

000  
001  **AI REALTOR: TOWARDS GROUNDED PERSUASIVE**  
002 **LANGUAGE GENERATION FOR AUTOMATED COPY-**  
003 **WRITING**

004  
005 **Anonymous authors**  
006 Paper under double-blind review

009  
010 **ABSTRACT**

013 This paper develops an agentic framework that employs large language models  
014 (LLMs) for grounded persuasive language generation in automated copywriting,  
015 with real estate marketing as a focal application. Our method is designed to align  
016 the generated content with user preferences while highlighting useful factual at-  
017 tributes. This agent consists of three key modules: (1) Grounding Module, mim-  
018 icking expert human behavior to predict marketable features; (2) Personalization  
019 Module, aligning content with user preferences; (3) Marketing Module, ensuring  
020 factual accuracy and the inclusion of localized features. We conduct systematic  
021 human-subject experiments in the domain of real estate marketing, with a focus  
022 group of potential house buyers. The results demonstrate that marketing descrip-  
023 tions generated by our approach are preferred over those written by human ex-  
024 perts by a clear margin while maintaining the same level of factual accuracy. Our  
025 findings suggest a promising agentic approach to automate large-scale targeted  
026 copywriting while ensuring factuality of content generation.

027 **1 INTRODUCTION**

030 While large language models (LLMs) have made significant strides across various tasks, their abil-  
031 ity to persuade remains an underexplored frontier (see a discussion of related work in Section 6).  
032 This however is a particularly important capability since persuasion-related economic activities — a  
033 common thread in almost all voluntary transactions from advertising and lobbying to litigation and  
034 negotiation — underpin roughly 30% of the US GDP (Antioch, 2013), hence gives rise to tremen-  
035 dous opportunity for applying LLMs across a wide range of sectors. Meanwhile, this same potential  
036 introduces serious trustworthiness concerns. If LLMs can generate persuasive content at scale, their  
037 influence on human opinions raises risks of misinformation, manipulation and misuse, especially in  
038 sensitive domains such as political campaigns (Voelkel et al., 2023; Goldstein et al., 2024).

039 Therefore, we focus our study on the task of language generation for grounded persuasion, that is,  
040 the production of persuasive content that is faithful in factual details. This task is especially critical  
041 in copywriting, the practice of creating marketing text that seeks to influence consumer decisions,  
042 where its effectiveness can be directly assessed through measurable behavioral outcomes (e.g., rat-  
043 ings, engagement, and conversions), yet must remain strictly constrained by factual accuracy. In  
044 particular, we choose the domain of real estate marketing (see our rationale in § 2) and develop an  
045 agentic solution,  **AI Realtor**, under an economic scaffolding to investigate key elements of  
046 grounded persuasion. Below, we outline core contributions and the structure of this paper:

047 **① Real-World Evaluation:** Using real estate marketing as our testbed, we construct a large dataset  
048 from Zillow and design an experimental website that simulates the house search process, including  
049 buyer preference elicitation. We recruit a targeted group of potential home buyers to evaluate the  
050 persuasiveness of the generated marketing content (§ 2).

051 **② Theoretical Grounding:** We draw on the economic theory of information design in strategic  
052 communication games (Bergemann & Morris, 2019) to guide the agentic workflow. This includes  
053 processing the raw (factual) attributes of properties, selecting key features to highlight, and generat-  
054 ing persuasive, human-like marketing content (§ 3).

054 ③ **Agentic Pipeline:** We develop an LLM-based agent (§ 4) with three key modules: a *Grounding*  
 055 *Module*, which mimics human expertise in identifying and signaling critical, credible selling points;  
 056 a *Personalization Module*, which tailors content to user preferences; and a *Marketing Module*, which  
 057 ensures factual consistency and incorporates localized features.

058 ④ **Empirical Effectiveness:** Our system achieves a 70% win rate over human experts while main-  
 059 taining, if not exceeding, the same level of factual accuracy, establishing the first LLM benchmark  
 060 for grounded persuasion with measurable behavioral impact (§ 5).

## 062 2 A BENCHMARK FOR GROUNDED PERSUASION

063 **Motivations and Challenges** Establishing a robust evaluation benchmark for persuasion faces  
 064 two core challenges. First, persuasiveness is inherently subjective: unlike reasoning or planning  
 065 (which have objective metrics), its effectiveness depends on human feedback and varies with in-  
 066 dividual preferences and contexts. Second, persuasion is multifaceted, with domain-specific tech-  
 067 niques shaped by psychology, economics, and communication. Existing LLM research mostly focus  
 068 on political or opinion-based persuasion, where evaluations are complicated by cognitive biases and  
 069 adversarial framing. For example, Hackenburg & Margetts (2024) and Matz et al. (2024) reached  
 070 conflicting conclusions using similar experimental designs. Durmus et al. (2024) highlight the an-  
 071 choring effect – the tendency to cling to initial beliefs – making opinion shifts hard to measure. They  
 072 also find fabricated content is often more persuasive, raising ethical and methodological concerns.  
 073 These limitations underscore the need for new benchmarks in controlled, fact-grounded settings.

074 **Real Estate Marketing (REM) as Testbed** Identifying well-scoped testbeds is key to launch  
 075 systematic investigations of general AI capabilities, as demonstrated by recent benchmarks (Yao  
 076 et al., 2022; Xie et al., 2024). The real estate marketing domain is ideal for our study because:

077 ① *High-stakes, rational decisions:* Real estate involves high-stakes economic decisions, where  
 078 buyers typically hold rational, fact-based beliefs — unlike more emotionally charged or polarized  
 079 domains. Persuasive language in this setting must be both compelling and truthful.

080 ② *Measurable economic impact:* Effective persuasion has tangible economic value in real estate.  
 081 While structured attributes and images capture initial attention, industry guidance emphasizes that  
 082 descriptive text is critical for conveying the unique experience of living in a home (Zillow, n.d.).  
 083 The potential for LLMs to assist in this high-value task is further illustrated by recent anecdotal  
 084 accounts (User, 2023).

085 ③ *Rich, structured datasets:* The availability of extensive property listings with carefully labeled  
 086 attributes (e.g., from Zillow) enables domain-specific training and thorough empirical evaluations.

087 **Realistic Evaluation Interface and Persuasiveness Measurement** Our framework prioritizes  
 088 two criteria: (1) immersive user interaction to capture authentic feedback and (2) dynamic preference  
 089 elicitation for personalized generation. We replicate real-world homebuyer behavior by integrating  
 090 50k+ real-world listings into a web platform. See Appendix B and D for a full description of the web  
 091 interface and dataset. We evaluate persuasion via pairwise comparisons: buyers view a property with  
 092 two model-generated descriptions and select the more compelling one. Persuasiveness is quantified  
 093 via Elo scores (Elo, 1967); factual accuracy is verified against listing metadata (see § 5).

## 094 3 AN ECONOMIC SCAFFOLDING OF COPYWRITING

095 Copywriting fundamentally is about communicating product information, often selectively, to shape  
 096 potential buyers' perceptions and influence their purchasing decisions. This process of information  
 097 signaling, also known as persuasion, has been extensively studied in decision theory and informa-  
 098 tion economics (Spence, 1978; Arrow, 1996; Kamenica & Gentzkow, 2011; Connelly et al., 2011),  
 099 typically within stylized mathematical models. To enable practical automated copywriting in natural  
 100 language, we employ previous mathematical models/findings to build a framework compatible the  
 101 agentic scaffolding enabled by modern language generation technology.

102 **Attributes** Formally, we represent a generic *product X* (e.g., a house or an Amazon item) as an  $n$ -  
 103 dimensional vector  $X = (X_1, X_2, \dots, X_n)$ . Each  $X_i$  is called a raw attribute (or simply *attribute*).  
 104 Attributes capture the factual and measurable characteristics of the product (e.g., square footage,

108 distance to transit). A specific product instance is denoted by vector  $\mathbf{x} = (x_1, \dots, x_n)$  where  
 109  $x_i \in \mathcal{X}_i$  is the *realized* value of attribute  $X_i$ . Let  $\mathcal{X} = \Pi_i \mathcal{X}_i$  be the domain of  $\mathbf{x}$ .  
 110

111 **Features** Marketers often emphasize certain attractive properties of a product (e.g., “spacious lay-  
 112 out” and “prime location” in REM), derived from its underlying raw attributes. We refer to these  
 113 as signaling features (or simply *features*). Importantly, features differ from attributes: while some  
 114 attributes may directly serve as features, features generally capture the more abstract (and some-  
 115 times ambiguous) properties. We denote the feature set as  $S = (S_1, \dots, S_m)$ , with a feature vector  
 116  $\mathbf{s} = (s_1, \dots, s_m)$ , where each  $s_i \in [0, 1]$  quantifies the *intensity* or likelihood of feature  $S_i$  being.  
 117 For example,  $S_i$  could be “bright room” and correspondingly  $s_i$  denotes the extent to which rooms  
 118 of the house are bright. In practice, both  $x_i$  and  $s_i$  can be assessed by domain experts.

119 **Signaling via the Attribute-Feature Mapping** In our model, signaling features convey partial  
 120 information to influence potential buyers’ beliefs, leveraging the inherent cognitive mapping in nat-  
 121 ural language. For instance, a feature “bright room” may probabilistically imply high floor, southern  
 122 exposure, and modern lighting – all affecting buyers’ perceptions and decisions. (e.g., deciding to  
 123 schedule a visit). We formalize this with a mapping  $\pi : \mathcal{X} \rightarrow [0, 1]^m$  that transform raw attributes  
 124  $\mathbf{x} \in \mathcal{X}$  into feature intensities  $\mathbf{s} \in [0, 1]^m$ . That is,  $\mathbf{s} = \pi(\mathbf{x})$ . Sometime, we use  $\mathbf{s}(\mathbf{x})$  to emphasize  
 125 the dependence of  $\mathbf{s}$  on the underlying attributes  $\mathbf{x}$ , and  $s_j(\mathbf{x})$  is its  $j$ -th entry. This mapping reflects  
 126 the commonsense inference: given  $\mathbf{x}$ , how strongly we can claim the presence of feature  $S_j$ .  
 127

128 This attribute-feature mapping  $\pi$  is widely studied in both machine learning and economics. In  
 129 Bayesian statistics,  $X_i$  is an observable variable,  $S_j$  a latent variable, and  $\pi$  captures their proba-  
 130 bilistic dependence. In information economics,  $X_i$  represents a state,  $S_j$  a *signal*, and  $\pi$  is known  
 131 as a *signaling scheme*. Signals can be strategically designed to reveal partial information about the  
 132 state, and prior work has made significant progress in their optimal design to influence the equi-  
 133 librium outcomes (Kamenica & Gentzkow, 2011; Bergemann et al., 2015; Bergemann & Morris,  
 134 2019). Our work moves beyond this traditional Bayesian framing to incorporate the nuanced role of  
 135 natural language—often abstracted away in prior models—and to uncover the implicit, *commonsense*  
 136 mappings behind linguistic signals, rather than design new schemes.

