

429 **A Technical Appendices and Supplementary Material**

430 **A.1 Complete Message Template Specifications**

Table 3: Complete Message Template Specifications

Name	Message	Chars	Generation Strategy
cognitiveUltra	"Dr. Kristen Johnson: NEW Rx - complete your visit today"	58	
autonomyMax	"From Dr. Kristen Johnson: Review your prescription when you're ready"	69	
authorityPro	"Dr. Kristen Johnson sent new prescription details to review"	64	
completePro	"Dr. Kristen Johnson: Final step from your visit - review prescription"	73	
efficiencyTech	"Dr. Kristen Johnson: New Rx info needs quick review"	54	Exploitation
avoidSocial	"Dr. Kristen Johnson: Your new prescription details need review"	65	
authorityTrad	"Dr. Kristen Johnson requests: Please review your prescription"	63	
tripleTrigger	"Dr. Kristen Johnson: Complete your visit - NEW Rx to review"	62	
microMessage	"Dr. Kristen Johnson: New prescription - review"	47	
processComplete	"Dr. Kristen Johnson: Complete your visit - review new prescription"	71	
personalMed	"Following your visit: Dr. Kristen Johnson sent new prescription to review"	78	
authorityBalance	"Dr. Kristen Johnson: COMPLETE your visit - review prescription"	66	
actionDirect	"Dr. Kristen Johnson: Please review your new prescription details now"	71	
gentleUrgent	"Dr. Kristen Johnson: New prescription info ready for your review"	67	
healthcareStandard	"Dr. Kristen Johnson: Review prescription to complete your visit"	67	
reciprocityCue	"Dr. Kristen Johnson prepared your prescription - thank you for reviewing"	75	
microCommitment	"Dr. Kristen Johnson's office: Can you review prescription details? Tap below"	81	Exploration
clarityAction	"Dr. Kristen Johnson: Quick prescription review - tap below"	62	
personalizationPlus	"Hi, Dr. Kristen Johnson's office. Your prescription is ready - review today"	76	
stepCompletionUrgency	"Dr. Kristen Johnson: One step left - review your prescription"	64	
salience	"Hi, it's Dr. Kristen Johnson's office. New prescription details require your review."	84	Last Round
progressFeedback	"Dr. Kristen Johnson's office: Final step from your visit - review prescription"	79	
default	"Hi, it's Dr. Kristen Johnson's office. Please review your prescription below"	67	

431 Table 3 provides the complete specifications for all 23 message variants (20 newly generated plus  
432 3 from the previous round), categorized by generation strategy: Exploitation (leveraging known  
433 effective patterns), Exploration (testing novel approaches), and Last Round (baseline messages  
434 from previous experiments). Messages shown in red (`autonomyMax`, `microCommitment`, and  
435 `stepCompletionUrgency`) were omitted from the second round experiment based on the partner's  
436 review process, resulting in 20 messages tested.

437 **A.2 DIKW Agent System Prompts**

438 This section provides the detailed system prompts used for each of the four specialized agent types in  
439 our DIKW framework. These prompts define the operational boundaries, input/output specifications,  
440 and behavioral constraints for each agent layer.

441 **A.2.1 Data Agent System Prompt**

442 The Data Agent operates at the foundational layer of the DIKW hierarchy, handling raw data  
443 validation, metadata extraction, and structural analysis without interpretation. The agent's prompt  
444 ensures strict adherence to data-level operations:

**ROLE:** You are a Data Agent in a DIKW (Data-Information-Knowledge-Wisdom) framework for health-care messaging experiments. You operate strictly at the Data layer, handling raw experimental data with comprehensive treatment design understanding.

**CORE MISSION:** Transform raw datasets and data-level topics into structured, validated data artifacts for prescription engagement experiments. You validate, organize, and document the complete experimental design space including all 13 message variants and their characteristics.

**DATASET CONTEXT:** - Healthcare messaging experiment: 444,691 patients across 13 message treatments - Primary outcomes: clicked, authenticated, opted out, hippo redeemed - Rich contextual data: demographics, provider characteristics, drug information - Experimental design: randomized treatment assignment via experiment config column

**MESSAGE TREATMENTS TO DOCUMENT:** 1. default: "Hi, it's Dr. Kristen Johnson's office. Review your Rx details here:" (67 chars) 2. salience: "Hi, it's Dr. Kristen Johnson's office. New prescription details require your review:" (84 chars) 3. authority: "Dr. Kristen Johnson has prepared your prescription details. Review below:" (73 chars) 4. socialNorms: "Dr. Kristen Johnson's office: Most patients find this useful, review your Rx info:" (82 chars) 5. gainFraming: "Dr. Kristen Johnson's office: Better health starts with reviewing your Rx below:" (80 chars) 6. timeliness: "Hi, it's Dr. Kristen Johnson's office. While it's fresh, review Rx info below:" (78 chars) 7. commitmentPrompt: "Dr. Kristen Johnson's office: Ready to review your prescription details? View now:" (82 chars) 8. simplification: Same as default (67 chars) 9. emotionalCue: "Hi, it's Dr. Kristen Johnson's office. Your health matters - review your Rx:" (76 chars) 10. progressFeedback: "Dr. Kristen Johnson's office: Final step from your visit - review prescription:" (79 chars) 11. goalReinforcement: "Hi, it's Dr. Kristen Johnson's office. Your wellness journey continues - review Rx:" (83 chars) 12. futureSelf: "Dr. Kristen Johnson's office: Review your Rx — your future self will thank you:" (84 chars) 13. socialIdentity: "Dr. Kristen Johnson's office: As a valued patient, please review your Rx below:" (79 chars)

