detailed descriptions of these metrics can be found in §C. Note that these metrics are also used 174 to measure the impact of watermark removal at-175 tacks on the quality of the modified text  $\tilde{T}'$ , relative 176 to the watermarked text T. 177

Attack	Category	Resource
Lowercasing		/
Contracting	Linguistic variation	/
Expanding		/
Misspelling		Common misspellings
Typoing	Levical editing	/
Synonymizing	Lexical cutting	WordNet
Swapping		/
Copy-pasting	Text-mixing	Non-watermarked text
Deep-paraphrasing	Paraphrasing	LLM-based paraphraser
Translating	rarapinasing	LLM-based translator
Black-box adv	versarial attack	Generic detector

Table 2: A taxonomy of watermark removal attacks.

Robustness. To evaluate a watermarker's attack resilience, WATERPARK measures the attack's impact on the watermarking effectiveness. Specifically, let  $\tilde{T}$  and  $\tilde{T}'$  respectively denote the watermarked text before and after the attack. WATER-PARK compares the detector's TPR, FPR, and AUC with respect to  $\tilde{T}$  and  $\tilde{T}'$ . Intuitively, a smaller difference indicates that the watermarker is less attack-sensitive.

178

179

180

181

182

183

184

185

186

187

189

190

191

192

193

194

195

196

197

199

200

201

203

204

205

206

207

209

210

211

212

# **4** Evaluation

We leverage WATERPARK to empirically assess representative LLM watermarkers, focusing on their attack robustness and other related criteria.

#### 4.1 **Experimental Setting**

Our evaluation considers all the watermarkers in Table 6 and attacks in Table 2. We employ two LLMs OPT-1.3B and Llama2-7B-chat-hf to represent small and large models respectively. We use two datasets C4 and HC3 to simulate diverse tasks such as question-answering and text completion.

We implement and configure all watermarkers and attacks in accordance with their official documentation. Table 5 summarizes the default parameter settings. To validate our implementation, we assess their effectiveness and fidelity, as reported in §E and §F, corroborating results from existing literature. A detailed description of each attack can be found in §D. We measure the TPRs of different watermarkers against each attack (with their FPRs fixed as 1%).

# 4.2 Robustness – Observational Study

We systematically assess the resilience of existing LLM watermarkers against representative watermark removal attacks. All the results are shown in Table 3.

Water	CLEAN	Linguistic variation		Lexical editing			Text-mixing				Paraphrasing				
marker		Contra	Expan	LowCase	Swap	Туро	Syno	Missp	CP1-10	CP3-10	CP1-25	CP3-25	DP-20	DP-40	Trans
TGRL	0.992	0.994	0.994	0.984	0.874	0.788	0.996	0.992	0.034	0.176	0.126	0.836	0.952	0.858	0.652
UG	0.964	0.963	0.965	0.879	0.956	0.884	0.965	0.956	0.131	0.219	0.192	0.534	0.940	0.916	0.658
UPV	0.776	0.772	0.814	0.685	0.667	0.355	0.764	0.695	0.050	0.073	0.242	0.450	0.296	0.057	0.006
RDF	0.996	0.993	0.993	0.983	0.951	0.961	0.990	0.993	0.094	0.396	0.430	0.974	0.988	0.962	0.154
UB	1.000	1.000	1.000	0.946	0.030	0.088	0.992	0.996	0.002	0.009	0.048	0.410	0.461	0.149	0.306
SIR	0.954	0.954	0.954	0.886	0.932	0.566	0.936	0.904	0.138	0.210	0.178	0.466	0.902	0.814	0.818
GO	0.998	0.998	0.998	0.988	0.920	0.852	0.996	1.000	0.181	0.589	0.620	0.992	0.971	0.847	0.928

Table 3: Attack resilience of LLM watermarkers. The intensity of red shading indicates higher values, while the intensity of blue shading indicates lower values, with 0.5 serving as the threshold between the two color gradients.

#### 4.2.1 Linguistic Variation Attack

This attack perturbs the linguistic features of the watermarked text, without changing its semantics.

i) Overall, most watermarkers show strong resilience against linguistic variation attacks. For instance, TGRL reaches close to 100% TPRs under all three attacks. ii) UPV is the only watermarker marginally susceptible to such attacks. This can be explained by the fact that its neural network-based detector primarily depends on implicit textual features, which appear to be sensitive to changes in linguistic characteristics.

## 4.2.2 Lexical Editing Attack

This attack modifies individual words while maintaining the watermarked text's semantics.

i) TRGL, UB, SIR and GO tend to be more vulnerable to lexical editing attacks, compared with RDF and UG watermarkers. Intuitively, as textdependent watermarkers (e.g., TRGL, UB, and GO) use previous tokens as the context for the next token in both the watermark generator and detector, the lexical editing thus causes a mismatch between the generator and detector. ii) UB exhibits much higher vulnerability to such attacks, compared with the others. For example, its TPR drops near zero under the typoing and swapping attacks. This may be explained as follows. Recall that, to achieve unbiasedness, UB applies "hard" perturbation on the next-token distribution (e.g., by rejecting half of the vocabulary); thus the disruption to the previous tokens tends to cause a more significant mismatch between the generator and detector, compared with other watermarkers that employ "soft" perturbation (e.g., distribution shift and transform). iii) The SIR shows notably inferior performance, especially when subjected to typoing attacks. While it employs a neural network to predict token-specific logit perturbations that are designed to be unbiased (no preference for particular tokens) and balanced (with perturbations summing to zero), its results are suboptimal, suggesting that it is fundamentally challenging to make accurate predictions under such perturbations. iv) The typoing attack is particularly effective, as it may cause significant tokenization errors. 252

253

254

255

257

258

259

261

262

263

264

265

266

267

268

269

270

271

272

273

274

275

276

277

278

279

281

282

287

290

#### 4.2.3 Text-Mixing Attack

This class of attacks "dilutes" the watermark by mixing the watermarked text with nonwatermarked text fragments. Here, to evaluate the resilience of different watermarkers against text-mixing attacks, we use the copy-pasting attack (Kirchenbauer et al., 2023b) as the concrete attack, which embeds the watermarked text into the context of non-watermarked, human-written text (generated under the same prompt). We use CP-n-m to denote the attack in which the modified text T' consists of n segments of watermarked texts, each of length m% of |T'|, and the rest as non-watermarked text.

i) Overall, most watermarkers experience significant TPR drops, especially under CP-1-10 that only preserves 10% of the watermarked text. ii) Among all the watermarkers, GO and RDF show significantly higher attack resilience. This can be attributed to their sampling strategy: both employ distribution transform, which generates the next token deterministically conditional on a given random permutation. Thus, GO and RDF tend to have stronger per-token signals than the other watermarkers that sample the next token from a given pool (e.g., green list). This observation is consistent with that in §E.2. iii) Meanwhile, UB and UPV are the most vulnerable to the copy-pasting attack, with close to zero TPRs under CP-1-10 and CP-1-25. This can be explained as follows. The model-assisted detector of UPV determines the given text as watermarked based on its aggregated features (rather than per-token statistics), while the

237

240

241

242 243

245

247

248

251

213

injected non-watermarked segments may greatly disrupt such features. Meanwhile, UB applies hard perturbation on the next-token distribution (e.g., by rejecting half of the vocabulary); thus the disruption to the previous tokens causes a significant mismatch between the generator and detector.

# 4.2.4 Paraphrasing Attack

291

292

296

301

313

314

316

318

319

320

321

322

324

325

326

329

331

332

334

337

340

This class of attacks employs an additional LLM (i.e., paraphraser) to re-write the given watermarked text  $\tilde{T}$  (while preserving its semantics) to evade the detector. Here, we consider Dipper (Krishna et al., 2023) as the paraphraser that rewrites  $\tilde{T}$  in one shot. Further, we also consider the translating attack uses a translator model Seamless-m4tv2-large (Communication et al., 2023) that first translates  $\tilde{T}$  to French and then translates it back.

i) UPV and UB exhibit higher vulnerability to the Dipper attack, compared to other watermarkers. UPV's TPR drops to around zero under DP-40, which aligns with our analysis in §4.2.3: the paraphrased text segments may greatly disrupt the aggregated textual features for UPV's model-assisted detector, while the disruption to the previous tokens may cause a substantial mismatch between UB's generator and detector, due to its rigid perturbation to the next-token distribution (e.g., rejecting half of the vocabulary). ii) In contrast, RDF and UG are especially robust against the Dipper attack. This can be attributed to their index-dependent and context-free designs, which are less sensitive to the change of previous tokens than text-dependent watermarkers (e.g., TGRL, SIR, and GO). iii) Interestingly, RDF is more vulnerable to the translating attack than the Dipper attack. This susceptibility arises from RDF's detection mechanism, which is highly sensitive to text length variations, a weakness readily exploited by the translating attack's tendency to produce shorter output text.