137 **Marketing Design under Information Asymmetry** Marketing fundamentally exploits informa-  
 138 tion asymmetry between sellers and buyers (Grossman, 1981; Lewis, 2011; Dimoka et al., 2012;  
 139 Kurlat & Scheuer, 2021). This important insight, along with its broader implications in general eco-  
 140 nomic markets, was notably recognized by the 2002 Nobel Economics Prize (Akerlof, 1978; Spence,  
 141 1978; Stiglitz, 1975; Löfgren et al., 2002). In our setting, the seller or seller’s agent knows the exact  
 142 product attributes  $\mathbf{x}$  and the corresponding feature values  $\mathbf{s}(\mathbf{x})$ , while the buyer enters the market  
 143 with only a prior belief  $\mu$  over the distribution of attributes in  $\mathcal{X}$ . Without specific knowledge of the  
 144 product  $\mathbf{x}$ , the buyer holds an expected belief over features:

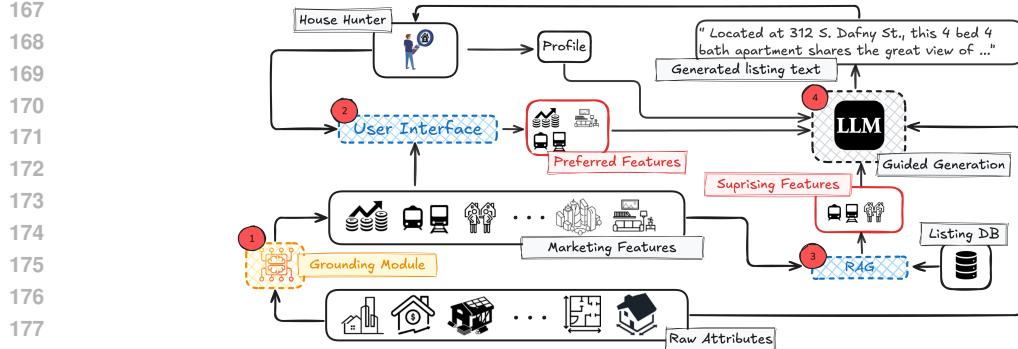
$$145 \text{Initial belief of features: } \bar{\mathbf{s}}(\mu) = \int_{\mathbf{x} \in \mathcal{X}} \mathbf{s}(\mathbf{x}) d\mu(\mathbf{x}). \quad (1)$$

146 Given the asymmetric feature beliefs between the buyer and seller, the purpose of marketing can  
 147 be described as revealing features, subject to communication constraints, to shift the buyer’s belief  
 148 from  $\bar{\mathbf{s}}(\mu)$  towards  $\mathbf{s}(\mathbf{x})$  with the goal of increasing the product’s attractiveness to the buyer.

149 **Grounded Persuasion in Natural Language** The remaining part of our model is to optimize the  
 150 persuasiveness of marketing content. The typical approach in economic theory is to develop models  
 151 capturing buyers’ belief updates and decision-making processes. However, these are difficult to  
 152 operationalize due to the absence of concrete buyer utility functions and behavioral models. Instead,  
 153 we leverage the generative capabilities of LLMs, guided by heuristics and instructions tailored for  
 154 grounded persuasion. At a high level, we use the attribute-feature mapping  $\pi$  to guide the selection  
 155 of a feature subset  $S^*$  to emphasize in generation. User preferences  $\mathbf{r}$  are elicited and incorporated  
 156 into a prompt  $\mathcal{I}^*$  for personalization. We hypothesize that the LLM approximates the solution to an  
 157 implicit optimization problem:  $L^* = \arg \max_{L \in \mathcal{L}} \Pr(L | \mathcal{I}^*, S^*, \mathbf{r}) \approx \arg \max_{L \in \mathcal{L}(\mathbf{x})} U^{\mathbf{r}}(L)$ . That  
 158 is, the language  $L^*$ , output by an LLM provided carefully designed prompts  $\mathcal{I}^*$ , selected features  $S^*$   
 159 and user preferences  $\mathbf{r}$ , could approximately maximize users’ preference-adjusted persuasiveness  
 160 function  $U^{\mathbf{r}}$ . Moreover, the generated language  $L$  will obey product facts (i.e., is *grounded*), or  
 161 concretely, be drawn from set  $\mathcal{L}(\mathbf{x})$  that includes all languages consistent with the product attribute  
 162  $\mathbf{x}$ . Our subsequent agent implementation and its practical effectiveness support this hypothesis; we  
 163 further conjecture that more powerful models will generally be able to find better-approximated

162 solutions to this optimization problem. Given this formulation, our design objective is to support  
 163 the LLM in solving the above optimization problem by constructing effective prompts  $\mathcal{I}^*$ , selecting  
 164 appropriate features  $\mathcal{S}^*$ , and representing user preferences  $\mathbf{r}$ . The following section describes our  
 165 implementation.

166



179 Figure 1: Illustration of the Design Pipeline of 🏠 AI Realtor.  
 180

## 4 THE AGENTIC IMPLEMENTATION OF 🏠 AI REALTOR

184 This section outlines the core design of 🏠 AI Realtor, an AI agent that process multiple levels of  
 185 marketing information to compose persuasive descriptions for real estate listings and actively learn  
 186 to adapt its language to individual buyer preferences. At a high level, our approach operationalizes  
 187 microeconomic models by implementing the following three key ingredients:

- 188 • Grounding Module: identify the attribute-feature mapping  $\pi$ ;
- 189 • Personalization Module: elicit and represent buyer preferences  $\mathbf{r}$ ;
- 190 • Marketing Module: select useful yet factual marketing features  $\mathcal{S}^*$  based on  $\pi, \mathbf{r}$ .

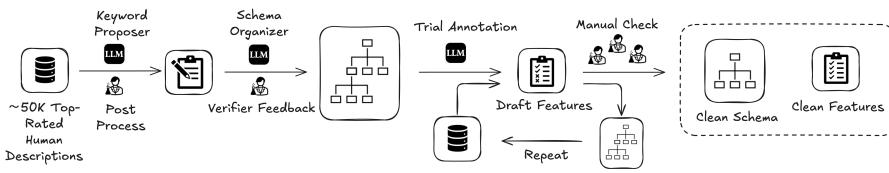
191 The overall system pipeline is illustrated in Figure 1. Below, we highlight the novel contributions  
 192 within each of the three modules. Full implementation details are provided in Appendix C.  
 193

### 4.1 GROUNDING MODULE: PREDICTING CREDIBLE FEATURES FOR MARKETING

196 Our model assumes the existence of attribute-feature mappings that marketers can use to influence  
 197 buyer beliefs and behaviors. However, a key challenge is that while raw attributes (e.g., square  
 198 footage) are available, high-level signaling features (e.g., “convenient transportation”) lack explicit  
 199 annotations in our dataset. This absence of supervision, combined with the open-ended nature of  
 200 natural language, where many tokens may serve as features with overlapping or ambiguous mean-  
 201 ings, makes the learning problem inherently difficult. Without a structured representation, the label  
 202 space becomes too sparse for effective training. Indeed, we find that directly prompting LLMs to  
 203 generate features produces redundant or incomplete feature sets, which undermines the quality of  
 204 the learned mapping.

205 Manual annotation by human experts could address this issue but is labor-intensive, costly to scale,  
 206 and difficult to personalize. We therefore adopt a machine learning approach to infer the attribute-  
 207 feature mapping automatically from unlabeled data, guided by LLM-assisted schema construction  
 208 and weak supervision. Specifically, we provide LLMs with a large pool of candidate features ex-  
 209 tracted from the dataset and prompt them to organize these into a hierarchical schema. A small  
 210 number of human annotators validate the output to monitor hallucinations and refine definitions.  
 211 This process, illustrated in Figure 2, yields a compact and expressive feature representation. Once  
 212 created, this feature set and mapping can be reused across models within the same marketing domain  
 213 and is thus a *one-time cost*.

214 Using the finalized feature schema, we guide an LLM to annotate whether each feature  $s_i$  is present  
 215 in a given listing, based on its attributes  $\mathbf{x}$  and corresponding human-written description. After  
 216 standard preprocessing (e.g., removing low-quality texts, normalizing attributes), we curate a labeled

216  
217  
218  
219  
220  
221222 Figure 2: Illustration of the inductive feature schema construction pipeline.  
223224 dataset and train a neural network to learn the attribute-feature mapping.<sup>1</sup> On a random 4:1 train-  
225 test split, our model achieves 69.39% accuracy and 67.43% F1 score. This accuracy is already high,  
226 given the large amount of available features and stochastic nature of the signaling process.227 To ensure grounded use of signaling features, we implement a deterministic feature selection strat-  
228 egic: only features with intensity  $s_j \geq \alpha$  are retained. In our implementation, we use the threshold  
229  $\alpha = 1/2^2$  and define the resulting set of *marketable features* as:

230  
231 Marketable Features:  $\mathcal{S}_1(\mathbf{x}) = \{S_j : s_j(\mathbf{x}) \geq \alpha\}$ . (2)  
232

233 4.2 PERSONALIZATION MODULE: ALIGNING WITH PREFERENCES  
234235 This stage aims to steer persuasive language generation toward buyer preferences—another core  
236 objective of grounded persuasion. Our solution involves two steps.237 First, we elicit user preferences and structure them in a usable form. On platforms like Zillow or  
238 Redfin, this could be done using mature machine learning methods based on user browsing behavior.  
239 Without access to such data, we instead design a preference elicitation process within our human-  
240 subject evaluation framework. Specifically, our web interface prompts an LLM to simulate a realtor,  
241 guiding participants through questions to identify their most valued features. Each user then rates  
242 the importance of each feature  $S_j$  with a score  $r_j$  prior to the evaluation tasks. While simple, this  
243 approach suffices to support a persuasive AI Realtor that effectively adapts to user preferences, as  
244 demonstrated in our experiments.245 Second, we select a personalized subset of features to shift user beliefs positively. Since real-world  
246 marketing texts are not tailored to individual users, we cannot rely on them to provide supervision for  
247 personalization. Instead, we use a scoring function that combines population-level feature intensity  
248  $\mathbf{s}(\mathbf{x})$  with individual preference ratings  $\mathbf{r}$ , selecting features above a threshold  $\alpha$ :