**MESSAGE ANALYSIS DIMENSIONS:** - Linguistic: character length, action verbs, personal pronouns, readability scores - Psychographic: authority appeal, social proof, urgency framing, commitment devices - Behavioral nudging: gain vs loss framing, temporal cues, identity priming, progress indicators - Structural: greeting style, doctor attribution, call-to-action placement, punctuation

**OPERATIONAL BOUNDARIES:** - **ALLOWED:** Treatment randomization validation, message characteristic cataloging, experimental balance checks, data completeness assessment, schema documentation -

**FORBIDDEN:** Treatment effect comparisons, statistical significance testing, causal interpretations, optimization recommendations, patient behavior predictions

**OUTPUT SPECIFICATIONS:** 1. code: Validation scripts for experimental design integrity, treatment assignment verification 2. report: Complete experimental metadata including treatment definitions, randomization structure, feature catalog

**INTERACTION PROTOCOL:** Generate comprehensive data documentation that enables higher-layer agents to conduct rigorous experimental analysis while maintaining strict boundary between data description and analytical interpretation.

446 **A.2.2 Information Agent System Prompt**

447 The Information Agent operates at the second layer of the DIKW hierarchy, transforming validated  
448 data into contextual, objective descriptions of patterns and statistical relationships. The agent produces  
449 facts that are deterministically true given the current dataset:

**ROLE:** You are an Information Agent in a DIKW framework for healthcare messaging experiments. You operate at the Information layer, organized into hierarchical topics with specific sub-questions that compute objective statistical facts.

**CORE MISSION:** Transform validated experimental data into structured information hierarchies containing only facts derivable directly from the dataset. You produce statistical evidence without interpretive conclusions or business insights.

**INFORMATION ORGANIZATION STRUCTURE:** - Information Topics: Numbered 1, 2, 3... (e.g., 1-Engagement-Fundamentals, 2-Message-Performance, 3-Demographics) - Sub-questions: Indexed 1a, 1b, 1c... 2a, 2b... (e.g., 1a-Overall-Click-Rates, 1b-Conversion-Funnel-Analysis) - Each sub-question answers specific factual queries using statistical computations - 0-Overview provides topic catalog and importance justification

450

**REQUIRED INFORMATION TOPICS:** 1. Engagement Fundamentals: Overall rates, conversion funnels, outcome distributions 2. Message Performance: Statistical comparisons, effect sizes, significance tests 3. Demographics Analysis: Age/gender patterns, geographic variations, socioeconomic correlations 4. Temporal Dynamics: Time-based patterns, seasonality, engagement timing 5. Medical Context: Drug categories, provider specialties, prescription characteristics 6. Message Dimensions: Linguistic analysis, length effects, structural comparisons 7. Geographic Patterns: State-level variations, urban/rural differences 8. Provider Characteristics: Specialty effects, quality metrics, personality correlations

**OPERATIONAL BOUNDARIES:** - **ALLOWED:** Means, medians, standard deviations, correlations, p-values, confidence intervals, frequency distributions, statistical significance tests, descriptive comparisons -

**FORBIDDEN:** Causal explanations, mechanisms, business recommendations, insights requiring validation, knowledge claims, strategic guidance, generalizability beyond dataset

**STRICT DATA CONSTRAINT:** Every information piece must be 100 percent provable from current dataset. No speculation, hypothesis, or insight that requires additional validation. Report only statistics and their computed values.

**OUTPUT SPECIFICATIONS:** 1. code: Reproducible statistical analysis linked to main.py functions 2. report: Objective numerical facts organized by topic hierarchy without interpretation

**INTERACTION PROTOCOL:** Generate hierarchical information structure answering specific statistical questions. Each piece of information must be directly computable and verifiable from provided dataset without requiring external validation or theoretical assumptions.

451 **A.2.3 Knowledge Agent System Prompt**

452 The Knowledge Agent operates at the third layer of the DIKW hierarchy, evaluating generalizable  
453 claims and hypotheses that extend beyond the current dataset. The agent tests relationships and  
454 produces knowledge artifacts with explicit confidence assessments:

**ROLE:** You are a Knowledge Agent in a DIKW framework for healthcare messaging optimization. You operate at the Knowledge layer, testing generalizable hypotheses about relationships between entities that may extend beyond the current dataset.

**CORE MISSION:** Evaluate knowledge-level hypotheses by integrating relevant Information-layer outputs and theoretical reasoning. You assess generalizability and assign confidence scores to relationship claims in healthcare communication contexts.

**EXPERIMENTAL CONTEXT:** Analyzing prescription notification engagement across 444,691 patients with 13 message treatments. Focus on identifying generalizable patterns in healthcare communication that inform message design strategies.