#### 4.2.5 Fidelity Preservation

Recall that besides their attack effectiveness, another key metric for watermark removal attacks is whether they can preserve the quality of original texts. We thus compare the semantics of watermarked text  $\tilde{T}$  and modified text T' using the metrics in §3.2. Figure 2 illustrates the quality preservation of different attacks on GO, with similar results on other watermarkers, with more results in §G.3.

Observe that most attacks preserve the semantics of the watermarked text  $\tilde{T}$  in the modified text  $\tilde{T}'$ , as measured by BERTScore and P-SP scores. In



Figure 2: Quality preservation of different attacks.

comparison, the copy-pasting (CP) attack causes more significant text-quality degradation than other attacks, in that it may disrupt the orders of watermarked and non-watermarked segments and insert duplicate segments. Also, note that most attacks emphasize the semantic similarity between  $\tilde{T}$  and  $\tilde{T}'$  rather than their lexical similarity (as measured by WER and BLEU scores).

# 4.3 Robustness – Causal Analysis

In addition to the observational studies, we further consider conducting causal analysis to understand the impact of individual design choices (e.g., samplers). However, this is challenging in our context because the various components of a watermarker are often highly interconnected and difficult to decouple. For example, the distribution-reweight strategy and the difference-based detection of UB are closely linked and cannot be easily replaced by other designs. To address this challenge, we select two watermarkers with their only difference in the design of one specific component (e.g., context dependency). The results are shown in Figure 3.



Figure 3: Watermarker robustness to multi-attacks. a) Context dependency: TGRL (text-dependent) and UG (context-free); b) Generation strategy: TGRL (distribution-shift) and GO (distribution-transform); c) Detection method: UPV (Model-based) and UPV-key (Score-based).

# 4.3.1 Context Dependency

We select TGRL and UG to represent textdependent and context-free designs respectively. TGRL uses the previous k tokens as a randomness seed to divide the vocabulary into red/green lists for the current token, whereas UG uses fixed red/green lists in generating all the tokens. In gen-

369

341

342

343

344

345

348

350

351

352

354

355

357

359

360

361

eral, UG performs similarly to TGRL. However, 370 when subjected to the Dipper attack at varying in-371 tensity levels, UG consistently outperforms TGRL. For instance, as we increase the attack intensity (e.g. DP-60-20, DP-60-40), which involves rearranging the text, UG maintains a relatively stable TPR with only minor reductions. In contrast, TGRL's TPR 376 significantly drops, particularly at DP-60-20. Recall that the Dipper attack extensively paraphrases and reorders the text, leading to substantial changes in consecutive tokens. Due to its text-dependent design, TGRL struggles to maintain consistency in random seeds between detection and generation. In contrast, UG's context-free design avoids this issue. This aligns with the results in UG (Zhao et al., 384 2023), showing that when subjected to paraphrasing attacks, UG consistently outperforms TGRL.

#### 4.3.2 Generation Strategy

390

391

400

401

402

403

404

405

406

407

408

409

410

411

TGRL and GO are both context-dependent watermarkers. As it relies on distribution shift generation, TGRL tends to be less robust against copypaste attacks. In contrast, although copy-paste attacks can significantly disrupt detection and dilute the watermark by inserting large amounts of text, due to its distribution transform design, GO can better maintain the watermark's concentration. As a result, GO achieves a higher detection rate of watermarked text after such attacks.

# 4.3.3 Detection Method

UPV employs model-based detection to detect watermarks. Alternatively, it can also use score-based detection similar to TGRL. The original paper claims that model-based detection is more resistant to paraphrasing attacks than score-based detection. However, our results indicate that this advantage is not significant. Furthermore, score-based UPV demonstrates stronger robustness against weaker attacks (*e.g.*, misspelling and swapping attacks). Using model-based detection introduces higher uncertainty compared to score-based detection and does not provide watermarkers with greater flexibility in resisting paraphrasing attacks.

Both RDF and GO utilize distribution-transform 412 generation and rely on score-based detection, with 413 RDF specifically employing edit score. RDF shows 414 415 marginally higher robustness against lexical editing attacks with fixed token length compared to GO. 416 However, it exhibits lower resilience against high-417 intensity text-mixing attacks (e.g., CP-3-10), where 418 token length varies considerably. 419

# 5 Discussion

Next, we examine the current practices of operating watermarkers in adversarial environments and explore potential improvements.

### 5.1 Specific vs. Generic Detector



Figure 4: Detection of watermarked texts by watermarker-specific and generic detectors ('1' or '0' indicate that the detector detects the given watermarked text as watermarked or non-watermarked).

For each watermarker, we mainly use its specific detector to detect watermarked texts. Here, we explore a generic, neural network-based detector as an alternative. To this end, we employ a pre-trained RoBERTa model and fine-tune it as a binary classifier using watermarked and nonwatermarked texts. We use OPT-1.3B as the underlying LLM and C4 as the reference dataset. The watermarked text is generated by the LLM and the watermarker jointly, whereas the non-watermarked text is produced by ChatGPT-3.5-Turbo using the same prompt to mimic human response. Table 4 shows the TPRs of watermarkers with generic detectors (with FPRs fixed as 1%).

Watermarker TGRL RDF UB UG UPV SIR GO

 TPR
 0.985 | 0.981 | 0.989 | 0.999 | 0.997 | 0.996 | 0.970 

 Table 4: TPRs of watermarkers with generic detectors (with FPRs fixed as 1%).

We compare the attack resilience of watermarkerspecific and generic detectors. For each watermarker, we apply the Dipper-40 attack and examine whether two detectors can effectively detect watermarked texts after the paraphrasing attack. Figure 4 depicts the confusion matrices of both detectors.

We have the following findings. i) The specific and generic detectors jointly achieve a high detection rate. Across all the watermarkers, the chance that both detectors fail to detect the watermarked texts (0-0) is below 0.07. ii) For RDF, the generic detector seems less effective than the specific detector, while for UB and UPV, the generic detector outperforms the specific detector by a large margin. 420

421 422 423

424

425

426

427

428

429

430

431

432

433

434

435

436

437

438

439

440

441

442

443

444

445

446

447

448

449

450

451

Recall that UB and UPV are highly sensitive to 453 the Dipper attack (see §4.2.4. Thus, employing a 454 generic detector alongside a watermarker-specific 455 detector can be an effective strategy for enhancing 456 the security of vulnerable watermarkers. However, 457 note that given the availability of generic detectors, 458 it is also feasible for the adversary to leverage such 459 detectors as an attack checker to adapt their attacks, 460 which we will discuss in §5.2.3. 461

### 5.2 Advanced Attack

462

463

464

465

466

467

468

469

470

471

472

473

474

475

476

477

478

479

480

481

482

483

484

485

486

487

488

489 490

491

492

493

494

In addition to the previous simple attacks, we investigate the effectiveness of more advanced attacks.

#### 5.2.1 Varying Attack Intensity



Figure 5: Resilience of watermarkers against increasingly intensive Dipper attacks.

One straightforward way to improve an attack's effectiveness is to increase its intensity, potentially at the cost of other metrics (e.g., text quality). Here, we consider the Dipper attack (Krishna et al., 2023) under varying intensity settings, denoted as DP-l-o, where lexical change (l) indicates that l% of the given text is paraphrased and order change (o) indicates that o% of the text is re-ordered. We compare the watermarkers' resilience under varying attack intensity, as shown in Figure 5.

i) As expected, most watermarkers observe TPR drops as the attack intensity increases. ii) Among these, UG demonstrates the most consistent resilience under varying attack intensity. This can be attributed to its context-free design: the same perturbation is applied to the next-token distribution across all the tokens, which is thus immune to the change of previous tokens. iii) Compared with textdependent watermarkers (e.g., TGRL and GO), an index-dependent watermarker (e.g., RDF) shows stronger resilience, especially under high attack intensity (e.g., DP-60-20), due to its weaker dependency on previous tokens. iv) UPV's performance is inconsistent; it struggles with low-intensity attacks (e.g., DP-20) but shows resilience to highstrength ones (e.g., DP-60). This inconsistency can be attributed to UPV's model-assisted detector and the inherent instability of its neural network, as confirmed by repeated experiments.



Figure 6: Resilience of watermarkers against individual (left) and combined (right) attacks.

#### 5.2.2 Combining Simple Attacks

We first explore whether combining two attacks improves the attack's effectiveness. Here, considering the most feasible combinations, we only combine two simple attacks from the "weak" linguistic variation and lexical editing attacks. Figure 6 illustrates the effectiveness of such combined attacks.

We have a set of interesting observations. i) UPV and SIR, which demonstrate resilience against all simple attacks, are highly vulnerable to all the combined attacks. For instance, the TPR of SIR drastically drops to below 0.3 under the contracting+typoing attack. ii) UB, which is vulnerable to the typoing and swapping attacks, is consequently vulnerable to all the combined attacks that involve typoing or swapping. iii) TGRL and GO, which are robust against all the simple attacks (including typoing and swapping), show significant vulnerability to the typoing+swapping attack. This can be explained by that as typoing and swapping respectively disrupt the tokenization and token-indexing, their combination may substantially amplify such effects. iv) RDF and UG are especially robust against the combined attacks. This can be attributed to their "weaker" context dependencies, which is consistent with the findings in §4.2.4.