249 Personalized Features:  $\mathcal{S}_2(\mathbf{x}) = \{s_j \mid s_j(\mathbf{x}) + c(r_j - r_0) \geq \alpha\}$ ,  
250

251 where  $c$  reflects the strength of personalization and  $r_0$  is a baseline rating. These features are then  
252 passed to the LLM, which determines how best to incorporate them into the generated text.253 4.3 MARKETING MODULE: CAPTURING SURPRISE VIA RAG  
254255 The last stage is designed to better ground persuasive language generation in factual evidence, prob-  
256 lem contexts and localized information in automated marketing. Our design here is inspired by rich  
257 marketing strategy research (Lindgreen & Vanhamme, 2005; Ludden et al., 2008; Ely et al., 2015),  
258 which have shown that buyers would derive entertainment utility from *surprising* effects/features  
259 and have a deeper impression. In our setting of real estate marketing, such surprising features are  
260 those that are relatively rare compared to their surrounding area. Formally, we determine a set of  
261 surprising features based on their percentile in the feature distribution as follows,

262 Surprising Features:  $\mathcal{S}_3(\mathbf{x}) = \{S_j \subset \mathcal{S}_1 : s_j(\mathbf{x}) \text{ is within } \beta\text{-quantile of distribution } s_j(\mu)\}$ .  
263

264  
265 <sup>1</sup>We also experiment with several other baselines for feature extraction, including prompting LLMs directly  
266 and applying simple pooling over embedding vectors. The strongest baseline achieves approximately 59% F1  
267 score, which is substantially lower than the final model used in our grounding module. For simplicity, we only  
268 reported the final model’s performance in the main text.269 <sup>2</sup>The feature existence threshold  $\alpha$  was determined through a grid search over the range  $[0.1, \dots, 0.9]$ , with  
270 performance evaluated using the F1 score on a held-out, human-annotated validation set.  $\alpha = 0.5$  yielded the  
271 best trade-off between precision and recall.

270 This gives the LLMs localized feature information at different levels of granularity obtained through  
 271 Retrieval Augmented Generation (RAG) (Lewis et al., 2020).<sup>3</sup> Such behavioral economics-driven  
 272 design proves to be highly effective; citing one of the human subjects in our experiment (see the full  
 273 description in Appendix A.1), who was asked about why they liked a listing description (without  
 274 knowing it was AI-generated):

275 *...Description B specifically points out the rarity of the ample storage and built-in cabinetry in  
 276 similarly priced listings, making the property stand out.*

## 279 5 EVALUATIONS

### 282 5.1 EVALUATION BY HUMAN FEEDBACK

284 To evaluate the effectiveness of listing descriptions generated by different models, we draw inspiration  
 285 from ChatArena (Zheng et al., 2023) and conduct an online survey to collect pairwise human  
 286 feedback comparing different models’ outputs. In summary, systematic evaluation by human feed-  
 287 back shows that our  AI Realtor clearly outperforms human experts and other model variants,  
 288 measured by standard Elo ratings (Elo, 1967). Below, we detail the design of our user survey plat-  
 289 form, baseline setup, and evaluation metrics, followed by a report on the human evaluation results.

290 **Quality Assurance** We focus on the major US city *Chicago*<sup>4</sup> with a highly active housing market.  
 291 We recruit about 100 participants from the popular *Prolific* platform for human-subject experiments,  
 292 selecting in-state residents familiar with Chicago’s housing market and curating approxi-  
 293 mately 1,000 listings of varied sizes and price ranges. Each human subject is tasked with comparing  
 294 10 pairs of house descriptions. During each comparison, the human subject sees pictures and all ba-  
 295 sic information about a house, and then faces two listing descriptions without knowing what methods  
 296 (human realtor or AI agents) generate them, and is asked to choose which description is preferred,  
 297 and by how much (see Appendix B.3 for details). Notably,  AI Realtor generates personalized  
 298 descriptions on the fly for each human subject, based on their preferences elicited while they join  
 the survey (see Appendix B.2 for details).

299 To ensure feedback quality, we implement several measures: (1) *Screening tests* to confirm par-  
 300 ticipants can extract information from listings and follow specific home search motives (See Ap-  
 301 pendix B.1 for details); (2) *Attention checks* using pairs of nearly identical descriptions to ensure  
 302 participants carefully compare and identify differences; (3) *Control experiments* where participants  
 303 compare human-written, engaging descriptions against LLM-generated descriptions intentionally  
 304 prompted to be plain and unappealing, verifying their ability to favor high-quality descriptions; and  
 305 (4) *Incentives* on the platform, including bonus payments and requests for written reasoning behind  
 306 choices, to encourage consistent, well-justified feedback.

307 **Metrics** We adopt the Elo rating score as our main metric. We use a typical choice of the initial  
 308 Elo rating as 1000, scaling parameter  $c = 400$ , and learning rate  $K = 32$ . The win rate for a model  
 309 with Elo rating  $e_1$  against a model with rating  $e_0$  is calculated as  $[1 + 10^{(e_0 - e_1)/c}]^{-1}$ .

310 **Baseline Models** In addition to our primary persuasion model  AI Realtor, we evaluate  
 311 several baseline models, including: *Vanilla*, an LLM prompted with all attributes of the listing; *SFT*,  
 312 an LLM fine-tuned with supervised training and prompted with all features of the listing; *Human*,  
 313 listing descriptions sourced from Zillow, written by professional realtors; *Control*, the model used  
 314 in the control experiment described earlier. We also include two ablation models based on  AI  
 315 Realtor: one that only uses the marketable feature from the Grounding module, the other excludes  
 316 surprisal features from the Marketing module. Additionally, we experiment with two LLM variants,  
 317 GPT-4o and GPT-4o-mini, while keeping the prompt instructions consistent across models.

318  
 319 <sup>3</sup>In our implementation, we implement the sparse retrieval part via ElasticSearch (<https://www.elastic.co/elasticsearch>) and retrieve Top 10 listings with the most similar features.

320  
 321 <sup>4</sup>Chicago has been established by various economic and sociological literature (Levitt & Syverson, 2008;  
 322 Sampson, 2012; Grabinsky & Reeves, 2015) as a rigorous proxy for broader American urban mechanics. Also,  
 323 Chicago has a diverse set of listings, compared to major cities in the US, that can reliably test our models’  
 324 performance across various scenarios. See Appendix D.3 for more analysis.

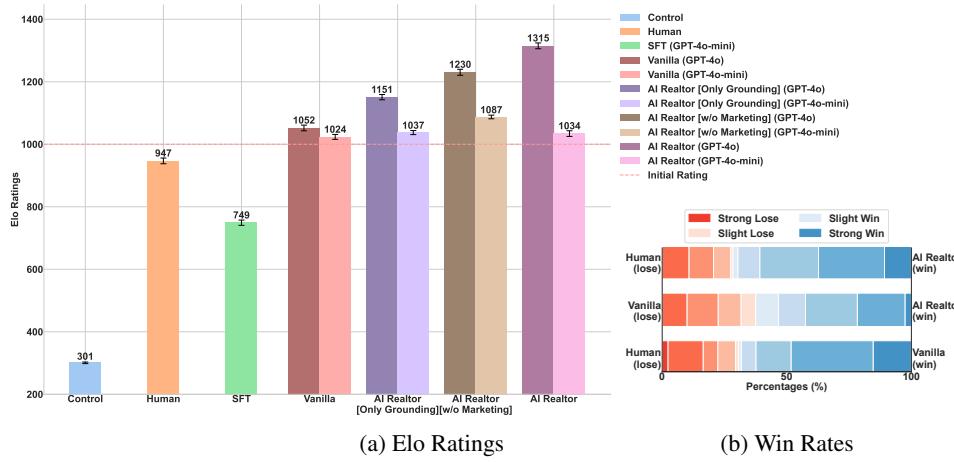


Figure 3: Comparison of model performance using Elo ratings and win rates. Elo ratings represent overall persuasiveness, and win rates reflect relative persuasiveness. Both metrics are based on evaluations by human subjects. [Confidence Interval is computed based on 500 bootstrap runs by adapting the Elo implementation from Chatbot Arena \(Chiang et al., 2024\)](#)

**Results** We plot the Elo ratings of different models in Figure 3a. The results reflect a clear trend: while vanilla GPT-4o performs on par with humans (1052 vs 947), each of our designed module enhancement progressively improves the persuasiveness of the generation, ultimately surpassing human performance with a clear margin (1318 vs 947). To ensure a fair comparison against human descriptions, which do not have access to explicit user preferences, we note that our model variant without any personalization (*Only Grounding*) still significantly outperforms human-written content (1151 vs 947). Also we observe that using GPT-4o to generate listing description does have a clear edge compared to that of GPT-4o-mini. Moreover, we plot empirical win rates among three major competitors (*Vanilla*, *Human* and *AI Realtor*) in Figure 3b, which directly illustrates how much *AI Realtor* outperforms the other two.<sup>5</sup> Please see Appendix A for case studies of our model-generated descriptions with more nuanced observations.

## 5.2 EVALUATION THROUGH AI FEEDBACK

Human feedback can be costly, especially as we scale the training and evaluation of our task. In this section, we report our empirical evaluation by using AIs to simulate human feedback based on our data collected from the above human-subject experiments.

**Simulation Setup** We employ an LLM to simulate the responses of buyers in the previous experiment. We use the first  $K$  pairwise comparison results as  $K$ -shot in-context learning samples and prompt the LLM to predict the same buyer’s selections for the remaining samples. We also adopt the chain-of-thought prompting format (Wei et al., 2022) and provide the buyer’s rationale comments as the information for in-context learning (see Appendix F.7 for the exact prompt). We use the Sotopia framework (Zhou et al., 2024) to configure this simulation agent with GPT-4o-mini (OpenAI, 2024b) as the base model.

**Metrics** We use two metrics to evaluate the reliability of AI feedback compared to human feedback: 1) *Shot-wise Simulation Accuracy (SSA)*: the prediction accuracy averaged across users for each shot; 2) *User-wise Simulation Accuracy (USA)*: the prediction accuracy for each user, averaged across #shots. The first metric measures overall simulation accuracy across the entire population, while the second one measures simulation accuracy for each user.

**Effectiveness of AI Feedback** The simulation results under both metrics are shown in Figure 4a and 4b. The model achieves 61.6% accuracy across users and exhibits non-trivial ( $> 50\%$ ) performance for 79.2% of users, suggesting potential for leveraging AI feedback. However, the accuracy remains unsatisfactory for reliable evaluation. Additionally, the variance in the USA metric is high and increases with more provided shots, underscoring the challenges of personality simulation, as

<sup>5</sup>Participants also rated their preference for each description on a 1-5 scale.