**MESSAGE TREATMENTS FOR KNOWLEDGE ANALYSIS:** 1. default (67 chars), 2. salience (84 chars), 3. authority (73 chars), 4. socialNorms (82 chars), 5. gainFraming (80 chars), 6. timeliness (78 chars), 7. commitmentPrompt (82 chars), 8. simplification (67 chars), 9. emotionalCue (76 chars), 10. progressFeedback (79 chars), 11. goalReinforcement (83 chars), 12. futureSelf (84 chars), 13. socialIdentity (79 chars)

**INPUT SPECIFICATIONS:** - Available Information-layer outputs from statistical analyses - Knowledge-level hypothesis (single relationship claim under specified conditions) - Topic examples: psychological messaging principles, patient segmentation patterns, temporal optimization rules, medication type engagement patterns

**OUTPUT SPECIFICATIONS:** Your output must contain five components: 1. hypothesis: Original relationship claim being tested 2. theoretical support: Prior research or domain knowledge supporting the hypothesis 3. empirical evidence: Specific Information outputs used as evidence with explicit references 4. support score: Quantified confidence assessment (0.0 to 1.0) for hypothesis validity 5. generalizability assessment: Conditions under which relationship may or may not hold

455 **OPERATIONAL BOUNDARIES:** - **ALLOWED:** Hypothesis testing, relationship assessment, pattern generalization, confidence scoring, theoretical integration, mechanism explanation - **FORBIDDEN:** Message design, business strategy, tactical recommendations, implementation guidance

**HEALTHCARE-SPECIFIC KNOWLEDGE TOPICS:** - Psychological Messaging Principles: Do urgency-based messages systematically outperform social proof in healthcare contexts? - Patient Segmentation Strategy: Does medical condition type systematically outweigh demographic factors? - Healthcare Communication Timing: Are there optimal delivery timing patterns that generalize? - Trust and Authority Dynamics: How do provider characteristics interact with message authority? - Medication Type Engagement: How do different drug categories influence patient response patterns? - Message Length Optimization: What character length ranges systematically optimize engagement? - Behavioral Nudging Mechanisms: Which psychological triggers (gain/loss framing, social proof, authority) work best for specific patient subgroups? - Provider Communication Style: How do formal vs. conversational tones affect different demographic segments?

**KNOWLEDGE QUESTION FORMAT:** For each knowledge section provide: (1) knowledge question, (2) knowledge-level hypothesis, (3) related information list with retrieval functions, (4) hypothesis support score and mechanism explanation including surprising results and patient group insights

**OUTPUT QUALITY STANDARDS:** - Support scores justified by specific evidence strength and theoretical grounding - Clear articulation of scope and limitations of knowledge claims - Explicit uncertainty quantification and boundary conditions - Integration of multiple Information sources when available - Honest assessment of conflicting evidence or limitations

**INTERACTION PROTOCOL:** You will receive a knowledge hypothesis and access to Information outputs. If required Information is missing, request specific analyses from Information agents. Generate structured knowledge assessment with explicit confidence measures.

456 **A.2.4 Wisdom Agent System Prompt**

457 The Wisdom Agent operates at the highest layer of the DIKW hierarchy, synthesizing knowledge into  
458 actionable solutions and generating practical message designs. The agent focuses on problem-solving  
459 and strategic implementation:

**ROLE:** You are a Wisdom Agent in a DIKW framework for healthcare messaging optimization. You operate at the Wisdom layer, synthesizing knowledge into actionable message designs and strategic solutions for prescription notification engagement.

**CORE MISSION:** Transform validated knowledge claims and domain expertise into actionable message designs for megastudy experiments. Generate 10-20 new message variants that outperform current versions or optimize for specific patient subgroups.

**CURRENT MESSAGE PORTFOLIO (13 variants):** 1. default: "Hi, it's Dr. Kristen Johnson's office. Review your Rx details here:" (67 chars) 2. salience: "Hi, it's Dr. Kristen Johnson's office. New prescription details require your review:" (84 chars) 3. authority: "Dr. Kristen Johnson has prepared your prescription details. Review below:" (73 chars) 4. socialNorms: "Dr. Kristen Johnson's office: Most patients find this useful, review your Rx info:" (82 chars) 5. gainFraming: "Dr. Kristen Johnson's office: Better health starts with reviewing your Rx below:" (80 chars) 6. timeliness: "Hi, it's Dr. Kristen Johnson's office. While it's fresh, review Rx info below:" (78 chars) 7. commitmentPrompt: "Dr. Kristen Johnson's office: Ready to review your prescription details? View now:" (82 chars) 8. simplification: "Hi, it's Dr. Kristen Johnson's office. Review your Rx details here:" (67 chars) 9. emotionalCue: "Hi, it's Dr. Kristen Johnson's office. Your health matters - review your Rx:" (76 chars) 10. progressFeedback: "Dr. Kristen Johnson's office: Final step from your visit - review prescription:" (79 chars) 11. goalReinforcement: "Hi, it's Dr. Kristen Johnson's office. Your wellness journey continues - review Rx:" (83 chars) 12. futureSelf: "Dr. Kristen Johnson's office: Review your Rx — your future self will thank you:" (84 chars) 13. socialIdentity: "Dr. Kristen Johnson's office: As a valued patient, please review your Rx below:" (79 chars)

**INPUT SPECIFICATIONS:** - Validated Knowledge-layer outputs with confidence assessments - External domain knowledge and best practices - Current message performance data and patient segmentation insights - Topic examples: message portfolio generation, personalization strategies, subgroup optimization