## 5.2.3 Adaptive Attack



Figure 7: Attacks leveraging surrogate detectors.

Given the availability of generic detectors, it is possible for the adversary to exploit such detectors to adapt their attacks. We consider a scenario as shown in Figure 7: the adversary performs a gradient-based attack (Guo et al., 2021) that itera521

522

523

524

525

526

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

529

531

532

533

534

540

541

542

544

545

546

548

549

550

551

552

554

tively modifies the watermarked text to evade the generic detector, and then forwards the modified text to the target, watermark-specific detector.



Figure 8: Watermark detection by watermarker-specific and generic detectors on gradient-based attacked samples ('1' or '0' indicate that the detector detects the given watermarked text as watermarked or non-watermarked).

We evaluate this attack's effectiveness using 500 watermarked texts from the C4 dataset. We use GBDA (Guo et al., 2021) as the adversarial attack and limit the steps of perturbations to 100. Compared to Dipper, GBDA makes more substantial textual modifications. The results are summarized in Figure 8. i) Leveraging the surrogate detector significantly improves the attack effectiveness: with the BERTScore and BLEU-1 between  $\tilde{T}$  and  $\tilde{T}'$  is about 0.764 and 0.188, respectively, slightly lower than the Dipper attack, the detection rates of most watermarkers drop below 10%. ii) Although some samples do not evade the generic detector, they evade the specific detector successfully (e.g., TGRL, UG, and UB). iii) UG and RDF exhibit greater robustness than the other watermarkers. Specifically, for RDF, 42% of the samples evade the generic detector but are still detected by the RDF-specific detector. This superior robustness is likely due to their weaker context dependencies.



Figure 9: Effectiveness of paraphrasing attacks with ChatGPT-3.5-Turbo as the paraphraser.

# 5.2.4 Leveraging Expert LLMs

We explore the question of "What if the adversary has access to highly capable LLMs?" Specifically, we implement another paraphrasing attack that employs ChatGPT-3.5-Turbo as the paraphraser. For each watermarker, we randomly sample 100 watermarked texts that are successfully detected, and query the ChatGPT API with the prompt: "Paraphrase the following text and keep the length simi*lar to the original text*/*n* [*the watermarked text*]", and then forward the paraphrased text to the detector. To further evaluate the impact of the paraphrasing strength, we also measure the attack effectiveness under multiple rounds of paraphrasing as suggested in (Zhang et al., 2023). Notably, the attack described in (Zhang et al., 2023) employs the T5 model to paraphrase the entire text up to 300 times. In contrast, our study limits paraphrasing to a maximum of 5 iterations. The difference between these threat models is due to the distinct objectives of the two studies: while (Zhang et al., 2023) aims to understand the lower bound of watermark robustness, we focus on evaluating watermark robustness against practical, resource-limited adversaries. The results are summarized in Figure 9.

555

556

557

558

559

560

561

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

591

592

593

594

595

597

598

599

601

602

603

604

605

We have the following findings. i) The detection rates of all the watermarkers drop below 0.3 after one round of GPT paraphrasing, indicating that using highly capable LLMs to paraphrase watermarked texts is a dominant attack that effectively nullifies most watermarkers. However, as specified in our threat model (§3.1), access to highly capable LLMs may fall outside the scope of our robustness assessment, as their availability negates the need for watermark removal attacks. ii) For watermarkers that survive the first round of paraphrasing (e.g., UG and GO), their detection rates quickly drop below 0.15 as the adversary applies multiple rounds of paraphrasing. This observation corroborates the findings in (Zhang et al., 2023).

## 6 Conclusion

In this paper, we systematize the existing LLM watermarkers and watermark removal attacks, mapping out the design space for various watermarking and attacking techniques. We then design and implement WATERPARK, the first open-source platform devoted to assessing the attack robustness of LLM watermarkers in a unified and comprehensive manner. Leveraging WATERPARK, we conduct a systematic evaluation of the robustness of existing watermarkers, addressing unresolved questions, revealing design trade-offs, and identifying opportunities for further improvement. Our findings shed light on the current LLM watermarking techniques, while WATERPARK serves as a valuable benchmark aiding future research.

702

703

704

705

706

657

658

659

660

661

662

# Limitations

606

610

611

612

615

616

617

618

619

621

625

627

638

640

We now examine the limitations of our study and identify promising directions for future research.

Watermarkers. This study primarily focuses on training-free, pre-generation watermarking, which is applicable to any given LLM and offers flexible control over multiple criteria (e.g., quality, effectiveness, and robustness). While we have included peer-reviewed watermarking methods to the best of our ability, many methods, such as training-time watermarking and post-generation watermarking, are not covered in this paper.

Threat Model. Notably, our evaluation is based on the following assumptions about adversary capabilities: i) the adversary can modify the watermarked text in a computationally efficient manner (e.g., synonym substitution), ii) the adversary uses LLMs less capable than the target LLM, and/or iii) the adversary can train a detector using given watermarked and non-watermarked texts. We argue that these assumptions are realistic and practical. Otherwise, using more capable LLMs could easily generate high-quality, non-watermarked texts to evade detection, while launching computationally expensive attacks would significantly increase the adversary's cost. Future research directions include exploring alternative threat models to expand the scope of our analysis.

**Causal Analysis of Design Choices.** While our watermarker taxonomy provides a clear classification framework, we focus on analyzing the holistic design choices of each watermarked (rather than individual design modules) to elucidate the underlying factors driving our experimental observations. This limitation stems from the relatively small number of watermarkers within each taxonomic category, precluding definitive conclusions about the vulnerability of specific taxonomic elements based on current experimental evidence.

645A more granular assessment of specific design646choices would ideally involve comprehensive ab-647lation studies, systematically modifying individual648design elements and comparing their performance649against baseline configurations. However, this ap-650proach faces significant practical challenges due651to the intricate inter-dependencies and tight cou-652pling among watermarker components. To partially653address this challenge, in §4.3, we strategically654compare two watermarkers that differ only in the655design of one specific component (e.g., context de-656pendency), enabling the examination of the effects

of specific design variations. We consider a more systematic causal analysis as future research.

**Parameter Tuning.** Our comparative analysis employs AUC curves based on default parameter settings across watermarkers. For benchmarking purposes, we evaluate True Positive Rate (TPR) at a fixed False Positive Rate (FPR) of 1%, aligning with established practices in comparative studies. This standardized approach enables systematic assessment of watermark detection effectiveness while maintaining consistent false positive control. While this approach provides a pragmatic framework for comparative evaluation, it underscores the importance of future research into comprehensive parameter optimization. Such studies could reveal the full performance potential of each method and yield deeper insights into their relative strengths and limitations.

# **Ethics Consideration**

In this paper, we conduct a systematic study of state-of-the-art LLM watermarkers, focusing on their robustness against watermark removal attacks. We aim to understand how various design choices affect watermarkers' resilience and identify best practices for operating watermarkers in adversarial environments.

**Stakeholder Considerations.** The primary stakeholders affected by this research include users and developers of LLM watermarkers, as well as the broader community relying on these techniques. By identifying the strengths/limitations of various watermarkers, our work could influence the perceived reliability and deployment of LLM watermarkers in critical applications. Conversely, exposing these vulnerabilities allows for the improvement of the current techniques, ultimately contributing to more secure and robust systems.

**Potential Harms and Benefits.** Exposing vulnerabilities in widely adopted security mechanisms can yield contrasting outcomes. Initially, it may erode trust in LLM watermarkers as reliable safeguards, potentially leading to their underuse and leaving systems exposed to unchecked risks. However, by illuminating these weaknesses, our research catalyzes the development of more robust techniques and fosters a nuanced comprehension of their constraints, ultimately fortifying the longterm security ecosystem.

Future Research and Mitigation. Our study has identified several potential countermeasures to

812

813

address the vulnerabilities we uncovered. These
recommendations are designed to steer future research and assist developers in bolstering the security of LLM watermarkers.