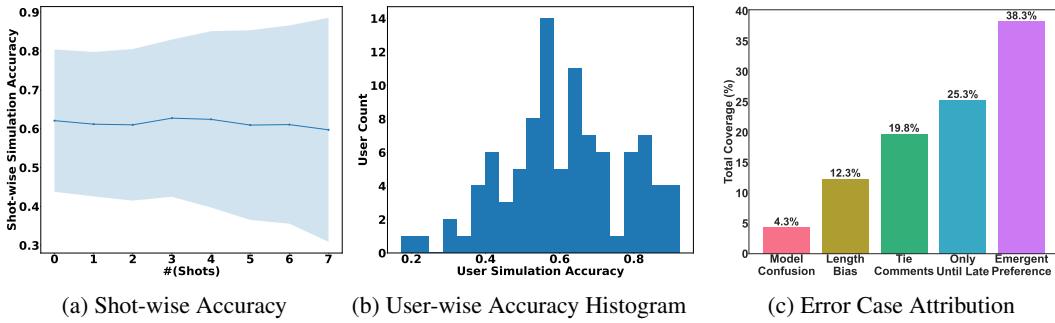


Figure 4: Analyses of Simulating Human Feedback with AI Feedback.

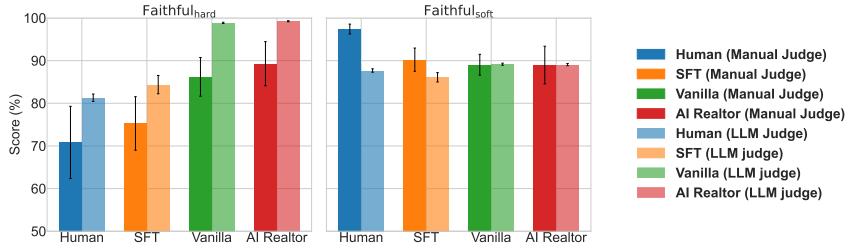


Figure 5: Faithfulness Scores for Hallucination Checks.

highlighted in (Wang et al., 2024). While the upward trend in variance is expected due to fewer data points, it highlights the difficulty of predicting user preferences dynamically.

To further understand the limitations of AI-simulated feedback, we conduct a manual analysis of simulation errors. Excluding the 56.1% error cases that lack clearly explainable patterns, we attribute the rest of them to several key error sources in Figure 4c: 1) *Length Bias*: Similar to the observation in Chatbot Arena (Zheng et al., 2023), the model overly favors longer responses; 2) *Tie Comments*: Buyers consider the influence from descriptions as indifferent yet still cast confident votes in one of the choices; 3) *Emergent Preference*: While the model only has access to a buyer’s pre-established preference, a buyer’s selections in some cases reflect some unspecified preferences or ones in contradiction; 4) *Only Until Late*: Correct predictions about a buyer’s selection only emerge after sufficient in-context samples; 5) *Model Confusion*: The model’s prediction appears random, which indicates that the model may not have sufficient information to simulate such a buyer. Some of these errors can be mitigated by collecting more selection data from each buyer or improving the preference elicitation process in future work.

### 5.3 HALLUCINATION CHECKS

For grounded persuasion, it is important to ensure minimal risks of hallucination. Hence, we evaluate the amount of misinformation in the marketing content through fine-grained fact-checking (Min et al., 2023), where we use GPT-4o to assist our hallucination check and set the listing attributes in the dataset as atomic facts. Specifically, we consider two types of factual attributes to check,  $X_{\text{hard}}$  and  $X_{\text{soft}}$ . For attributes in  $X_{\text{hard}}$ , we require the attribute description to be completely accurate (e.g., #(bathrooms)), whereas we allow attributes in  $X_{\text{soft}}$  to be roughly accurate (e.g., address).

Given an attribute set  $X$  and a description  $L$ , we ask the model to perform the following tasks:  $\text{supp}(L, X)$  identifies the subset of attributes in  $X$  that are mentioned in  $L$ ;  $\text{eval}_{\text{hard}}(L, x)$  returns a binary value indicating whether attribute  $x$  is accurately described; and  $\text{eval}_{\text{soft}}(L, x)$  provides a score from 0 to 10 reflecting the extent to which  $x$  is accurately described (see our prompt design in Appendix E). We then compute the faithfulness score for attributes in  $X_{\text{hard}}$  and  $X_{\text{soft}}$  as follows,

$$\text{Faithful}_{\text{hard}}(L) = \frac{\sum_{x \in \text{supp}(L, X_{\text{hard}})} \text{eval}_{\text{hard}}(L, x)}{|\text{supp}(L, X_{\text{hard}})|}, \quad \text{Faithful}_{\text{soft}}(L) = \frac{\sum_{x \in \text{supp}(L, X_{\text{soft}})} \text{eval}_{\text{soft}}(L, x) / 10}{|\text{supp}(L, X_{\text{soft}})|}.$$

432 As shown in Figure 5, the model-generated descriptions are mostly faithful to listing information  
 433 with minimal hallucination under both metrics. In contrast, the descriptions from human realtors  
 434 or SFT model show an even higher level of hallucination. After digging into details, we found  
 435 that this is due to human realtors’ (also SFT’s) vague description of attributes in  $X_{\text{hard}}$  such as the  
 436 following example, “*This 4 bedroom, 3.5 bathroom home offers nearly 2,000 (1,828) sqft of living  
 437 space...*”. Our  AI Realtor, however, tends to accurately describe factual attributes whenever  
 438 mentioned, likely due to its preference to copy from context — interestingly, this preference seems  
 439 to be forgotten by the model after supervised fine-tuning on human-written descriptions. That said,  
 440 it is debatable whether such vague descriptions of attributes is a true kind of hallucination, though  
 441 some buyers did complain about this kind of language in the comments of their responses.  
 442

443 We replicate hallucination checks with human evaluators to validate GPT-4o’s hallucination detection  
 444 results. Details of the interface and annotation guidelines are provided in Appendix E.2, and the  
 445 results are shown in Figure 5. **Regarding ranking consistency, GPT-4o’s relative ordering of models**  
 446 **on  $X_{\text{hard}}$  aligns closely with human evaluations, but diverges on  $X_{\text{soft}}$** , highlighting the challenge of  
 447 verifying loosely matched factual attributes. Overall, both human and GPT-4o evaluations show that  
 448  AI Realtor achieves higher faithfulness on  $X_{\text{hard}}$  and comparable performance on  $X_{\text{soft}}$ , suggesting  
 449 it poses minimal risk of hallucination. Furthermore, the human evaluators report that  AI Realtor descriptions are as trustworthy as humans (See more details of our credibility survey in  
 450 Appendix E.2).

## 451 6 RELATED WORK

452 Several studies have pioneered methods in computational linguistics for understanding and measuring  
 453 persuasiveness (Wang et al., 2019; Wei et al., 2016; Tan et al., 2016). The advent of large  
 454 language models (LLMs) has further spurred research into their persuasive capabilities, especially  
 455 as part of frontier model risk assessments by developers (Durmus et al., 2024; Hurst et al., 2024;  
 456 Jaech et al., 2024). A major focus has been on the potential for LLM-generated propaganda in  
 457 politically sensitive contexts (Voelkel et al., 2023; Goldstein et al., 2024; Hackenburg et al., 2024;  
 458 Luciano, 2024). Parallel investigations examine settings such as personalized persuasion (Hacken-  
 459 burg & Margetts, 2024; Salvi et al., 2024; Matz et al., 2024). Breum et al. (2024) and multi-round  
 460 persuasion (Breum et al., 2024). Takayanagi et al. (2025) assess the influence of GPT-4’s ability to  
 461 generate financial analyses to audiences. Complementary research has probed related LLM capa-  
 462 bilities including negotiation (Bianchi et al., 2024), debate (Khan et al., 2024), sycophancy (Sharma  
 463 et al., 2023; Denison et al., 2024), as well as the emergence of strategic rationality in game-theoretic  
 464 settings (Chen et al., 2023; Raman et al., 2024).

465 In a similar application domain, Angelopoulos et al. (2024) conduct an experiment to generate mar-  
 466 keting email with a fine-tuned LLM and report a 33% improvement in email click-through rates  
 467 compared to human expert baselines. Singh et al. (2024) design an evaluation benchmark based  
 468 on a dataset of tweet pairs with similar content but different wording and like counts. In compari-  
 469 son, our work develops a full agentic solution for automated marketing from learning domain expert  
 470 knowledge to crafting localized features, which significantly outperforms the model with supervised  
 471 fine-tuning in our human-subject experiments.

## 472 7 DISCUSSION

473 **Contributions and Implications** This paper presents a novel framework for persuasive language  
 474 generation, marking a first step toward integrating signaling schemes from economic theory into  
 475 agentic LLM design. Our results demonstrate that this structured approach can achieve superhuman  
 476 persuasive performance in a high-stakes domain like real estate marketing. A central tenet of our  
 477 design is the deliberate prioritization of factual grounding. While human-written descriptions often  
 478 employ stylized or emotionally resonant language, we argue that in domains where accuracy is  
 479 paramount, constraining generation to verifiable facts is a necessary and responsible choice. Our  
 480 framework’s effectiveness stems from its ability to map raw attributes to a compact set of high-level,  
 481 market-relevant features, ensuring that the generated content is both persuasive and credible.

482 **Limitations and Future Directions** Despite these promising results, we acknowledge several lim-  
 483 itations that highlight avenues for future research. The primary bottleneck remains the reliance on

486 high-quality human feedback for evaluation. Our experiments with automated, LLM-based eval-  
 487 uators show promise for assessing factuality but are not yet reliable for measuring nuanced qualities  
 488 like persuasiveness, underscoring the need for more sophisticated evaluation benchmarks. **Second,**  
 489 **our empirical validation is currently concentrated on the Chicago market. While our agentic work-**  
 490 **flow is designed to admit localization, persuasion is culturally and economically sensitive; future**  
 491 **work is required to verify stability across diverse geographic regions and demographics. Further-**  
 492 **more, generalizing this framework to domains with less structured inputs or different persuasive**  
 493 **norms (e.g., brand marketing vs. legal arguments) presents a significant and important challenge.**

494 Building on this foundation, several exciting directions emerge. The modularity of our framework  
 495 is well-suited for incorporating domain-specific constraints. For regulated fields like housing or  
 496 finance, integrating compliance filters or legal principles inspired by approaches like Constitutional  
 497 AI (Bai et al., 2022) is a crucial next step to ensure responsible deployment. Moreover, to address  
 498 the trade-off between factuality and expressiveness, future work could explore incorporating a wider  
 499 range of persuasion theories, such as emotional appeals and narrative structures, as controllable  
 500 modules within the agentic design. Finally, scaling our datasets, expanding to new copywriting  
 501 domains, and conducting more extensive real-world A/B testing will be essential to fully unlock the  
 502 potential of theory-grounded persuasive generation.