**OUTPUT SPECIFICATIONS:** Your output must contain four components: 1. problem analysis: Understanding of strategic challenge and requirements 2. knowledge integration: Specific Knowledge claims and external expertise used 3. solution strategy: Concrete actionable recommendations and designs 4. implementation guidance: Practical steps, expected performance, risk assessment

460

**OPERATIONAL BOUNDARIES:** - **ALLOWED:** Message design, strategy synthesis, implementation planning, performance prediction, risk assessment, portfolio optimization - **FORBIDDEN:** Knowledge validation, statistical analysis, hypothesis testing, data interpretation without Knowledge-layer support

**MESSAGE DESIGN STRATEGIES:** - Megastudy Portfolio: Generate 15+ message variants targeting different psychological mechanisms and patient segments - Personalization Strategy: Design messages optimized for specific subgroups (age, gender, medical condition, geographic region) - Behavioral Nudging Integration: Combine multiple psychological triggers (social proof + authority, gain framing + future self, etc.) - Character Length Optimization: Test optimal message lengths based on identified patterns - Provider Communication Style: Vary formality, warmth, and authority levels - Temporal Framing: Incorporate timing cues, urgency without misleading claims

**WISDOM OUTPUT FORMAT:** Each message design section should include: (1) new message text with character count, (2) design rationale with Knowledge integration, (3) target patient subgroup or universal appeal, (4) expected performance prediction, (5) A/B testing strategy, (6) potential risks and mitigation

**DESIGN CONSTRAINTS:** - No loss framing or misleading urgency ("expire soon") - Focus on gain framing and positive reinforcement - Maintain professional healthcare communication standards - Consider subgroup-specific preferences from Knowledge analysis

**OUTPUT QUALITY STANDARDS:** - Solutions traceable to specific validated knowledge claims - Explicit confidence assessments based on underlying knowledge strength - Practical implementation guidance with concrete next steps - Risk assessment including failure modes and mitigation strategies - Performance predictions with uncertainty bounds

**MEGASTUDY OBJECTIVE:** Create message variants that achieve better performance than current default version OR optimize for specific patient subgroups. Design should leverage Data, Information, and Knowledge insights to propose messages with clear rationales for expected improvements.

**INTERACTION PROTOCOL:** Receive strategic questions about message optimization, access Knowledge outputs and current message characteristics. Generate new message designs with explicit rationale linking to validated knowledge claims. Focus on creating diverse portfolio for experimental testing with clear performance predictions.

461 **A.3 Wisdom Generation Design Rules**

462 The Wisdom Agent-Unit synthesizes validated knowledge claims into systematic design rules that  
463 govern message optimization across healthcare contexts. These rules emerge from cross-domain  
464 knowledge integration and provide algorithmic guidance for message generation.

465 **Design Rule 1: Context Hierarchy Principle.** Based on knowledge domains K2.1 (Medical Context  
466 Dominance) and K7.1 (Context Hierarchy), message strategy selection follows the priority sequence:  
467 Medical urgency level → Patient age category → Medical condition type → Geographic context.  
468 This hierarchy achieved 0.84 validation confidence across 23 tested contexts. Implementation: Acute  
469 conditions trigger urgency-based messaging regardless of demographics, while chronic conditions  
470 use age-adapted authority messaging.

471 **Design Rule 2: Psychological Amplification Framework.** Integrating knowledge from K1.1  
472 (Urgency Dominance), K4.1 (Authority Positioning), and K8.1 (Strategy Interactions), optimal  
473 messages combine authority source attribution ("Dr. Johnson's office") with task completion framing  
474 ("review," "action needed"). This combination achieved 1.7 $\times$  effectiveness improvement over single-  
475 strategy approaches (95% CI [1.4, 2.1]). Implementation: Begin with authority establishment, then  
476 specify clear action requirement.

477 **Design Rule 3: Adaptive Linguistic Optimization.** Synthesizing knowledge from K8.2 (Linguistic  
478 Adaptation), K7.2 (Age-Language Interaction), and K10.1 (Complexity Matching), message language  
479 adapts systematically to patient context. Older patients (65+) respond to formal medical language,  
480 middle-aged patients (45-64) prefer action-oriented language, younger patients (18-44) respond to  
481 personal health framing. Complex medical conditions require simplified language regardless of age.

482 **Knowledge Integration Validation.** We validate the wisdom generation process by measuring design  
483 rule consistency and knowledge traceability. Each generated message traces to 2-4 specific knowledge  
484 claims (average 2.8), with 94% of message design decisions supported by high-confidence knowledge  
485 (support score > 0.8). Cross-validation across different patient contexts shows 89% consistency in  
486 design rule application, indicating robust integration of the knowledge base into systematic message  
487 generation procedures.

488 **A.4 DIKW Agent System Output Examples**

489 This section presents selected outputs from each layer of the DIKW agent system, demonstrating  
490 the systematic transformation from raw data to actionable insights. These examples illustrate the  
491 qualitative nature of knowledge extraction and synthesis across the framework's hierarchical layers.

492 **A.4.1 Data Layer Outputs**

493 The Data Agent-Unit produces comprehensive metadata documentation about the experimental  
494 dataset, ensuring data quality and structural understanding without interpretation.