#### 711 References

712

713

715

717

718

721

727

728

731

733

734

735

736

737

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

759

- Scott Aaronson and Hendrik Kirchner. Watermarking gpt outputs. https://www.scottaaronson.com/ talks/watermark.ppt.
- Sahar Abdelnabi and Mario Fritz. 2021. Adversarial watermarking transformer: Towards tracing text provenance with data hiding. In 2021 IEEE Symposium on Security and Privacy (SP), pages 121–140. IEEE.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (*FAccT*).
- Miranda Christ, Sam Gunn, and Or Zamir. 2023. Undetectable watermarks for language models. Cryptology ePrint Archive, Paper 2023/763.
- Seamless Communication, Loïc Barrault, Yu-An Chung, Mariano Cora Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, Christopher Klaiber, Pengwei Li, Daniel Licht, Jean Maillard, Alice Rakotoarison, Kaushik Ram Sadagopan, Guillaume Wenzek, Ethan Ye, Bapi Akula, Peng-Jen Chen, Naji El Hachem, Brian Ellis, Gabriel Mejia Gonzalez, Justin Haaheim, Prangthip Hansanti, Russ Howes, Bernie Huang, Min-Jae Hwang, Hirofumi Inaguma, Somya Jain, Elahe Kalbassi, Amanda Kallet, Ilia Kulikov, Janice Lam, Daniel Li, Xutai Ma, Ruslan Mavlyutov, Benjamin Peloquin, Mohamed Ramadan, Abinesh Ramakrishnan, Anna Sun, Kevin Tran, Tuan Tran, Igor Tufanov, Vish Vogeti, Carleigh Wood, Yilin Yang, Bokai Yu, Pierre Andrews, Can Balioglu, Marta R. Costa-jussà, Onur Celebi, Maha Elbayad, Cynthia Gao, Francisco Guzmán, Justine Kao, Ann Lee, Alexandre Mourachko, Juan Pino, Sravya Popuri, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, Paden Tomasello, Changhan Wang, Jeff Wang, and Skyler Wang. 2023. Seamlessm4t: Massively multilingual & multimodal machine translation.
- Compilatio. Cheating in the age of chatgpt: findings and solutions for preserving academic integrity. https: //www.compilatio.net/en/blog/cheating-cha tgpt.
  - Prithiviraj Damodaran. 2021. Parrot: Paraphrase generation for NLU.
  - Jacoband Devlin, Ming-Weiand Chang, Kentonand Lee, and Kristina Toutanova. 2019. BERT: Pre-training of

deep bidirectional transformers for language understanding. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL).* 

- Jaiden Fairoze, Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody, and Mingyuan Wang. 2023. Publicly detectable watermarking for language models. *ArXiv e-prints*.
- Jaiden Fairoze, Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody, and Mingyuan Wang. 2023. Publicly detectable watermarking for language models. *arXiv preprint arXiv:2310.18491*.
- Pierre Fernandez, Antoine Chaffin, Karim Tit, Vivien Chappelier, and Teddy Furon. 2023. Three bricks to consolidate watermarks for large language models. In 2023 IEEE International Workshop on Information Forensics and Security (WIFS), pages 1–6. IEEE.
- Eva Giboulot and Furon Teddy. 2024. Watermax: breaking the llm watermark detectability-robustnessquality trade-off. *arXiv preprint arXiv:2403.04808*.
- Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. 2021. Gradient-based adversarial attacks against text transformers. *ArXiv e-prints*.
- Abe Bohan Hou, Jingyu Zhang, Tianxing He, Yichen Wang, Yung-Sung Chuang, Hongwei Wang, Lingfeng Shen, Benjamin Van Durme, Daniel Khashabi, and Yulia Tsvetkov. 2023. Semstamp: A semantic watermark with paraphrastic robustness for text generation. *arXiv preprint arXiv:2310.03991*.
- Zhengmian Hu, Lichang Chen, Xidong Wu, Yihan Wu, Hongyang Zhang, and Heng Huang. 2024. Unbiased watermark for large language models. In *Proceedings of the International Conference on Learning Representations (ICLR).*
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. Atlas: Few-shot learning with retrieval augmented language models. *ArXiv e-prints*.
- Nikola Jovanović, Robin Staab, and Martin Vechev. 2024. Watermark stealing in large language models. *arXiv preprint arXiv:2402.19361*.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023a. A watermark for large language models. In *Proceedings of the IEEE Conference on Machine Learning* (*ICML*).
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Manli Shu, Khalid Saifullah, Kezhi Kong, Kasun Fernando, Aniruddha Saha, Micah Goldblum, and Tom Goldstein. 2023b. On the reliability of watermarks for large language models. *ArXiv e-prints*.
- Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. 2023. Paraphrasing evades detectors of ai-generated text, but retrieval

814

815

- 826 827 828 829 830 831 832 833
- 835 836 837 838 839 840 841 842 843 844 845 846
- 8
- 852 853 854
- 8
- 860 861
- 8

864 865 866

867 868

869 870 is an effective defense. In *Proceedings of the Advances in Neural Information Processing Systems* (*NeurIPS*).

- Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. 2023. Robust distortion-free watermarks for language models. *ArXiv e-prints*.
- John Kudo, Takuand Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP).
- Taehyun Lee, Seokhee Hong, Jaewoo Ahn, Ilgee Hong, Hwaran Lee, Sangdoo Yun, Jamin Shin, and Gunhee Kim. 2023. Who wrote this code? watermarking for code generation. *arXiv preprint arXiv:2305.15060*.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019. Textbugger: Generating adversarial text against real-world applications. In *Proceedings of the Network and Distributed System Security Symposium* (NDSS).
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Alexander Cosgrove, Christopher D Manning, Christopher Re, Diana Acosta-Navas, Drew Arad Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue WANG, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Andrew Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. Holistic evaluation of language models. Transactions on Machine Learning Research.
- Aiwei Liu, Leyi Pan, Xuming Hu, Shu'ang Li, Lijie Wen, Irwin King, and Philip S Yu. 2023a. An unforgeable publicly verifiable watermark for large language models. In *Proceedings of the International Conference on Learning Representations (ICLR).*
- Aiwei Liu, Leyi Pan, Xuming Hu, Shiao Meng, and Lijie Wen. 2024. A semantic invariant robust watermark for large language models. In *Proceedings of the International Conference on Learning Representations (ICLR).*
- Aiwei Liu, Leyi Pan, Yijian Lu, Jingjing Li, Xuming Hu, Lijie Wen, Irwin King, and Philip S Yu. 2023b.
  A survey of text watermarking in the era of large language models. arXiv preprint arXiv:2312.07913.
- Yixin Liu, Hongsheng Hu, Xuyun Zhang, and Lichao Sun. 2023c. Watermarking text data on large language models for dataset copyright protection. arXiv preprint arXiv:2305.13257.

Nils Lukas, Edward Jiang, Xinda Li, and Florian Kerschbaum. 2021. Sok: How robust is image classification deep neural network watermarking? In *Proceedings of the IEEE Symposium on Security and Privacy (S&P).*  871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918

919

920

921

922

- Ingo Lütkebohle. Gptzero. https://gptzero.me/.
- George A. Miller. 1995. Wordnet: a lexical database for english. *Commun. ACM*, 38(11):39–41.
- Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D. Manning, and Chelsea Finn. 2023. Detectgpt: Zero-shot machine-generated text detection using probability curvature.
- Travis Munyer and Xin Zhong. 2023. Deeptextmark: Deep learning based text watermarking for detection of large language model generated text. *arXiv preprint arXiv:2305.05773*.
- Openai. Openai chatgpt blog. https://openai.com /blog/chatgpt.
- Leyi Pan, Aiwei Liu, Zhiwei He, Zitian Gao, Xuandong Zhao, Yijian Lu, Binglin Zhou, Shuliang Liu, Xuming Hu, Lijie Wen, et al. 2024. Markllm: An opensource toolkit for llm watermarking. *arXiv preprint arXiv:2405.10051*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL).*
- Julien Piet, Chawin Sitawarin, Vivian Fang, Norman Mu, and David Wagner. 2023. Mark my words: Analyzing and evaluating language model watermarks. *ArXiv e-prints*.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. MAUVE: Measuring the gap between neural text and human text using divergence frontiers. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*.
- Jie Ren, Han Xu, Yiding Liu, Yingqian Cui, Shuaiqiang Wang, Dawei Yin, and Jiliang Tang. 2023. A robust semantics-based watermark for large language model against paraphrasing. *arXiv preprint arXiv:2311.08721*.
- Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. 2023. Can ai-generated text be reliably detected? *ArXiv e-prints*.
- Ryoma Sato, Yuki Takezawa, Han Bao, Kenta Niwa, and Makoto Yamada. 2023. Embarrassingly simple text watermarks. *arXiv preprint arXiv:2310.08920*.
- Karanpartap Singh and James Zou. 2023. New evaluation metrics capture quality degradation due to llm watermarking. *arXiv preprint arXiv:2312.02382*.

- 924 925 931 932 934 937 940 941 942 943 945 947 949 950 951 953 954 957

- 960 961 962
- 963

- 968 969
- 970
- 974 975
- 976 977
- 978 979
- 982

- Zhensu Sun, Xiaoning Du, Fu Song, and Li Li. 2023. Codemark: Imperceptible watermarking for code datasets against neural code completion models. In Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, pages 1561–
- 1572. Zhensu Sun, Xiaoning Du, Fu Song, Mingze Ni, and Li Li. 2022. Coprotector: Protect open-source code against unauthorized training usage with data poisoning. In Proceedings of the ACM Web Conference
- 2022, pages 652-660. Ruixiang Tang, Qizhang Feng, Ninghao Liu, Fan Yang, and Xia Hu. 2023. Did you train on my dataset? towards public dataset protection with cleanlabel backdoor watermarking. ACM SIGKDD Explorations

Newsletter, 25(1):43-53.