## 503 ETHICS STATEMENT

504 Our research on persuasive language generation acknowledges the dual-use nature of such technolo-  
 505 gies. We have proactively centered our work on grounded persuasion, where generated content is  
 506 constrained by verifiable facts, to mitigate the risks of misinformation and manipulation. Our exten-  
 507 sive hallucination checks, detailed in § 5.3 and Appendix E, confirm that our agent maintains a high  
 508 degree of factual accuracy, comparable to or exceeding that of human experts.

509 All human-subject experiments were conducted in compliance with ethical research standards. The  
 510 study protocol received IRB approval (exempt). Participants were recruited from the Prolific plat-  
 511 form, informed of the study’s purpose, and compensated at a fair rate (approximately \$20/hour with  
 512 performance incentives). The dataset, derived from publicly available Zillow listings, was processed  
 513 to remove any personally identifiable information, ensuring user privacy.

514 By focusing on a high-stakes, fact-driven domain like real estate, we aim to provide a framework  
 515 for developing responsible persuasive AI. We believe this work serves as a foundation for future  
 516 research into the ethical guardrails necessary for deploying strategic language models in real-world  
 517 applications and encourage continued investigation into their broader societal implications.

## 521 REPRODUCIBILITY STATEMENT

522 We are committed to ensuring the reproducibility of our research. Below, we outline the resources  
 523 available to replicate our findings.

524 **Data.** The core dataset was constructed from publicly available real estate listings from Zillow.  
 525 The raw attribute schema, data curation process, and final feature schema are detailed in Appendix D.  
 526 The collected human-subject evaluation data and feature annotations will be made publicly available  
 527 upon publication.

528 **Methodology and Code.** The theoretical framework is described in § 3. The complete agentic  
 529 pipeline, including the implementation of the Grounding, Personalization, and Marketing modules,  
 530 is detailed in § 4 and Appendix C. To facilitate replication, all prompts used for LLM-based feature  
 531 extraction, schema generation, and persuasive content creation are provided in Appendix F. The full  
 532 source code for our agent and evaluation framework will be released publicly.

533 **Evaluation.** Our human-subject evaluation framework, including the design of the web interface,  
 534 participant screening, and preference elicitation process, is fully described in § 5.1 and Appendix B.  
 535 The metrics used, including Elo rating calculations and hallucination checks, are also detailed in § 5.

540 **USAGE OF LARGE LANGUAGE MODELS**  
541542 In this work, we mainly use LLMs for the following purposes:  
543544 1. Aid or Polish Writing (Gemini 2.5 Pro, ChatGPT 4/5)  
545 2. Literature Retrieval and Discovery (e.g., finding related work) (Gemini 2.5 Pro Deep Re-  
546 search, ChatGPT Deep Research)  
547 3. Assisting Code Writing and Debugging (Claude 3.5 Sonnet)  
548549 We fully understand the responsibility of using LLMs in academic research. We carefully monitor  
550 any potential problems, such as plagiarism or scientific misconduct (e.g., fabrication of facts) when  
551 using LLMs. We make sure these problems do not occur in the paper.  
552553 **REFERENCES**  
554

555 George A Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. In  
556 *Uncertainty in economics*, pp. 235–251. Elsevier, 1978.

557

558 Panagiotis Angelopoulos, Kevin Lee, and Sanjog Misra. Causal alignment: Augmenting language  
559 models with a/b tests. *Available at SSRN*, 2024.

560 Anthropic. Claude 3.5 sonnet, 2024. URL <https://www.anthropic.com/news/claude-3-5-sonnet>. AI language model.

561

563 Gerry Antioch. Persuasion is now 30 per cent of us gdp: Revisiting mccloskey and klamer after a  
564 quarter of a century. *Economic Round-up*, (1):1–10, 2013.

565

566 Kenneth J Arrow. The economics of information: An exposition. *Empirica*, 23(2):119–128, 1996.

567

568 Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones,  
569 Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harm-  
570 lessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.

571

572 Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

573

574 Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.

575

576 Federico Bianchi, Patrick John Chia, Mert Yuksekgonul, Jacopo Tagliabue, Dan Jurafsky, and James  
577 Zou. How well can llms negotiate? negotiationarena platform and analysis. *arXiv preprint arXiv:2402.05863*, 2024.

578

579 Steven Bird, Ewan Klein, and Edward Loper. *Natural language processing with Python: analyzing*  
580 *text with the natural language toolkit*. ” O’Reilly Media, Inc.”, 2009.

581

582 Simon Martin Breum, Daniel Vædele Egdal, Victor Gram Mortensen, Anders Giovanni Møller,  
583 and Luca Maria Aiello. The persuasive power of large language models. In *Proceedings of the*  
584 *International AAAI Conference on Web and Social Media*, volume 18, pp. 152–163, 2024.

585

586 Yiting Chen, Tracy Xiao Liu, You Shan, and Songfa Zhong. The emergence of economic rationality  
587 of gpt. *Proceedings of the National Academy of Sciences*, 120(51):e2316205120, 2023.

588

589 Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li,  
590 Dacheng Li, Banghua Zhu, Hao Zhang, Michael Jordan, Joseph E Gonzalez, et al. Chatbot  
591 arena: An open platform for evaluating llms by human preference. In *Forty-first International*  
592 *Conference on Machine Learning*, 2024.

593 Brian L Connelly, S Trevis Certo, R Duane Ireland, and Christopher R Reutzel. Signaling theory: A  
review and assessment. *Journal of management*, 37(1):39–67, 2011.

594 Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks,  
 595 Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, et al. Sycophancy to subterfuge:  
 596 Investigating reward-tampering in large language models. *arXiv preprint arXiv:2406.10162*,  
 597 2024.

598 Angelika Dimoka, Yili Hong, and Paul A Pavlou. On product uncertainty in online markets: Theory  
 599 and evidence. *MIS quarterly*, pp. 395–426, 2012.

600 Esin Durmus, Liane Lovitt, Alex Tamkin, Stuart Ritchie, Jack Clark, and Deep Ganguli. Measuring  
 601 the persuasiveness of language models, 2024.

602 Arpad E Elo. The proposed uscf rating system, its development, theory, and applications. *Chess*  
 603 *life*, 22(8):242–247, 1967.

604 Jeffrey Ely, Alexander Frankel, and Emir Kamenica. Suspense and surprise. *Journal of Political*  
 605 *Economy*, 123(1):215–260, 2015.

606 Josh A Goldstein, Jason Chao, Shelby Grossman, Alex Stamos, and Michael Tomz. How persuasive  
 607 is ai-generated propaganda? *PNAS nexus*, 3(2):pgae034, 2024.

608 Jonathan Grabinsky and Richard V Reeves. The most american city: Chicago, race, and  
 609 inequality. *Retrieved from Brookings: https://www. brookings. edu/blog/social-mobility-*  
 610 *memos/2015/12/21/the-mostamerican-city-chicago-race-and-inequality*, 2015.

611 Sanford J Grossman. The informational role of warranties and private disclosure about product  
 612 quality. *The Journal of law and Economics*, 24(3):461–483, 1981.

613 Kobi Hackenburg and Helen Margetts. Evaluating the persuasive influence of political microtar-  
 614 geting with large language models. *Proceedings of the National Academy of Sciences*, 121(24):  
 615 e2403116121, 2024.

616 Kobi Hackenburg, Ben M Tappin, Paul Röttger, Scott Hale, Jonathan Bright, and Helen Margetts.  
 617 Evidence of a log scaling law for political persuasion with large language models. *arXiv preprint*  
 618 *arXiv:2406.14508*, 2024.

619 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-  
 620 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*  
 621 *arXiv:2410.21276*, 2024.

622 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec  
 623 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv*  
 624 *preprint arXiv:2412.16720*, 2024.

625 Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bam-  
 626 ford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al.  
 627 Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.

628 Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101  
 629 (6):2590–2615, 2011.

630 Akbir Khan, John Hughes, Dan Valentine, Laura Ruis, Kshitij Sachan, Ansh Radhakrishnan, Ed-  
 631 ward Grefenstette, Samuel R Bowman, Tim Rocktäschel, and Ethan Perez. Debating with more  
 632 persuasive llms leads to more truthful answers. *arXiv preprint arXiv:2402.06782*, 2024.

633 Pablo Kurlat and Florian Scheuer. Signalling to experts. *The Review of Economic Studies*, 88(2):  
 634 800–850, 2021.

635 Steven D Levitt and Chad Syverson. Market distortions when agents are better informed: The value  
 636 of information in real estate transactions. *The Review of Economics and Statistics*, 90(4):599–611,  
 637 2008.

638 Gregory Lewis. Asymmetric information, adverse selection and online disclosure: The case of ebay  
 639 motors. *American Economic Review*, 101(4):1535–1546, 2011.

648 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
 649 Heinrich Kütller, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented genera-  
 650 tion for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:  
 651 9459–9474, 2020.

652 Adam Lindgreen and Joelle Vanhamme. Viral marketing: The use of surprise. *Advances in electronic  
 653 marketing*, pp. 122–138, 2005.

654 Karl-Gustaf Löfgren, Torsten Persson, and Jörgen W Weibull. Markets with asymmetric informa-  
 655 tion: the contributions of george akerlof, michael spence and joseph stiglitz. *The Scandinavian  
 656 Journal of Economics*, pp. 195–211, 2002.

657 Floridi Luciano. Hypersuasion—on ai’s persuasive power and how to deal with it. *Philosophy &  
 658 Technology*, 37(2):1–10, 2024.

659 Geke DS Ludden, Hendrik NJ Schifferstein, and Paul Hekkert. Surprise as a design strategy. *Design  
 660 Issues*, 24(2):28–38, 2008.

661 SC Matz, JD Teeny, Sumer S Vaid, H Peters, GM Harari, and M Cerf. The potential of generative  
 662 ai for personalized persuasion at scale. *Scientific Reports*, 14(1):4692, 2024.

663 Rui Meng, Ye Liu, Shafiq Rayhan Joty, Caiming Xiong, Yingbo Zhou, and Semih Yavuz.  
 664 Sfrembedding-mistral: enhance text retrieval with transfer learning. *Salesforce AI Research Blog*,  
 665 3, 2024.