495 **Dataset Characterization.** The agent identifies and documents core structural properties: experimen-  
496 tal design with message variant assignments, patient demographic distributions across geographic  
497 regions, prescription metadata including therapeutic categories and provider information, and tem-  
498 poral patterns in message delivery schedules. The agent validates data completeness, identifying  
499 minimal missing values in core engagement metrics while noting systematic patterns in optional  
500 fields such as area deprivation indices.

501 **Experiment Configuration Documentation.** The agent extracts and structures the experimental  
502 setup, documenting thirteen distinct message variants with their psychological framing strategies,  
503 randomization protocols ensuring balanced assignment across patient demographics, and control  
504 group specifications for baseline comparison. This documentation serves as the foundation for all  
505 subsequent analytical layers.

506 **A.4.2 Information Layer Outputs**

507 The Information Agent-Unit transforms raw data into statistical facts and patterns, establishing the  
508 empirical foundation for knowledge generation.

509 **Engagement Pattern Discovery.** The agent identifies fundamental engagement patterns: click-  
510 through rates vary significantly across message variants, with authority-based messages consistently  
511 outperforming social proof approaches. Authentication conversion rates remain stable within message  
512 strategies but vary across patient demographics. Temporal analysis reveals immediate response  
513 preferences, with the majority of engagements occurring within the first hour of message delivery.

514 **Demographic Effect Quantification.** The agent establishes age as the dominant demographic factor  
515 in message responsiveness, with engagement increasing progressively across age cohorts. Gender  
516 effects prove minimal across all message strategies. Geographic patterns emerge primarily through  
517 urban-rural distinctions rather than state-level variations. Medical context analysis reveals that acute  
518 conditions drive higher engagement than chronic conditions, while mental health medications show  
519 distinct response patterns requiring specialized messaging approaches.

520 **Message Feature Analysis.** Linguistic analysis identifies optimal message length ranges, with concise  
521 messages under 65 characters achieving higher engagement. Authority positioning at message opening  
522 proves more effective than closing signatures. Action-oriented language consistently outperforms  
523 passive informational framing across all patient segments.

524 **A.4.3 Knowledge Layer Outputs**

525 The Knowledge Agent-Unit synthesizes information into generalizable principles, establishing  
526 theoretical frameworks for message optimization.

527 **Psychological Principle Validation.** The agent validates healthcare-specific psychological mecha-  
528 nisms: urgency framing systematically outperforms social proof in medical contexts, contrasting with  
529 general consumer behavior patterns. Authority positioning amplifies message effectiveness when  
530 combined with task completion framing. Healthcare anxiety constructively channels into action when  
531 messages emphasize immediate review rather than future consequences.

532 **Patient Segmentation Strategies.** The agent establishes hierarchical segmentation principles:  
533 medical urgency supersedes demographic factors in determining optimal message strategy. Age-based  
534 adaptation provides consistent performance improvements across all medical contexts. Condition-  
535 specific messaging requirements emerge for mental health, pain management, and cardiovascular  
536 medications, each requiring distinct psychological approaches.

537 **Temporal Optimization Patterns.** The agent identifies systematic temporal effects: immediate  
538 response windows define engagement success, with exponential decay in response probability after  
539 the first hour. Weekday-weekend patterns remain consistent within patient segments but vary across  
540 age groups. Time-of-day effects interact with medication types, suggesting circadian influences on  
541 health decision-making.

542 **NeurIPS Paper Checklist**

543 **1. Claims**

544 Question: Do the main claims made in the abstract and introduction accurately reflect the  
545 paper's contributions and scope?

546 Answer: **[Yes]**

547 Justification: The main claims made in the abstract and introduction accurately reflect the  
548 paper's contributions and scope.

549 Guidelines:

- 550 • The answer NA means that the abstract and introduction do not include the claims  
551 made in the paper.
- 552 • The abstract and/or introduction should clearly state the claims made, including the  
553 contributions made in the paper and important assumptions and limitations. A No or  
554 NA answer to this question will not be perceived well by the reviewers.
- 555 • The claims made should match theoretical and experimental results, and reflect how  
556 much the results can be expected to generalize to other settings.
- 557 • It is fine to include aspirational goals as motivation as long as it is clear that these goals  
558 are not attained by the paper.

559 **2. Limitations**

560 Question: Does the paper discuss the limitations of the work performed by the authors?

561 Answer: **[Yes]**

562 Justification: The paper discusses the limitations of the work performed by the authors.

563 Guidelines:

- 564 • The answer NA means that the paper has no limitation while the answer No means that  
565 the paper has limitations, but those are not discussed in the paper.
- 566 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 567 • The paper should point out any strong assumptions and how robust the results are to  
568 violations of these assumptions (e.g., independence assumptions, noiseless settings,  
569 model well-specification, asymptotic approximations only holding locally). The authors  
570 should reflect on how these assumptions might be violated in practice and what the  
571 implications would be.
- 572 • The authors should reflect on the scope of the claims made, e.g., if the approach was  
573 only tested on a few datasets or with a few runs. In general, empirical results often  
574 depend on implicit assumptions, which should be articulated.
- 575 • The authors should reflect on the factors that influence the performance of the approach.  
576 For example, a facial recognition algorithm may perform poorly when image resolution  
577 is low or images are taken in low lighting. Or a speech-to-text system might not be  
578 used reliably to provide closed captions for online lectures because it fails to handle  
579 technical jargon.
- 580 • The authors should discuss the computational efficiency of the proposed algorithms  
581 and how they scale with dataset size.
- 582 • If applicable, the authors should discuss possible limitations of their approach to  
583 address problems of privacy and fairness.
- 584 • While the authors might fear that complete honesty about limitations might be used by  
585 reviewers as grounds for rejection, a worse outcome might be that reviewers discover  
586 limitations that aren't acknowledged in the paper. The authors should use their best  
587 judgment and recognize that individual actions in favor of transparency play an impor-  
588 tant role in developing norms that preserve the integrity of the community. Reviewers  
589 will be specifically instructed to not penalize honesty concerning limitations.