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and finetuned chat models. ArXiv e-prints.
  - Shangqing Tu, Yuliang Sun, Yushi Bai, Jifan Yu, Lei Hou, and Juanzi Li. 2023. Waterbench: Towards holistic evaluation of watermarks for large language models. arXiv preprint arXiv:2311.07138.
  - Lean Wang, Wenkai Yang, Deli Chen, Hao Zhou, Yankai Lin, Fandong Meng, Jie Zhou, and Xu Sun. 2024. Towards codable text watermarking for large language models. In Proceedings of the International Conference on Learning Representations (ICLR).
  - Johnand Wieting, Kevinand Gimpel, Grahamand Neubig, and Taylor Berg-kirkpatrick. 2022. Paraphrastic representations at scale. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP).
  - Yihan Wu, Zhengmian Hu, Junfeng Guo, Hongyang Zhang, and Heng Huang. 2024. A resilient and accessible distribution-preserving watermark for large language models. In Forty-first International Conference on Machine Learning.

Xiaojun Xu, Yuanshun Yao, and Yang Liu. 2024. Learning to watermark llm-generated text via reinforcement learning. arXiv preprint arXiv:2403.10553.

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1008

1009

1010

1011

1012

1014

1015

- Xi Yang, Kejiang Chen, Weiming Zhang, Chang Liu, Yuang Qi, Jie Zhang, Han Fang, and Nenghai Yu. 2023. Watermarking text generated by black-box language models. arXiv preprint arXiv:2305.08883.
- KiYoon Yoo, Wonhyuk Ahn, and Nojun Kwak. 2023. Advancing beyond identification: Multi-bit watermark for language models. arXiv preprint arXiv:2308.00221.
- KiYoon Yoo, Wonhyuk Ahn, and Nojun Kwak. 2024. Advancing beyond identification: Multi-bit watermark for language models. In Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).
- Hanlin Zhang, Benjamin L. Edelman, Danilo Francati, Daniele Venturi, Giuseppe Ateniese, and Boaz Barak. 2023. Watermarks in the sand: Impossibility of strong watermarking for generative models. ArXiv e-prints.
- Ruisi Zhang, Shehzeen Samarah Hussain, Paarth Neekhara, and Farinaz Koushanfar. 2023. Remarkllm: A robust and efficient watermarking framework for generative large language models. ArXiv e-prints.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *Proceedings of* the International Conference on Learning Representations (ICLR).
- Xuandong Zhao, Prabhanjan Ananth, Lei Li, and Yu-Xiang Wang. 2023. Provable robust watermarking for ai-generated text. ArXiv e-prints.
- Xuandong Zhao, Sam Gunn, Miranda Christ, Jaiden 1017 Fairoze, Andres Fabrega, Nicholas Carlini, Sanjam 1018 Garg, Sanghyun Hong, Milad Nasr, Florian Tramer, 1019 et al. 2024. Sok: Watermarking for ai-generated 1020 content. arXiv preprint arXiv:2411.18479. 1021

1028

1029

1030

1031

1032

1033

1034

1035

1038

1039

1040

1041

1042

1043

1046

1047

1048

1049

1050

1052

# A Parameter Setting

Table 5 lists the default setting of the parameters of each watermarker in our evaluation. Note that  $\gamma$  and  $\delta$  are the gamma and delta used in the watermarker, and n denotes the window size.

Watermarker	Parameter	Setting
TCPI	$\gamma$	0.25
TOKL	δ	2.0
UG	$\gamma$	0.5
00	δ	2.0
	δ	1.5
	n	10
CTWL	message code length	20
	encode ratio	10.0
	message strategy	vanilla
	$\gamma$	0.5
	δ	2.0
UPV	n	3
	bit number	16
	layers	9
UB	watermark type	delta
	δ	1.0
SIR	n	10
	watermark type	context
GO	n	3
RDF	number of random sequences	50

Table 5: Default parameter setting of watermarkers.

# **B** Taxonomies

We present a taxonomy of LLM watermarkers, as summarized in Table 6. The current watermarkers can be categorized based on the information carried by the watermarks (e.g., one-bit versus multibit) and their key design choices, including context dependency, generation strategy, and detection method. Next, we mainly focus on the key design factors.

# **B.1** Context Dependency

The watermarker applies a perturbation  $\Delta_t$  to the LLM's next-token distribution  $p(x_t|x_{< t})$  as  $\tilde{p}(x_t|x_{< t}) = p(x_t|x_{< t}) + \Delta_t$ , from which the next token  $x_t$  is generated. The perturbation  $\Delta_t$  often depends on the given context. Three types of context dependencies are typically used in the existing watermarkers.

**Index-dependent watermark.** The watermarker produces a pseudo-random number  $r_t$  by applying a keyed hash function

$$r_t = f_k(t), \tag{2}$$

that only depends on the index t of the next token;  $r_t$  is then used to generate the perturbation  $\Delta_t$  (Kuditipudi et al., 2023).

**Text-dependent watermark.** The watermarker considers the previous tokens  $x_{<t}$  as the context

window to generate  $\Delta_t$  (Kirchenbauer et al., 2023a; Aaronson and Kirchner; Hu et al., 2024). For instance, GO (Aaronson and Kirchner) computes the hash of the concatenation of previous w tokens as  $r_t$ :

$$t = f_k(x_{t-w} \| \dots \| x_{t-1}), \qquad (3)$$

1053

1054

1055

1056

1057

1058

1060

1061

1062

1063

1064

1065

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1093

1094

1095

1096

1097

1098

1099

1100

1101

while TGRL (Kirchenbauer et al., 2023a) suggests using the min hash of the previous token:

$$r_t = \min(f_k(x_{t-w}), \dots, f_k(x_{t-1}))$$
(4)

Similarly, UB (Hu et al., 2024) concatenates the previous w tokens to generate the context code for reweighting the logits. UPV (Liu et al., 2023a) and SIR (Liu et al., 2024) use neural networks to generate  $\Delta_t$  based on the previous tokens.

**Context-free watermark.** The watermarker applies a universal perturbation  $\Delta_t$  across the next-token distributions of all the tokens without considering their contexts (Zhao et al., 2023).

#### **B.2** Generation Strategy

r

By applying the perturbation  $\Delta_t$  to the LLM's nexttoken distribution  $p(x_t|x_{< t})$ , the watermarker generates the new distribution  $\tilde{p}(x_t|x_{< t})$  to sample the next token  $x_t$ . The existing sampling strategies can be categorized as follows.

Distribution shift. TGRL (Kirchenbauer et al., 2023a) modifies  $p(x_t|x_{< t})$  by adding a shift  $\delta$  to the logits of "green-list" tokens (the remaining as "red-list" tokens) as the modified distribution  $\tilde{p}(x_t|x_{\leq t})$ . A token x is considered as green listed if  $\pi_{r_t}(x) < \gamma d$  where  $\pi_{r_t}$  is a permutation seeded by  $r_t$ ,  $\gamma$  is a parameter to control the size of the green list, and d is the number of vocabulary size. Similarly, UG (Zhao et al., 2023) uses a fixed redgreen split over the vocabulary, showing greater robustness than TGRL against edit-distance-bounded attacks due to its "hard" split. UPV (Liu et al., 2023a) selects the top-k tokens from  $p(x_t|x_{< t})$ , applies a neural network to predict the green-list tokens from these tokens, and adds  $\delta$  to the logits of green-list tokens. Unlike the other strategies, the distribution-shift strategy preserves the diversity of generated tokens; however, it can not be made indistinguishable, as  $\tilde{p}(x_t|x_{< t})$  and  $p(x_t|x_{< t})$  are inherently distinguishable.

**Distribution reweight.** Similar to distribution shift, this strategy alters the next-token distribution but uniquely perturbs the logit of each token. For instance, SIR (Liu et al., 2024) trains a neural network to predict the perturbation to logit of



Table 6: A taxonomy of LLM watermarkers. References: TGRL (Kirchenbauer et al., 2023a), UG (Zhao et al., 2023), UPV (Liu et al., 2023a), SIR (Liu et al., 2024), RDF (Kuditipudi et al., 2023), UB (Hu et al., 2024), DIP (Wu et al., 2024), GO (Aaronson and Kirchner), CTWL (Wang et al., 2024), MPAC (Yoo et al., 2024).

each token, which is unbiased (no preference over specific tokens) and balanced (with total perturbation summing up to 0). UB (Hu et al., 2024) advocates unbiased watermarking such that the expectation of the reweighted distribution agrees with the original distribution. It proposes two reweighting schemes:  $\delta$ -reweighting uniformly samples a token from  $p(x_t|x_{< t})$  and changes its probability to 1;  $\gamma$ -reweighting shuffles all the tokens, rejects the first half, and double the probabilities of the remaining half. DIP (Wu et al., 2024) also uses a distribution-reweight generation strategy similar to UB but does not need to access the LM during detection.