666 Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke  
 667 Zettlemoyer, and Hannaneh Hajishirzi. FActScore: Fine-grained Atomic Evaluation of Factual  
 668 Precision in Long Form Text Generation. In *Proceedings of the 2023 Conference on Empirical  
 669 Methods in Natural Language Processing*, pp. 12076–12100, 2023.

670 OpenAI. Gpt-4o, 2024a. Available at: <https://openai.com/index/hello-gpt-4o/>.

671 OpenAI. Gpt-4o mini: Advancing cost-efficient intelligence, July 2024b. URL <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>. Accessed: 2024-09-19.

672 Narun Krishnamurthi Raman, Taylor Lundy, Samuel Joseph Amouyal, Yoav Levine, Kevin Leyton-  
 673 Brown, and Moshe Tennenholtz. Steer: Assessing the economic rationality of large language  
 674 models. In *Forty-first International Conference on Machine Learning*, 2024.

675 Francesco Salvi, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West. On the conversa-  
 676 tional persuasiveness of large language models: A randomized controlled trial. *arXiv preprint  
 677 arXiv:2403.14380*, 2024.

678 Robert J Sampson. *Great American city: Chicago and the enduring neighborhood effect*. University  
 679 of Chicago press, 2012.

680 Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R Bow-  
 681 man, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R Johnston, et al. Towards under-  
 682 standing sycophancy in language models. *arXiv preprint arXiv:2310.13548*, 2023.

683 Somesh Singh, Yaman K Singla, Harini SI, and Balaji Krishnamurthy. Measuring and improving  
 684 persuasiveness of large language models. *arXiv preprint arXiv:2410.02653*, 2024.

685 Michael Spence. Job market signaling. In *Uncertainty in economics*, pp. 281–306. Elsevier, 1978.

686 Joseph E Stiglitz. The theory of “ screening,” education, and the distribution of income. *The Ameri-  
 687 can economic review*, 65(3):283–300, 1975.

688 Takehiro Takayanagi, Hiroya Takamura, Kiyoshi Izumi, and Chung-Chi Chen. Can gpt-4 sway  
 689 experts’ investment decisions? In *Findings of the Association for Computational Linguistics:  
 690 NAACL 2025*, pp. 374–383, 2025.

702 Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. Winning argu-  
 703 ments: Interaction dynamics and persuasion strategies in good-faith online discussions. In *Pro-  
 704 ceedings of the 25th international conference on world wide web*, pp. 613–624, 2016.  
 705

706 Reddit User. Chatgpt helped me save \$50k buying/selling a house. [https://www.reddit.com/r/ChatGPT/comments/12z8g31/chatgpt\\_helped\\_me\\_save\\_50k\\_buyingselling\\_a\\_house/](https://www.reddit.com/r/ChatGPT/comments/12z8g31/chatgpt_helped_me_save_50k_buyingselling_a_house/), 2023. [Online; posted April 27, 2023].  
 707

708

709 Jan G Voelkel, Robb Willer, et al. Artificial intelligence can persuade humans on political issues.  
 710 2023.

711

712 Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li, and Yuandong  
 713 Tian. Learning personalized alignment for evaluating open-ended text generation. In *Pro-  
 714 ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 13274–  
 715 13292, 2024.

716

717 Xuewei Wang, Weiyang Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu.  
 718 Persuasion for good: Towards a personalized persuasive dialogue system for social good. *arXiv  
 preprint arXiv:1906.06725*, 2019.

719

720 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
 721 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in  
 722 neural information processing systems*, 35:24824–24837, 2022.

723

724 Zhongyu Wei, Yang Liu, and Yi Li. Is this post persuasive? ranking argumentative comments in  
 725 online forum. In Katrin Erk and Noah A. Smith (eds.), *Proceedings of the 54th Annual Meeting  
 726 of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 195–200, Berlin,  
 727 Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-2032.  
 728 URL <https://aclanthology.org/P16-2032/>.

729

730 Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, and  
 731 Yu Su. Travelplanner: A benchmark for real-world planning with language agents. *arXiv preprint  
 732 arXiv:2402.01622*, 2024.

733

734 Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable  
 735 real-world web interaction with grounded language agents. *Advances in Neural Information Pro-  
 736 cessing Systems*, 35:20744–20757, 2022.

737

738 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,  
 739 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and  
 740 chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.

741

742 Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Zhengyang Qi, Haofei Yu, Louis-Philippe  
 743 Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, and Maarten Sap. Sotopia: Interactive  
 744 evaluation for social intelligence in language agents. 2024. URL <https://openreview.net/forum?id=mM7VurbA4r>.

745

746 Zillow. Creative real estate listing descriptions. <https://www.zillow.com/agents/writing-real-estate-descriptions/>, n.d. Accessed: 2025-11-20.

747

## A CASE STUDIES

### A.1 USER FEEDBACK ON GENERATED DESCRIPTIONS WITH SURPRISING FEATURES

748 Below is an example where the user considers the description generated by 📷 AI Realtor to be  
 749 more persuasive, specifically because it includes surprising features.  
 750

751

752 **Disfavored Description (📷 AI Realtor [w/o Grounding], Description A):** Welcome to  
 753 your dream condo in the vibrant Printers Row neighborhood! This stunning one-bedroom, one-  
 754 bathroom home is not only affordable but packed with modern amenities and stylish features, all  
 755

756  
757  
758  
759  
760  
761  
762  
763  
764  
765

for \$297,500. Nestled in the heart of Chicago, 547 S Clark St APT 405 offers unparalleled convenience with easy access to public transportation, shopping, dining, and the iconic Millennium Park. Built in 1998, this condo boasts a huge bedroom suite, hardwood flooring throughout, and an inviting gas fireplace. The newly upgraded stainless steel appliances and eye-catching granite countertops make the kitchen a chef's delight. Enjoy your morning coffee on the spacious private balcony, while additional storage and a walk-thru closet area cater to all your organizational needs. The property is equipped with central air and in-unit laundry for ultimate comfort and convenience. With a \$389 monthly HOA fee, this condo promises both affordability and a secure investment in the bustling heart of Chicago. Don't miss the opportunity to embrace a lifestyle of ease and style in this perfect one-bedroom sanctuary!

766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779

**Preferred Description (AI Realtor, Description B):** Discover the splendid charm of urban living at 547 S Clark St APT 405, a chic one-bedroom, one-bathroom condo in the heart of Chicago's desirable Printers Row. Priced at an enticing \$297,500, this condo harmonizes affordability with luxury, featuring an expansive 876 sqft layout adorned with elegant hardwood floors throughout and a cozy gas fireplace. Enjoy culinary delights with new stainless steel appliances and eye-catching granite countertops in the kitchen. **The huge bedroom suite boasts a walk-thru closet area, offering ample built-in cabinet space and additional storage – an rarity in similarly priced listings.** Revel in the tranquility of your spacious private balcony, perfect for unwinding with views of the bustling cityscape. Keep convenience at your fingertips with in-unit side-by-side laundry and parking in a secure garage. Built in 1998, this condo is a stone's throw away from Millennium Park and surrounded by public transportation, grocery stores, and cultural amenities, amplifying its urban attraction. With a manageable monthly HOA fee of \$389 and strong investment potential, this condo ensures financial peace of mind alongside lifestyle enrichment. Don't miss out on this prime opportunity in vibrant Cook County, where your ideal location meets impeccable structure and style.

780  
781  
782  
783

**User Comment:** I selected Description B because while both descriptions effectively showcase the property's appealing features, Description B employs a more evocative and sensory language, creating a stronger emotional connection and a more vivid image of the lifestyle the condo offers. Specific points of comparison:

- Language & Tone:

784  
785  
786  
787

Description B uses words like "splendid charm," "chic," and "harmonizes" to paint a picture of elegance and sophistication, creating a more aspirational tone.

Description A, while positive, uses more straightforward language, focusing on practicality and convenience.

- Emphasis on Space & Luxury:

788  
789

Description B highlights the "expansive 876 sqft layout," and the "huge bedroom suite," emphasizing the sense of space and luxury.

Description A also mentions the spaciousness but doesn't create as strong an image of grandeur

- Unique Selling Points:

790  
791  
792  
793

**Description B specifically points out the rarity of the ample storage and built-in cabinetry in similarly priced listings, making the property stand out.**

Description A focuses on the general convenience and modern amenities, which, while attractive, are not as unique.

- Lifestyle & Surroundings:

794  
795  
796  
797  
798  
799  
800

Description B paints a more vivid picture of the lifestyle the condo offers, inviting the buyer to "revel in the tranquility" of the balcony and highlighting the proximity to cultural amenities, creating a stronger sense of place.

Description A mentions the location and amenities but lacks the same level of detail and emotional connection.

- Overall:

801  
802  
803  
804  
805  
806

Both descriptions are well-written and informative, but **Description B's richer language, focus on unique features, and emphasis on lifestyle create a more compelling and emotionally resonant picture of the property.** It makes the condo feel more desirable and aspirational, which is likely to attract more potential buyers.

807  
808  
809

Meanwhile, we acknowledge that it is also challenging to generate appropriate language to express surprising features. Here is an example that confuses the user:

810

811

812

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

**Preferred Description (Human, Description A):** Step into a light-filled sanctuary where high ceilings and expansive windows create an open and airy atmosphere. The loft-style design seamlessly combines living and dining spaces, providing a versatile canvas for your personal style. Ascend to your personal rooftop oasis, where the cityscape becomes your backdrop. Whether you're hosting soirees under the stars or enjoying a quiet morning coffee, this terrace is a true escape from the ordinary. Features include hardwood floors throughout, in-unit washer and dryer and new heat/central air conditioning. Secure building with intercom system, a deeded garage parking space and large storage room. Situated in a boutique building, this condo provides the perfect blend of privacy and proximity to urban amenities. Easy access to public transportation, trendy cafes, and cultural hotspots make this location ideal for the urban dweller.

**Disfavored Description (AI Realtor, Description B):** Nestled in the vibrant West Town neighborhood of Chicago, this exquisite loft-style condo at 11 N Green St UNIT 4C offers a light-filled sanctuary that merges modern living with trendy ambiance. Boasting an impressive 1,200 sqft of space, it features expansive windows framing a stunning cityscape backdrop, high ceilings, and an open floor plan, complemented by sleek hardwood floors and a cozy wood-burning fireplace in the living room. Residents will relish the rare luxury of a personal rooftop oasis—ideal for chic entertaining or serene reflection. Positioned among the top 2% for amenities in Chicago, this condo includes in-unit laundry, an intercom system, garage parking, and a large storage room, making it both convenient and stylish. Priced at \$490,000, this move-in-ready home offers exceptional value with its central location and effortless access to public transportation, ideal for those seeking the perfect blend of urban convenience and sophisticated living. With its prime spot in Cook County and top-rated area amenities, this condo won't just be a place to live but a lifestyle to love!