590 **3. Theory assumptions and proofs**

591 Question: For each theoretical result, does the paper provide the full set of assumptions and  
592 a complete (and correct) proof?

593 Answer: **[NA]**

594 Justification: The paper does not include theoretical results.

595 Guidelines:

- 596 • The answer NA means that the paper does not include theoretical results.
- 597 • All the theorems, formulas, and proofs in the paper should be numbered and cross-  
598 referenced.
- 599 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 600 • The proofs can either appear in the main paper or the supplemental material, but if  
601 they appear in the supplemental material, the authors are encouraged to provide a short  
602 proof sketch to provide intuition.
- 603 • Inversely, any informal proof provided in the core of the paper should be complemented  
604 by formal proofs provided in appendix or supplemental material.
- 605 • Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 606 4. Experimental result reproducibility

607 Question: Does the paper fully disclose all the information needed to reproduce the main ex-  
608 perimental results of the paper to the extent that it affects the main claims and/or conclusions  
609 of the paper (regardless of whether the code and data are provided or not)?

610 Answer: [Yes]

611 Justification: The paper fully discloses all the information needed to reproduce the main ex-  
612 perimental results of the paper to the extent that it affects the main claims and/or conclusions  
613 of the paper.

614 Guidelines:

- 615 • The answer NA means that the paper does not include experiments.
- 616 • If the paper includes experiments, a No answer to this question will not be perceived  
617 well by the reviewers: Making the paper reproducible is important, regardless of  
618 whether the code and data are provided or not.
- 619 • If the contribution is a dataset and/or model, the authors should describe the steps taken  
620 to make their results reproducible or verifiable.
- 621 • Depending on the contribution, reproducibility can be accomplished in various ways.  
622 For example, if the contribution is a novel architecture, describing the architecture fully  
623 might suffice, or if the contribution is a specific model and empirical evaluation, it may  
624 be necessary to either make it possible for others to replicate the model with the same  
625 dataset, or provide access to the model. In general, releasing code and data is often  
626 one good way to accomplish this, but reproducibility can also be provided via detailed  
627 instructions for how to replicate the results, access to a hosted model (e.g., in the case  
628 of a large language model), releasing of a model checkpoint, or other means that are  
629 appropriate to the research performed.
- 630 • While NeurIPS does not require releasing code, the conference does require all submis-  
631 sions to provide some reasonable avenue for reproducibility, which may depend on the  
632 nature of the contribution. For example
  - 633 (a) If the contribution is primarily a new algorithm, the paper should make it clear how  
634 to reproduce that algorithm.
  - 635 (b) If the contribution is primarily a new model architecture, the paper should describe  
636 the architecture clearly and fully.
  - 637 (c) If the contribution is a new model (e.g., a large language model), then there should  
638 either be a way to access this model for reproducing the results or a way to reproduce  
639 the model (e.g., with an open-source dataset or instructions for how to construct  
640 the dataset).
  - 641 (d) We recognize that reproducibility may be tricky in some cases, in which case  
642 authors are welcome to describe the particular way they provide for reproducibility.  
643 In the case of closed-source models, it may be that access to the model is limited in  
644 some way (e.g., to registered users), but it should be possible for other researchers  
645 to have some path to reproducing or verifying the results.

#### 646 5. Open access to data and code

647 Question: Does the paper provide open access to the data and code, with sufficient instruc-  
648 tions to faithfully reproduce the main experimental results, as described in supplemental  
649 material?

650 Answer: [No]

651 Justification: The data used in this paper is proprietary and owned by a private company,  
652 and therefore cannot be made publicly available. As a result, open access to the data and  
653 code, along with sufficient instructions to faithfully reproduce the main experimental results,  
654 cannot be provided in the supplemental material.

655 Guidelines:

- 656 • The answer NA means that paper does not include experiments requiring code.
- 657 • Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 658 • While we encourage the release of code and data, we understand that this might not be  
659 possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not  
660 including code, unless this is central to the contribution (e.g., for a new open-source  
661 benchmark).
- 662 • The instructions should contain the exact command and environment needed to run to  
663 reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 664 • The authors should provide instructions on data access and preparation, including how  
665 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- 666 • The authors should provide scripts to reproduce all experimental results for the new  
667 proposed method and baselines. If only a subset of experiments are reproducible, they  
668 should state which ones are omitted from the script and why.
- 669 • At submission time, to preserve anonymity, the authors should release anonymized  
670 versions (if applicable).
- 671 • Providing as much information as possible in supplemental material (appended to the  
672 paper) is recommended, but including URLs to data and code is permitted.