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

**Distribution transform.** Another line of watermarkers apply randomized transform on  $p(x_t|x_{< t})$ to sample  $x_t$ . For instance, RDF-EXP (Kuditipudi et al., 2023) and GO (Aaronson and Kirchner) use the Gumbel-max trick and apply the exponential transform. Let  $p(x_t|x_{< t}) = \{p_i\}_{i=1}^d$  be the distribution over the next token  $x_t$ . Then  $x_t$  is sampled as:

$$x_t = \arg\max_{i < d} r_i^{1/p_i} \tag{5}$$

where  $r_i$  is generated by the pseudo-random functions in Eq. 2 and Eq. 3. Similarly, RDF-IST (Kuditipudi et al., 2023) applies inverse transform over  $p(x_t|x_{< t})$ . With  $\pi_k$  as a random permutation seeded by the secret key k, the next token  $x_t$  is selected as:

$$x_t = \pi_k \left( \min_{j \le d} \sum_{i=1}^j p_{\pi_k(i)} \ge r_t \right) \tag{6}$$

1132which is the smallest index in the inverse permu-1133tation such that the CDF of the next token dis-1134tribution exceeds  $r_t$ . The distribution-transform1135strategy does not alter the next-token distribution,1136thus preserving the original text distribution (i.e.,1137indistinguishability).

#### **B.3** Detection Method

The watermarker's detector determines whether a given text  $T = (x_1, \ldots, x_n)$  is watermarked or not. Next, we categorize the existing detection methods as follows. 1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1158

1159

1160

1161

1164

**Score-based detection.** The detector computes the random value  $r_i$  for each position, the per-token statistics  $s(x_i, r_i)$ , and a score over such statistics, which is then subjected to a one-tailed statistical test to determine whether the text is watermarked. The per-token statistics vary with the concrete generators. For instance, GO (Aaronson and Kirchner) defines  $s(x_i, r_i) = -\log(1 - h_{r_i}(x_i))$ , while TGRL (Kirchenbauer et al., 2023a) defines  $s(x_i, r_i) = 1$  if  $x_i$  is in the green list and 0 otherwise. One simple way to aggregate the per-token statistics is to compute their sum (Aaronson and Kirchner; Kirchenbauer et al., 2023a; Zhao et al., 2023; Liu et al., 2024):

$$S = \sum_{i=1}^{n} s(x_i, r_i)$$
(7) 1157

However, as the random values may be misaligned with the tokens (e.g., due to editing), a more robust way is to compute the alignment score (e.g., edit score (Kuditipudi et al., 2023; Wang et al., 2024)):

$$S = s^{\psi}(n,n)$$
 1162

s.t. 
$$s^{\psi}(i,j) = \min \begin{cases} s^{\psi}(i-1,j-1) + s(x_i,r_j), \\ s^{\psi}(i,j-1) + \psi, \\ s^{\psi}(i-1,j) + \psi \end{cases}$$
(8)

where  $\psi$  is the "edit-cost" parameter.

Differential-based detection. This line of de-<br/>tectors also relies on the score of a given text. How-<br/>ever, the score is computed by comparing the given11651167<br/>text with the non-watermarked text generated by1167

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1209

1210

1211

1212 1213

1214

1215

1216

1217

the same LLM. For instance, UB (Hu et al., 2024) computes the log-likelihood ratio (LLR) score:

$$s(i) = \log \frac{\tilde{p}(x_i|x_{< i})}{p(x_i|x_{< i})} \tag{9}$$

and its more robust maximin variant. However, note that these detectors naturally require accessing the original LLM, which is not always feasible.

Model-assisted detection. Instead of computing the per-token statistics, which are often subject to watermark removal attacks, one may also train a model to predict whether the given text is watermarked. UPV (Liu et al., 2023a) trains a neural network, which shares the same embedding layers with its generator, to detect watermarked texts. Similarly, one may develop a generic detector by training it to distinguish watermarked and non-watermarked texts. We explore this option in §4.

#### С **Fidelity Metrics**

In this paper, to evaluate the impact of watermarking on text quality, we employs the following metrics to measure the difference between the original text T and the watermarked text T. Below, we provide a detailed description of each metric.

WER (word error rate) measures the percentage of mismatched tokens between  $\tilde{T}$  and T relative to the total number of tokens in T. This is a lexical metric used to measure the fraction of tokens that have been modified.

BLEU (Papineni et al., 2002) measures the lexical similarity of T and  $\tilde{T}$  by calculating the proportion of n-grams matched in T and T.

BERTScore (Zhang et al., 2020) measures the token-level similarity of T and T by leveraging the pre-trained contextual embeddings from BERT (Devlin et al., 2019) and matching tokens in T and T using cosine similarity.

P-SP (Wieting et al., 2022) evaluates the semantic similarity of T and T using the cosine similarity of their encodings. In particular, the encoding of T (or T) is calculated by averaging the embeddings of its subword units generated by Sentence-Piece (Kudo, 2018).

MAUVE (Pillutla et al., 2021) compares the distributions of T and T by computing an information divergence curve within a quantized embedding space. The area under this divergence curve provides a scalar summary of the trade-off between Type I (T places high mass in areas where T has low mass) and Type II (vice versa) errors.

Note that these metrics are also used to mea-1218 sure the impact of watermark removal attacks on the quality of the modified text  $\tilde{T}'$ , relative to the watermarked text T. 1221

#### D **Details of Watermark Removal Attacks**

We also present a taxonomy of existing watermark removal attacks according to their underlying perturbation and required resources, as summarized in Table 2.

Linguistic variation attack. This class of attacks perturb the linguistic features of the watermarked text, without changing its semantics. We consider the set of perturbations in HELM (Liang et al., 2023), which simulate natural linguistic variations encountered in human interactions with text typing interfaces: i) Lowercasing converts all the words to lower-cases, which potentially affects the interpretation of proper nouns or emphases. ii) Contracting/expanding replaces phrases with their contracted/expanded forms (e.g., "I am" to "I'm" and vice versa), which may impact the tokenizer.

Lexical editing attack. This class of attacks modifies individual words, aiming to maintain the original text's semantics. Specifically, we consider the following editing operations: i) *Misspelling*, similar to text-bugger (Li et al., 2019), replaces words with their common misspellings (plural forms also considered); ii) Typoing replaces certain letters in a word with others; iii) Synonymizing replaces words with their synonyms using Word-Net (Miller, 1995); and iv) *Swapping* randomly exchanges the positions of two words within the text, which alters the text structure while potentially preserving the overall semantics.

Text-mixing attack. This class of attacks aims to "dilute" the watermark by mixing the watermarked text with non-watermarked text fragments. Specifically, the copy-pasting attack (Kirchenbauer et al., 2023b) embeds the watermarked text into the context of non-watermarked, human-written text. Note that the influence of non-watermarked text can be controlled by setting the fractions of watermarked and non-watermarked text fragments (i.e., the mixing weights).

Paraphrasing attack. This class of attacks relies on an additional LLM (i.e., paraphraser) to re-write the given watermarked text to evade the detector. For instance, a light paraphraser (Sankar Sadasivan et al., 2023) (e.g., T5-based paraphraser (Damodaran, 2021)) can paraphrase the

1219 1220

1222

1223 1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

1242

1243

1244

1245

1246

1247

1248

1249

1251

1252

1253

1254

1255

1256

1257

1259

1260

1261

1262

1263

1264

1265

1304

1305

1306

1308

1309

1310

1312

1313

1314

1315

1316

1268

watermarked text sentence-by-sentence, while a more capable paraphraser (e.g., DIPPER (Krishna et al., 2023)) paraphrases the watermarked text in one-shot, also enabling to control lexical diversity and token order diversity.

Similar to paraphrasing, the *translating* attack uses a translator LLM (e.g., Seamless (Communication et al., 2023)) to cycle the watermarked text through multiple languages (e.g., from English to French and back to English). This process can significantly alter the sentence structure and phrasing.

In this class of attacks, we assume the adversary has access to a generic detector that is trained to distinguish watermarked and non-watermarked texts. The adversary then perturbs the watermarked text based on this surrogate detector. We explore this attack type in §5.

# E Watermark Effectiveness

We first evaluate the effectiveness of different watermarkers. Following prior work (Kirchenbauer et al., 2023a; Piet et al., 2023; Kuditipudi et al., 2023), for each watermarker, we sample 1,000 prompts and use the LLM in combination with the watermarker to generate the watermarked texts; meanwhile, we select human responses to the same prompts as the non-watermarked texts. We then measure the accuracy of the watermarker's detector in distinguishing the watermarked and nonwatermarked texts.

# E.1 Overall Effectiveness

We measure the overall effectiveness of each method through the lens of the ROC curve. Figure 10 (a-d) summarizes the overall effectiveness of existing watermarkers across different models and datasets. We have the following interesting observations.