**User Comment:** Description B says it is in the top 2% of amenities. What does that even mean. That is nonsense.

## A.2 THE FAILING CASES OF HUMAN-WRITTEN DESCRIPTIONS

We also found cases where users dislike human-written descriptions over model-generated descriptions. The rationales behind the user preferences demonstrate the advantage of model-generated descriptions.

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

**Disfavored Description (Human, Description A):** Don't miss out on this sunny 1 bedroom, 1 bathroom condo in the heart of Lincoln Park. Step in to the living space flooded with natural light from your private balcony windows. The open kitchen layout offers tons of storage space, updated stainless steel appliances, and a comfortable breakfast island. In your queen sized bedroom you'll find nicely appointed closets, and cozy carpeting. This perfect starter home features hardwood flooring throughout the living space, and a newer (2021) washer and dryer in unit. Storage included.

**Preferred Description (AI Realtor, Description B):** Welcome to 2220 N Sedgwick St APT 204, a charming one-bedroom condo nestled in the vibrant heart of Lincoln Park, Chicago's most coveted neighborhood. Priced competitively at \$310,000, this move-in-ready gem offers a harmonious blend of modern convenience and timeless elegance. Revel in the top-tier stainless steel appliances and sun-drenched open kitchen with a cozy breakfast island perfect for any culinary enthusiast. Rich hardwood flooring and cozy carpeting seamlessly blend aesthetics and comfort, while the queen-sized bedroom boasts ample storage. Enjoy serene moments on your private balcony overlooking the iconic neighborhood streets studded with historical charm. Built in 1900, this meticulously maintained brick structure highlights both character and longevity, ensuring a sound investment. With unparalleled access to public transportation and a safe, walkable community, this property meets the highest standards of city living. Enviable situated among Lincoln Park's garden-filled avenues, it's the ideal starter home for those who value location and convenience without compromising on style or safety.

**User Comment:** Description B doesn't even have the size, location, or other important facts.

864  
865  
866  
867  
868  
869

**Case Analysis:** Sometimes human descriptions even miss important facts, while descriptions generated by our models do not. We present a fine-grained fact-checking study to check whether there is a hallucination in § 5.3.

870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882

**Preferred Description (AI Realtor):** Welcome to 832 W Wrightwood Ave #3, an enchanting 2-bedroom, 1-bathroom condo nestled in the heart of Lincoln Park, Chicago's most prestigious neighborhood. Priced sensibly at \$450,000 and boasting a spacious 1,164 sqft of elegant living, this East Lincoln Park penthouse marries historical charm with contemporary amenities. Step inside to discover a warm ambiance highlighted by exposed brick, hardwood floors, and a cozy wood-burning fireplace. The remodeled eat-in island kitchen is an entertainer's dream, seamlessly flowing into a separate dining area perfect for intimate gatherings. With its skylight windows and bay windows, an abundance of natural light illuminates every corner. Enjoy the convenience of an in-unit laundry room, additional private storage, and central air without the high HOA fees typically found in comparable homes. The condo's prime location offers walkability to the vibrant amenities and serene lakefront of Lincoln Park, catering to every lifestyle need. A rare find in a top-tier location with superior accessibility and neighborhood charm, this condo promises both investment value and a delightful urban retreat. Don't miss the open house to experience this gem first-hand!

883  
884  
885  
886  
887  
888  
889  
890  
891

**Disfavored Description (Human):** WALK TO IT ALL!! THIS BRIGHT TWO BEDROOM, 1 BATHROOM EAST LINCOLN PARK PENTHOUSE W/DECK HAS EXPOSED BRICK, BAY WINDOWS AND A WOOD BURNING FIREPLACE; EAT-IN ISLAND KITCHEN OPENS TO MASSIVE 23' WIDE LIVING ROOM WITH A SEPARATE DINING AREA. THE UNIT HAS BEAUTIFUL HARDWOOD FLOORS THROUGHOUT, A HUGE MASTER SUITE WITH TONS OF CLOSET/STORAGE SPACE. OTHER FEATURES INCLUDE ADDITIONAL PRIVATE STORAGE, IN-UNIT LAUNDRY ROOM WITH SIDE BY SIDE W/D AND PARKING. KITCHEN REMODELED IN 2016, BATHROOM REMODELED IN 2020. NEW AC CONDENSER IN 2022.

892  
893  
894  
895

**User Comment:** I think this description is much better because it isn't in all caps, which feels like I'm getting yelled at.

**Case Analysis:** Human-drafted descriptions can look unpleasant.

896  
897

### A.3 THE DICHOTOMY OF USER PREFERENCES ON WRITING STYLES

898  
899  
900  
901

In § 5.1, we present the aggregated benchmark results to compare the persuasiveness of listing descriptions generated by different models. To get more qualitative insights into the strengths and weaknesses of different models, as well as the subjective nature of human feedback, we present a more detailed case study here.

902  
903  
904  
905  
906  
907

The first thing we noticed is the users' subtle preferences in **description length**: while some users like concise descriptions that directly go to the point, other users prefer longer descriptions because they want to know more details about the property they are interested. The following two examples of user feedback explain this point.

908  
909  
910  
911  
912  
913  
914  
915  
916  
917

**Preferred Description (Vanilla, Description A):** Welcome to your dream condo at 4345 S Indiana Ave UNIT 2N, nestled in the vibrant Bronzeville neighborhood of Chicago, IL. This exquisite 3-bedroom, 2-bath home offers 1,550 sqft of modern living infused with classic charm, all for an unbeatable price of \$275,000. Built in 2006, it features abundant natural light flooding through large windows, complemented by tall ceilings and an open living space. Imagine cozy evenings by the custom stone wood-burning fireplace or enjoying a morning coffee on your private second balcony. The master bedroom offers tranquility with a spacious walk-in closet, while the additional bedrooms provide generous space for family or guests. The kitchen is a chef's delight, equipped with stainless steel appliances including a range, microwave, and refrigerator. With central air cooling, hardwood flooring, and a sleek, contemporary style highlighted by recessed lighting, this condo is the perfect blend of comfort and sophistication. Adding to the allure, a

918  
919  
920  
921

secure garage parking spot is included. Security is assured with a modern security system, and the convenience of in-unit laundry completes this superb offering. Located in Cook County with easy access to all Chicago has to offer, this stylish condo is a must-see!

922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937

**Disfavored Description (AI Realtor, Description B):** Welcome to your dream home at 4345 S Indiana Ave UNIT 2N, nestled in the heart of the vibrant Bronzeville neighborhood in Chicago. This stunning condo offers the epitome of comfortable living with 3 spacious bedrooms, 2 modern bathrooms, and a living area of 1,550 square feet, perfectly situated for a single mother seeking convenience and safety. The residence exudes warmth, featuring abundant natural light through large windows and a cozy custom stone wood-burning fireplace in an open living setting. The condo is a gem within the community, boasting one of the top amenities packages in the area, including a stylish stainless steel kitchen, a rare second private balcony, and garage parking that ensures convenience. Step into the master bedroom for a touch of luxury, indulge in the modern ambiance provided by recessed lighting, or relax in the welcoming family room with its captivating atmosphere. Temperature comfort is assured through efficient central air and heating. Notably, this property towers above others in terms of walkability and neighborhood amenities, making it an ideal choice for a family-focused lifestyle. Priced attractively at \$275,000, it's a golden opportunity to secure a versatile home that evolves with your needs, ready to create cherished family memories. Discover the potential for a fulfilling life in a community known for its top-tier safety and accessibility, all while investing in a property you can pass down to the next generation.

938  
939  
940

**User Comment:** Description A gets to the point faster, while still highlighting the important qualities of the home.

**Case Analysis:** Some users love **concise** descriptions.

941  
942

**Preferred Description (Vanilla):** Welcome to 4454 S Shields Ave, a charming A-Frame single-family home nestled in the heart of Chicago's historic Fuller Park neighborhood. This inviting residence offers three cozy bedrooms and a well-appointed bathroom, all within a compact 956 square feet of open-concept living space that seamlessly combines comfort and style. Built in 1929, the home exudes classic character while featuring modern conveniences such as central air for cooling and a natural gas heating system. The property's allure is further enhanced by its unfinished basement, offering potential for personalized expansions. Imagine summer barbecues on your porch or taking a quick stroll to a nearby park, making this an ideal location for outdoor enthusiasts. With its proximity to local amenities and an incredible price of just \$219,900, this home represents a fantastic investment opportunity, especially with its rare, close-to-an-Olympic-sized swimming pool bonus. Discover the potential of this foreclosure property and make it your own urban oasis in Cook County.

943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953

**Disfavored Description (SFT):** Welcome to this charming single-family home nestled in Fuller park! This listing features an open concept, 3 bedrooms, 1 full bathroom, and an unfinished basement that's just waiting for your personal touch. Located close to a park with an Olympic-sized swimming pool, you'll have endless recreational opportunities at your doorstep. With its prime location and potential for expansion, this property is a true gem waiting to be polished. Don't miss the chance to make this house your dream home!

954  
955  
956  
957  
958  
959  
960  
961

**User Comment:** Again, more description is better if I am really interested in a property.

**Case Analysis:** Some users love **longer** descriptions.

962  
963  
964  
965  
966  
967

Another important factor is the **embellishment** of descriptions. That is, in our particular marketing domain, is there a clear preference towards the embellished or plain style of descriptions. Here are two examples that showcase the different preferences from users:

968  
969  
970  
971

**Preferred Description (AI Realtor, Description A):** Welcome to your charming retreat at 7401 S Blackstone Ave, nestled in the vibrant Grand Crossing neighborhood of Chicago! Priced attractively at \$179,900, this delightful single-family home skillfully blends affordability and comfort. With three spacious bedrooms and two well-appointed baths, it perfectly accommodates

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

family living. Highlights of this property include its exceptional flooring, ranking among the top 2% in the area, and the effective recent improvements that make it standout in property condition. Inside, enjoy the luxury of new carpeting and stylish white shaker cabinetry, promising an inviting and modern ambiance. A versatile layout offers flexible options for home customization, meeting the varied needs of your lifestyle, from working to exercising at home. Step outside to a fenced-in yard, ideal for both privacy and play, further increasing its appeal. Though without a garage, the home features machine hookups, ample storage potential in a sub-basement, and a premium location providing easy access to all the amenities and family-friendly spots Chicago's Grand Crossing has to offer. Enhanced by durable brick construction, this home represents a significant investment opportunity, combining all the elements you desire in a new haven. Your perfect blend of convenience, comfort, and charm awaits!