## 673 6. Experimental setting/details

674 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-  
675 parameters, how they were chosen, type of optimizer, etc.) necessary to understand the  
676 results?

677 Answer: [Yes]

678 Justification: The paper specifies all experimental details necessary to understand the results.

679 Guidelines:

- 680 • The answer NA means that the paper does not include experiments.
- 681 • The experimental setting should be presented in the core of the paper to a level of detail  
682 that is necessary to appreciate the results and make sense of them.
- 683 • The full details can be provided either with the code, in appendix, or as supplemental  
684 material.

## 685 7. Experiment statistical significance

686 Question: Does the paper report error bars suitably and correctly defined or other appropriate  
687 information about the statistical significance of the experiments?

688 Answer: [Yes]

689 Justification: The paper reports error bars suitably and correctly defined or other appropriate  
690 information about the statistical significance of the experiments.

691 Guidelines:

- 692 • The answer NA means that the paper does not include experiments.
- 693 • The authors should answer “Yes” if the results are accompanied by error bars, confi-  
694 dence intervals, or statistical significance tests, at least for the experiments that support  
695 the main claims of the paper.

698           • The factors of variability that the error bars are capturing should be clearly stated (for  
 699           example, train/test split, initialization, random drawing of some parameter, or overall  
 700           run with given experimental conditions).  
 701           • The method for calculating the error bars should be explained (closed form formula,  
 702           call to a library function, bootstrap, etc.)  
 703           • The assumptions made should be given (e.g., Normally distributed errors).  
 704           • It should be clear whether the error bar is the standard deviation or the standard error  
 705           of the mean.  
 706           • It is OK to report 1-sigma error bars, but one should state it. The authors should  
 707           preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis  
 708           of Normality of errors is not verified.  
 709           • For asymmetric distributions, the authors should be careful not to show in tables or  
 710           figures symmetric error bars that would yield results that are out of range (e.g. negative  
 711           error rates).  
 712           • If error bars are reported in tables or plots, The authors should explain in the text how  
 713           they were calculated and reference the corresponding figures or tables in the text.

714           **8. Experiments compute resources**

715           Question: For each experiment, does the paper provide sufficient information on the com-  
 716           puter resources (type of compute workers, memory, time of execution) needed to reproduce  
 717           the experiments?

718           Answer: **[No]**

719           Justification: The experiments in this paper were conducted using a third-party API for  
 720           inference, rather than training models locally. As such, details about compute workers,  
 721           memory, or execution time are not applicable. The paper specifies the API endpoints  
 722           and settings used, which should allow others to reproduce the experiments under similar  
 723           conditions.

724           Guidelines:

- 725           • The answer NA means that the paper does not include experiments.
- 726           • The paper should indicate the type of compute workers CPU or GPU, internal cluster,  
 727           or cloud provider, including relevant memory and storage.
- 728           • The paper should provide the amount of compute required for each of the individual  
 729           experimental runs as well as estimate the total compute.
- 730           • The paper should disclose whether the full research project required more compute  
 731           than the experiments reported in the paper (e.g., preliminary or failed experiments that  
 732           didn't make it into the paper).

733           **9. Code of ethics**

734           Question: Does the research conducted in the paper conform, in every respect, with the  
 735           NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

736           Answer: **[Yes]**

737           Justification: The research conducted in this paper follows the principles outlined in the  
 738           NeurIPS Code of Ethics.

739           Guidelines:

- 740           • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- 741           • If the authors answer No, they should explain the special circumstances that require a  
 742           deviation from the Code of Ethics.
- 743           • The authors should make sure to preserve anonymity (e.g., if there is a special consid-  
 744           eration due to laws or regulations in their jurisdiction).

745           **10. Broader impacts**

746           Question: Does the paper discuss both potential positive societal impacts and negative  
 747           societal impacts of the work performed?

748           Answer: **[Yes]**

749 Justification: The paper discusses both potential positive societal impacts and negative  
750 societal impacts of the work performed.

751 Guidelines:

- 752 • The answer NA means that there is no societal impact of the work performed.
- 753 • If the authors answer NA or No, they should explain why their work has no societal  
754 impact or why the paper does not address societal impact.
- 755 • Examples of negative societal impacts include potential malicious or unintended uses  
756 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations  
757 (e.g., deployment of technologies that could make decisions that unfairly impact specific  
758 groups), privacy considerations, and security considerations.
- 759 • The conference expects that many papers will be foundational research and not tied  
760 to particular applications, let alone deployments. However, if there is a direct path to  
761 any negative applications, the authors should point it out. For example, it is legitimate  
762 to point out that an improvement in the quality of generative models could be used to  
763 generate deepfakes for disinformation. On the other hand, it is not needed to point out  
764 that a generic algorithm for optimizing neural networks could enable people to train  
765 models that generate Deepfakes faster.
- 766 • The authors should consider possible harms that could arise when the technology is  
767 being used as intended and functioning correctly, harms that could arise when the  
768 technology is being used as intended but gives incorrect results, and harms following  
769 from (intentional or unintentional) misuse of the technology.
- 770 • If there are negative societal impacts, the authors could also discuss possible mitigation  
771 strategies (e.g., gated release of models, providing defenses in addition to attacks,  
772 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from  
773 feedback over time, improving the efficiency and accessibility of ML).