Most watermarkers are highly effective in generating and subsequently detecting watermarked texts on OPT-1.3B as shown in Figure 10 (a-b). For instance, RDF, UB, and GO all attain AUC scores above 0.99 over both C4 and HC3. Recall that C4 and HC3 represent the text completion and question-answering tasks respectively. The observation indicates that most watermarkers tend to be highly effective for relatively less capable LLMs such as OPT-1.3B, while the concrete dataset/task have a limited impact on their performance. We further validate this hypothesis under the setting of a fixed FPR. As shown in Figure 10 (e-f), we fix the FPRs of all the methods to be 0.01 and measure their TPRs. Observe that all the methods achieve above 0.9 TPR, with a marginal difference across C4 and HC3.

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350

1351

1352

1353

1354

1355

1356

1357

1359

1360

1361

1362

1363

1364

1365

1367

Meanwhile, most methods observe marginal performance drops on Llama2-7B, as shown in Figure 10 (c-d). For instance, compared with its performance on OPT-1.3B, the AUC of SIR drops by 0.11 and 0.08 on C4 and HC3 respectively. This observation aligns with previous research (Kuditipudi et al., 2023), indicating that watermarkers are more effective on OPT compared with Llama2. This phenomenon can be partly understood through the following explanation. In contrast of less capable LLMs (e.g., OPT-1.3B), Llama2-7B typically produces texts of lower perplexity. Since most watermarkers inject watermarks by slightly altering the next-token distribution, the lower perplexity in Llama2's outputs hampers the effectiveness of such perturbations. Moreover, it is observed that the performance of various methods varies significantly across different datasets. For instance, UG's AUC differs by 0.28 between C4 and HC3, while UB's AUC differs by 0.11. This observation is further supported by the TPR measures at fixed FPRs (fixed as 1%), as shown in Figure 10 (e-f). Our findings suggest that the concrete dataset/task tends to have a larger impact on watermarkers over more capable LLMs.

# E.2 Impact of Text Length

We evaluate how the (non-)watermarked text length (i.e., the number of tokens) impacts the performance of different methods. Specifically, we measure the TPR of each method with its FPR fixed as 1%. In the following, we set OPT-1.3B and C4 as the default LLM and dataset. Figure 11 summarizes the results.

Observe that as expected, the TPRs of all the methods improve as the text length grows from 1 to 200 tokens. As the text length exceeds 100 tokens, most methods reach TPRs close to 100%. Meanwhile, different methods show varying sensitivity to the text length. For instance, GO and RDF attain 100% TPRs with only 20 tokens, while UG reaches only around 50% TPR under the same setting. This can be explained as follows. Both GO and RDF use distribution transform-based samplers, which, conditional on given randomness (e.g., random permutation), generate the next token deterministically. Meanwhile, other methods randomly sample the next token from a given pool (e.g., green lists).



Figure 10: Overall effectiveness of different watermarkers in generating and detecting watermarked texts: ROC of (a) OPT-C4, (b) OPT-HC3, (c) Llama2-C4, and (d) Llama2-HC3; TPR (with FPR fixed as 0.01) on (e) C4 and (f) HC3.



Figure 11: TPRs of watermarkers with respect to text length (with FPRs fixed as 1%).

Thus, GO and RDF tend to have stronger signals per token for watermark detection.

#### E.3 Impact of Temperature

The temperature  $\tau$  is a key parameter that affects a watermarker's generative dynamics: intuitively, a higher  $\tau$  makes the sampling over the next-token distribution  $\tilde{p}(x_t|x_{< t})$  more random. Here we evaluate how the setting of  $\tau$  in each watermaker's generator may impact its effectiveness. Note that, unlike other watermarkers, RDF and GO do not generate the next-token distribution  $\tilde{p}(x_t|x_{< t})$  explicitly, we thus exclude them from the evaluation.



Figure 12: TPRs of watermarkers with respect to the temperature setting (with FPRs fixed as 1%).

Figure 12 compares how the TPRs of different watermarkers vary with the setting of  $\tau = 0.7, 1.0$ , and 1.3 (with FPRs fixed as 1%). Observe that the

performance of most watermarkers marginally improves with  $\tau$ , which corroborates prior work (Piet et al., 2023). For instance, the TPR of TGRL increases by about 0.05 as  $\tau$  varies from 0.7 to 1.3. Interestingly, in contrast, the TPR of UPV decreases as  $\tau$  grows. This can be explained as follows. UPV employs a neural network as the detector that depends on general textual features, which tends to be more sensitive to increasing randomness, compared with other watermarkers that rely on specific watermark signals (e.g., green/red-listed tokens).

### F Watermark Fidelity



ers.

We evaluate the impact of different watermarkers on the text quality. We compare the original text T and watermarked text  $\tilde{T}$  using the metrics (detailed in §3.2) of BERTScore (Zhang et al., 2020), P-SP (Wieting et al., 2022), and MAUVE (Pillutla et al., 2021). The results are summarized in Figure 13.

i) A majority of watermarkers well preserve the semantics of original texts, as indicated by their high BERTScore and MAUVE scores. Note that the P-SP scores of all the watermarkers are relatively lower than their BERTScore and MAUVE scores. This is due to their different emphases: P-SP measures the average similarity between the 1408 tokens in T and T, while BERTScore calculates the 1409 maximum similarity between the tokens in T and T. 1410 ii) Meanwhile, RDF, GO, and UB are less effective 1411 in preserving the quality of original texts, which 1412 can be attributed to their additional constraints of 1413 indistinguishability: the expectation of the water-1414 marker's next-token distribution is identical to the 1415 LLM's next-token distribution (i.e., indistinguisha-1416 bility). This observation suggest that there exists an 1417 inherent trade-off between the desiderata of quality 1418 and indistinguishability. 1419

## G Additional Results

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

# G.1 Multi-bit Watermarking

While our study focuses on one-bit watermarkers, for completeness, we also evaluate CTWL (Wang et al., 2024), a multi-bit watermarker. In contrast of one-bit watermarkers that encode only a single bit of information (i.e., whether a given text is watermarked), a multi-bit watermarker can encode multiple bits of information into the watermarked text, such as the generating model, the date of generation, and other details. However, despite its larger information capacity, we find that a multi-bit watermarker is typically less robust compared to one-bit watermarkers, as illustrated in Figure 14 and 15.



Figure 14: TPRs of CTWL (with FPRs fixed as 0.01) on different LLMs (OPT and Llama2) and datasets (C4 and HC3).



Figure 15: TPRs of CTWL (with FPRs fixed as 0.01) against various attacks.

The test results in the basic encoding and detection scenario highlight the sensitivity of CTWL to different language models. Although it performs 1437 well with the OPT model across both C4 and HC3 1438 datasets, its TPR on Llama2 is extremely low, and 1439 it becomes completely ineffective when tested on 1440 Llama2 using the HC3 dataset. This is due to the 1441 lower model perplexity of Llama2. We conduct 1442 additional experiments to compare the model per-1443 plexity on the same WikiText dataset, and the re-1444 sults show that Llama2 has a perplexity about 6.15, 1445 while OPT has about 12.43. Lower model per-1446 plexity results to larger fluctuations in the logits 1447 produced by the model, makes it more challeng-1448 ing for watermark injection (e.g., increase smaller 1449 logits to exceed larger ones). 1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

When facing the attacks, CTWL is more vulnerable than one-bit methods. It is particularly vulnerable to the copy-pasting attack, which can nearly disable the method as the TPR drops to near zero. Additionally, CTWL is highly susceptible to typoing, swapping, translating, lowercasing, and Dipper attacks, which generally do not affect many one-bit methods as severely. Despite the capacity of multi-bit methods to embed more information, their high sensitivity to language model variations and various attacks is a crucial limitation that needs to be addressed in future research.

### G.2 Robustness of Watermarkers on Llama2

Our evaluation in §4 mainly uses OPT-1.3B as the underlying LLM. Here, we present the results on Llama2-7B, a more capable LLM. The results can be found at Table 7

#### G.2.1 Linguistic Variation Attack

Comparing with the OPT results, most of the watermarkers show TPR drops on Llama2. One exception is GO, which still maintains resilience across all three attacks.

### G.2.2 Lexical Editing Attack

We have the following observation in Llama2: i) 1474 Context-free watermarkers (e.g., UG) show signif-1475 icantly higher attack resilience against swapping 1476 attacks. This may be due to the fact that the at-1477 tack alters the word and sentence order, which 1478 affects the detection of both text-dependent and 1479 index-dependent watermarkers. ii) Similar to OPT, 1480 text-dependent, (hard) distribution-reweighting wa-1481 termarkers tend to be more vulnerable to lexical 1482 editing attacks. iii) The typoing and swapping at-1483 tacks show significantly higher attack effectiveness 1484 on Llama2, similar to our observations on OPT. 1485

Water	CLEAN	AN Linguistic Variation		Textual Integrity			Content Manipulation				Paraphrase				
marker		Contra	Expan	LowCase	Swap	Туро	Syno	Missp	CP1-10	CP3-10	CP1-25	CP3-25	DP-20	DP-40	Trans
TGRL	0.896	0.864	0.892	0.75	0.408	0.34	0.818	0.78	0.014	0.074	0.088	0.348	0.536	0.334	0.196
UG	0.814	0.791	0.823	0.697	0.843	0.357	0.75	0.755	0.025	0.054	0.050	0.190	0.771	0.633	0.344
UPV	0.588	0.502	0.598	0.494	0.388	0.018	0.512	0.392	0.005	0.002	0.002	0.007	0.362	0.355	0.042
RDF	0.790	0.735	0.725	0.622	0.333	0.888	0.738	0.824	0.021	0.045	0.045	0.307	0.471	0.303	0.010
UB	0.946	0.926	0.940	0.512	0.016	0.030	0.614	0.650	0.000	0.000	0.000	0.022	0.143	0.036	0.094
SIR	0.566	0.066	0.078	0.078	0.076	0.010	0.064	0.050	0.034	0.026	0.042	0.048	0.040	0.081	0.100
GO	0.996	0.996	0.996	0.959	0.643	0.556	0.989	0.987	0.058	0.228	0.286	0.825	0.900	0.623	0.814

Table 7: Resilience of LLM watermarkers against attacks on Llama2.