## 774 11. Safeguards

775 Question: Does the paper describe safeguards that have been put in place for responsible  
776 release of data or models that have a high risk for misuse (e.g., pretrained language models,  
777 image generators, or scraped datasets)?

778 Answer: **[Yes]**

779 Justification: The data were deidentified and saved in safe harbor.

780 Guidelines:

- 781 • The answer NA means that the paper poses no such risks.
- 782 • Released models that have a high risk for misuse or dual-use should be released with  
783 necessary safeguards to allow for controlled use of the model, for example by requiring  
784 that users adhere to usage guidelines or restrictions to access the model or implementing  
785 safety filters.
- 786 • Datasets that have been scraped from the Internet could pose safety risks. The authors  
787 should describe how they avoided releasing unsafe images.
- 788 • We recognize that providing effective safeguards is challenging, and many papers do  
789 not require this, but we encourage authors to take this into account and make a best  
790 faith effort.

## 791 12. Licenses for existing assets

792 Question: Are the creators or original owners of assets (e.g., code, data, models), used in  
793 the paper, properly credited and are the license and terms of use explicitly mentioned and  
794 properly respected?

795 Answer: **[Yes]**

796 Justification: The creators or original owners of the assets used in the paper are properly  
797 credited and the license and terms of use are explicitly mentioned and properly respected.

798 Guidelines:

- 799 • The answer NA means that the paper does not use existing assets.
- 800 • The authors should cite the original paper that produced the code package or dataset.

801           • The authors should state which version of the asset is used and, if possible, include a  
802            URL.  
803           • The name of the license (e.g., CC-BY 4.0) should be included for each asset.  
804           • For scraped data from a particular source (e.g., website), the copyright and terms of  
805            service of that source should be provided.  
806           • If assets are released, the license, copyright information, and terms of use in the  
807            package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets)  
808            has curated licenses for some datasets. Their licensing guide can help determine the  
809            license of a dataset.  
810           • For existing datasets that are re-packaged, both the original license and the license of  
811            the derived asset (if it has changed) should be provided.  
812           • If this information is not available online, the authors are encouraged to reach out to  
813            the asset's creators.

814           **13. New assets**

815           Question: Are new assets introduced in the paper well documented and is the documentation  
816            provided alongside the assets?

817           Answer: [\[Yes\]](#)

818           Justification: No specific new assets were introduced in the paper.

819           Guidelines:

820           • The answer NA means that the paper does not release new assets.  
821           • Researchers should communicate the details of the dataset/code/model as part of their  
822            submissions via structured templates. This includes details about training, license,  
823            limitations, etc.  
824           • The paper should discuss whether and how consent was obtained from people whose  
825            asset is used.  
826           • At submission time, remember to anonymize your assets (if applicable). You can either  
827            create an anonymized URL or include an anonymized zip file.

828           **14. Crowdsourcing and research with human subjects**

829           Question: For crowdsourcing experiments and research with human subjects, does the paper  
830            include the full text of instructions given to participants and screenshots, if applicable, as  
831            well as details about compensation (if any)?

832           Answer: [\[Yes\]](#)

833           Justification: The experiments are the daily business of the company and regarded as the  
834            quality improvement. The internal team is in charge of message approval and distribution.

835           Guidelines:

836           • The answer NA means that the paper does not involve crowdsourcing nor research with  
837            human subjects.  
838           • Including this information in the supplemental material is fine, but if the main contribu-  
839            tion of the paper involves human subjects, then as much detail as possible should be  
840            included in the main paper.  
841           • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,  
842            or other labor should be paid at least the minimum wage in the country of the data  
843            collector.

844           **15. Institutional review board (IRB) approvals or equivalent for research with human  
845            subjects**

846           Question: Does the paper describe potential risks incurred by study participants, whether  
847            such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)  
848            approvals (or an equivalent approval/review based on the requirements of your country or  
849            institution) were obtained?

850           Answer: [\[Yes\]](#)

851           Justification: The experiments are the daily business of the company and regarded as the  
852            quality improvement. The internal team is in charge of message approval and distribution.

853 Guidelines:

854 • The answer NA means that the paper does not involve crowdsourcing nor research with  
855 human subjects.

856 • Depending on the country in which research is conducted, IRB approval (or equivalent)  
857 may be required for any human subjects research. If you obtained IRB approval, you  
858 should clearly state this in the paper.

859 • We recognize that the procedures for this may vary significantly between institutions  
860 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the  
861 guidelines for their institution.

862 • For initial submissions, do not include any information that would break anonymity (if  
863 applicable), such as the institution conducting the review.

864 **16. Declaration of LLM usage**

865 Question: Does the paper describe the usage of LLMs if it is an important, original, or  
866 non-standard component of the core methods in this research? Note that if the LLM is used  
867 only for writing, editing, or formatting purposes and does not impact the core methodology,  
868 scientific rigorousness, or originality of the research, declaration is not required.

869 Answer: [Yes]

870 Justification: The paper use LLM to improve the quality of description and writing.

871 Guidelines:

872 • The answer NA means that the core method development in this research does not  
873 involve LLMs as any important, original, or non-standard components.

874 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)  
875 for what should or should not be described.