#### G.2.3 Text-mixing Attacks

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1503

1505

With the exception of watermarkers based on distribution transform (e.g., RDF and GO), which shows greater resilience to text-mixing attacks, the other methods are virtually undetectable, which is consistent with the observations in OPT

#### G.2.4 Paraphrasing Attack

All methods show a significant TPR decrease, only GO presents a remarkable resilience. RDF performs well in OPT and does not present a significant resilience on Llama2.

#### G.3 Fidelity Preservation of attacks

Figure 16 illustrates the quality preservation of different attacks on GO, with similar results on other watermarkers,



Figure 16: Quality preservation of different attacks on GO.

# H Additional Related Work

We survey the relevant literature in the following categories: i) detection of LLM-generated texts, ii) LLM watermarking, iii) attacks on LLM watermarking, and iv) evaluation of LLM watermarkers.

**Detection of LLM-generated texts.** The advances in LLMs also give rise to their possible misuses. There is thus a pressing need for the capability of distinguishing LLM- and human-generated texts. Initial work attempts to either train classifiers using LLM- and human-generated texts(Mitchell et al., 2023) or to leverage intrinsic characteristics of LLM-generated texts (e.g., perplexity and variability in length, complexity, and information density)(Lütkebohle). Yet, with LLMs becoming increasingly capable, the difference between LLM-and human-generated texts is narrowing, making such approaches less effective.

1506

1507

1508

1509

1510

1511

1512

1513

1514

1515

1516

1517

1518

LLM watermarking In response, LLM water-1519 marking emerges as a promising alternative, which 1520 instruments the LLM generative process with sta-1521 tistical signals that can be subsequently detected. 1522 The existing LLM watermarking techniques can 1523 be categorized based on the stages in which they 1524 are applied(Liu et al., 2023b): i) training-time wa-1525 termarking(Liu et al., 2023c; Tang et al., 2023; 1526 Sun et al., 2022, 2023; Xu et al., 2024), ii) wa-1527 termarking during logit-generation(Kirchenbauer 1528 et al., 2023a; Zhao et al., 2023; Hu et al., 2024; Liu 1529 et al., 2023a; Liu et al., 2024; Fairoze et al., 2023; 1530 Ren et al., 2023; Fernandez et al., 2023; Wang et al., 1531 2024; Yoo et al., 2023; Lee et al., 2023; Giboulot 1532 and Teddy, 2024), iii) watermarking during token-1533 sampling (Kuditipudi et al., 2023; Aaronson and 1534 Kirchner; Hou et al., 2023), and iv) post-generation 1535 watermarking(Zhang et al., 2023; Yoo et al., 2024; 1536 Yang et al., 2023; Munyer and Zhong, 2023; Sato 1537 et al., 2023; Abdelnabi and Fritz, 2021). This study 1538 mainly focuses on training-free, pre-generation wa-1539 termarking, which applies to any given LLMs and 1540 provides flexible control over multiple criteria (e.g., 1541 quality, effectiveness, and robustness). 1542

The primary focus of the previous studies is to1543distinguish between human-written text and text1544generated by LLMs. In this paper, we assume that1545

the attacker lacks access to non-watermarked text produced by recent LLM. Instead, the attacker can only utilize other LLMs to generate text that mimics human responses and get the watermarked texts from recent LLM. The generic detector, which is within the adversary's capability, is specifically designed to distinguish between watermarked and non-watermarked texts(e.g., other LLMs' texts or human writing texts).

1546

1547

1548

1549

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1562

1563

1564

1565

1566

1568

1569

1570

1572

1574

1575

1576

1577

1578

1580

1581

1582

1583

1584

1585

1586

1587

1589

1590

1592 1593

1594

1595

1597

Attacks on LLM watermarking. One critical property of an LLM watermarker is its robustness against potential attacks. A variety of attacks can be applied to LLM watermarking, ranging from removing the embedded watermark to uncovering the green/red lists. For instance, Dipper(Krishna et al., 2023) is a widely used paraphrasing attack to evaluate the robustness of LLM watermarkers (Kirchenbauer et al., 2023a; Zhao et al., 2023) against watermark removal attacks; the watermark stealing attack(Jovanović et al., 2024) is proposed to identify green-list tokens and to replace them with redlist tokens, targeting TGRL(Kirchenbauer et al., 2023a) and UG(Zhao et al., 2023) to further launch spoofing or removal attacks. This study primarily focuses on watermark removal attacks as they can target any LLM watermarkers and have profound implications in practice (e.g., disinformation, academic cheating, and automated phishing).

Evaluation of LLM watermarkers. As LLM watermarkers become prevalent, recent work attempts to benchmark the performance of various watermarkers. However, current studies either focus solely on the effectiveness of watermarking or have limited assessments of robustness. For instance, WaterBench(Tu et al., 2023) compares the effectiveness of TGRL(Kirchenbauer et al., 2023a) and UG(Zhao et al., 2023) under varying hyper-parameter settings (e.g., prompt length); LLM-judger(Singh and Zou, 2023) uses GPT-3.5-Turbo as a judge to evaluate the effectiveness of RDF(Kuditipudi et al., 2023) and TGRL(Kirchenbauer et al., 2023a) and employs a binary classifier based on MLP to distinguish between watermarked and non-watermarked texts; MarkMyWords(Piet et al., 2023) compares the effectiveness of four watermarkers including TGRL(Kirchenbauer et al., 2023a), GO(Aaronson and Kirchner), RDF(Kuditipudi et al., 2023), and UW(Christ et al., 2023) and only evaluates the robustness of TGRL against watermark removal attacks.

To our best knowledge, this is the first study

dedicated to the robustness of LLM watermarkers1598against watermark removal attacks. We aim to1599understand how different design choices impact1600watermarkers' attack robustness and to identify best1601practices for operating watermarkers in adversarial1602environments.1603

1605

1606

1608

1611

1612

1613

1614

1615

1616

1617

1618

1619

1620

1621

1622

1623

1624

1627

1628

1629

1631

1632

1633

1634

1635

1636

1637

1639

1640

1641

# I Evaluation Guidelines for Future LLM Watermarking Research

Next, we propose a set of guidelines for evaluating the robustness of LLM watermarkers. These guidelines incorporate our findings in §4 and §5, providing a minimal checklist to claim the robustness of an LLM watermarker.

LLMs and tasks. Our experiments show that the referenced watermarkers show varying robustness across different LLMs and datasets. We speculate that there exists an intricate interplay between the watermarking mechanism, the LLM's capability, and the task's complexity. We thus recommend experimenting on i) LLMs with varying capability (e.g., as measured by perplexity), ii) datasets for different tasks (e.g., summarization and questionanswering), and iii) their combinations.

Attacks. Notably, using highly capable LLMs or applying computationally expensive rewriting can easily generate highly-quality, non-watermarked texts to evade detection; however, such attacks negate the need for watermark removal attacks in the first place. We thus recommend focusing on computationally efficient attacks such as linguistic variation, lexical editing, and lightweight paraphrasing, as well as their combinations, which reflects the risks of watermark removal attacks in practical settings.

**Robustness.** It is often critical to properly set the decision threshold for a watermarker (and also the attacks) to fully assess its robustness(Lukas et al., 2021), which unfortunately is often missing in the original papers. To overcome this issue, we recommend i) measuring the overall effectiveness (TPR) in terms of ROC (across different threshold settings), ii) measuring the TPR under a fixed FPR (e.g., 0.01), and iii) considering varying attack intensity.

Fidelity. It is notoriously challenging to mean-<br/>ingfully measure the quality of text data(Pillutla<br/>et al., 2021). We recommend employing a variety<br/>of metrics (e.g., BERTScore, P-SP, MAUVE) to<br/>comprehensively measure the quality retention of<br/>watermarkers as well as attacks. In addition, one1642<br/>1643

1648	may also leverage external tools (e.g., more ad-
1649	vanced LLMs such as GPT-4) or human evaluation
1650	to provide a more accurate assessment if feasible.