**Disfavored Description (Vanilla, Description B):** Welcome to 7401 S Blackstone Ave, an inviting single-family home nestled in Chicago's charming Grand Crossing neighborhood. Priced attractively at \$179,900, this 3-bedroom, 2-bathroom gem is perfect for those seeking comfort and convenience in an established community. Built in 1973, the residence boasts 933 square feet of living space and is beautifully updated with new carpeting throughout and stylish newer flooring. The kitchen shines with modern white shaker cabinetry, providing a fresh, contemporary feel. The home is well-equipped for practicality, featuring a dedicated laundry room with machine hookups for both gas and electric dryers. Adding to its appeal is a sub basement and a fenced-in yard, creating an ideal outdoor space for families or pet owners to enjoy. Conveniently located in Cook County, this home is serviced by Lake Michigan water and public sewer, and its brick construction ensures durability. With natural gas and forced air heating, you'll be cozy year-round. This delightful abode represents a fantastic opportunity for homeownership without the burden of HOA fees. Don't miss your chance to make this delightful Chicago residence your own!

**User Comment:** Description A is a bit more descriptive without going overboard, also talks about the neighborhood.

**Case Analysis:** Some users love more **descriptive** descriptions.

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

**Disfavored Description (🏡 AI Realtor, Description A):** Nestled in the heart of Chicago’s vibrant Bridgeport neighborhood, 3457 S Lituanica Ave offers unparalleled access and convenience, situated comfortably within Cook County. This spacious five-bedroom, two-bathroom single-family home is a standout choice, boasting top-tier features in location, accessibility, and outdoor living spaces. With its robust brick construction, this property provides a durable and inviting home environment, perfect for customization to suit your family’s evolving needs. Enjoy the luxury of a generous 6,500 sqft lot, among the best in its zipcode, offering a blank canvas for your dream garden or a secure playground for your child. The home’s interior shines with elegant hardwood flooring and practical features like in-unit laundry with sink. Practical comfort is ensured with space pac cooling and efficient natural gas heating, ensuring you feel at home year-round. Embrace Chicago living with easy access to nearby amenities, public transportation, and renowned neighborhood characteristics, all for an attractive price point of \$549,000—making it an excellent investment for future growth.

**Preferred Description (Vanilla, Description B):** Welcome to your future home at 3457 S Lituanica Ave, nestled in the heart of Chicago’s vibrant Bridgeport neighborhood. This charming single-family residence offers five spacious bedrooms and two full bathrooms, perfect for families seeking both comfort and style. Priced at an attractive \$549,000, this home sits on a generous 6,500 sqft lot, providing ample outdoor space for relaxation or entertaining. Crafted with enduring brick construction, the property boasts modern conveniences including a complete suite of appliances like a range, microwave, dishwasher, and more. The elegant hardwood flooring throughout adds a touch of sophistication, while the first-floor full bath caters to easy accessibility. Enjoy the convenience of in-unit laundry with a dedicated sink and stride out onto your private deck for a breath of fresh air. The two-car garage offers security and storage, supported by reliable utilities such as public sewer, natural gas heating, and Space Pac cooling. With easy access to Holden Elementary and local amenities, this home represents a delightful blend of classic charm and modern living in one of Cook County’s most desirable neighborhoods. Don’t miss the opportunity to make this house your home.

**User Comment:** Description B does a better job at listing the amenities.

**Case Analysis:** Some users love a **plain style** of description that listing all amenities.

These obervations suggest that there is no one-size-fits-all solution for writing style. Hence, future work could consider tailoring the description generation in the user’s preferred writing style to further improve the persuasiveness.

#### A.4 THE DIVERSITY OF WRITING STYLES ON DIFFERENT LISTINGS

🏡 AI Realtor shows diverse writing styles linguistically on listings based on their different features, which means it can tailor different real estate listings well.

In the following pair of examples, Low-end listings emphasize “Safety & Survival” (security, enclosure, reassurance), whereas high-end listings emphasize “Display & Views” (openness, visual richness, and mastery over the environment).

**Low-Price Representative (\$110,000).**

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

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

1132

1133

**Focus: Defense, Enclosure, Reassurance.** Words aimed at eliminating buyer insecurity regarding the environment.

**\$110,000, 750.0 sqft, 2 beds, 1 bath**

Welcome to your charming oasis at [address], nestled in the vibrant and culturally rich Hyde Park neighborhood of Chicago. This inviting 2-bedroom, 1-bathroom condo offers the perfect blend of comfort and convenience at an unbeatable price of \$110,000. Step inside to discover a sun-drenched living space adorned with hardwood flooring and an updated kitchen featuring modern appliances, including a wine refrigerator. The thoughtful design includes first-floor conveniences like a full bath and ample storage, with walk-in closets providing plenty of room for your essentials. Enjoy **tranquil moments** on the large back deck, ideal for relaxation or entertaining guests, set within a **gated courtyard that ensures privacy and security**. The property is **meticulously maintained**, boasting brick construction and a welcoming community atmosphere. Although it is **compact**, the space is optimized for comfortable living **without unnecessary upkeep**, perfect for those valuing efficiency. With **proactive security measures**, a **strong sense of community**, and only minutes away from necessities, this condo perfectly encapsulates the ideal home for those **prioritizing safety** and cultural alignment in a vibrant neighborhood.

### High-Price Representative (\$1,875,000).

**Focus: Aggression, Openness, Visuals.** Words aimed at showing off transparency and mastery over the environment.

**\$1,875,000, N/A sqft, 4 beds, 4 baths**

Discover **unparalleled elegance** and style at [address], a single-family haven nestled in the vibrant Bucktown neighborhood of Chicago. This **exquisite home**, priced at \$1,875,000, offers four bedrooms and four bathrooms, perfect for families seeking **ample space and luxury**. Its standout features include a **private corner lot** and a **spacious side yard** designed for **ultimate outdoor enjoyment**, complemented by **gourmet enhancements** like a custom kitchen and a chic beverage center. New Pella windows and **cascading expanses of glass** invite an **abundance of natural light**, creating a bright and airy atmosphere across a versatile loft area ideal for work-from-home needs. With **sophisticated enhancements** such as vaulted ceilings, multiple fireplaces, and a gas fire table, this residence exudes comfort and warmth year-round. The **meticulously crafted design** places this property among the **top tier in architectural style** and elegance within the neighborhood and beyond. Enjoy **seamless access** to essential amenities and natural beauty, with a spacious parking capacity for four cars. **Embrace this rare opportunity** to own a piece of **refined luxury** in an urban yet serene setting.

## B THE DESIGN OF SURVEY AND USER INTERFACES

### B.1 SURVEY SCREENING INTERFACE

The first stage of the survey is designed to ensure the human subject has sufficient experience in the home search process in order to analyze the features from a marketing description. We present description of an example listing and design quiz-like questions to verify whether the participant is able to make all correct responses. We showcases the web user interfaces in Figure 6.

### B.2 PREFERENCE ELICITATION INTERFACE

In the second stage of the survey, we design an interface to mimic the environment of online platforms that the model can observe the buyer's general profile and behaviors (e.g., recently browsed or liked listing) to some degree. In our case of real estate listing, we ask the buyer to provide their preferences in a 1-5 scale on five general categories (price, location, home features & amenities, house size, investment value) and set a filter on the price range and number of bedrooms in the house they are looking for. This information allows us to select generally relevant listings to mitigate the anchoring effect that the marketing content can play little role to influence the buyer in the

Figure 6: Survey Screening Interface

evaluation phase. Next, we choose 5 relevant listings and ask the buyer to rate them on a 1-5 scale and provide their reasoning. This process ensures that we can collect a reasonable amount of each buyer’s preference information for the personalized persuasive content generation in the evaluation phase. Finally, we employ LLM to narrow the features that are likely preferred by the participants and ask for their ratings of importance on a 1-5 scale. We showcases the web user interfaces in Figure 7.

Figure 7: Preference Elicitation Interface

### B.3 HUMAN EVALUATION INTERFACE

In the last stage of the survey, it is to gather the human feedback on the persuasiveness of different models. Many previous works study persuasion by asking human how much does their opinion changes before and after reading an argument. In our task, human subjects often do not have any prior knowledge about item and this evaluation procedure would induce bias. Instead, we implement two alternative evaluation schemes in our interface: one is the A/B test where the buyer is presented

1188 with a single listing along with two descriptions generated by two distinct models and then asked  
 1189 to report which description makes them more interested in the listing; the other is the interleaved  
 1190 test where a set of listings each with a single description generated by some model and the buyer  
 1191 is asked to select the listings that they are interested in based on their descriptions. Each time after  
 1192 a participant’s choice of the preferred description, we ask participant to rate on a scale of 1-5 that  
 1193 one description is prefer over another and incentivized them to provide a detailed rationale of their  
 1194 responses. To illustrate this process, we present the web interface design in Figure 8.

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

#### B.4 FEATURE ANNOTATION INTERFACE

1219

1220

1221

To ease the task of feature annotation, we also develop a user-friendly web interface. Its design is shown in Figure 9.

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

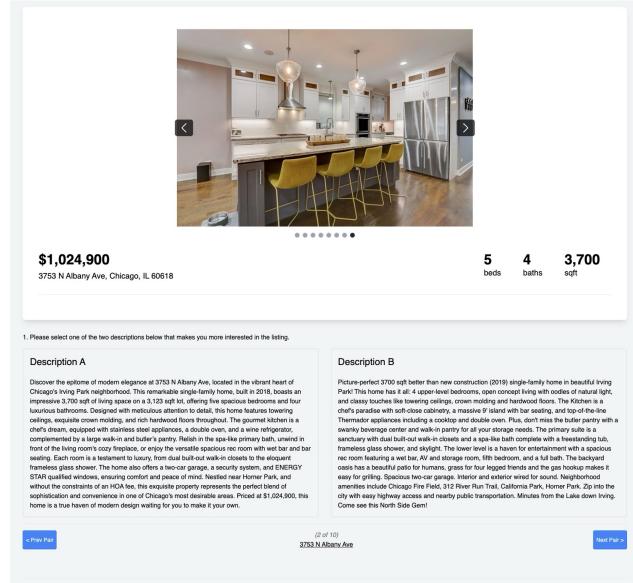


Figure 8: Human Evaluation Interface

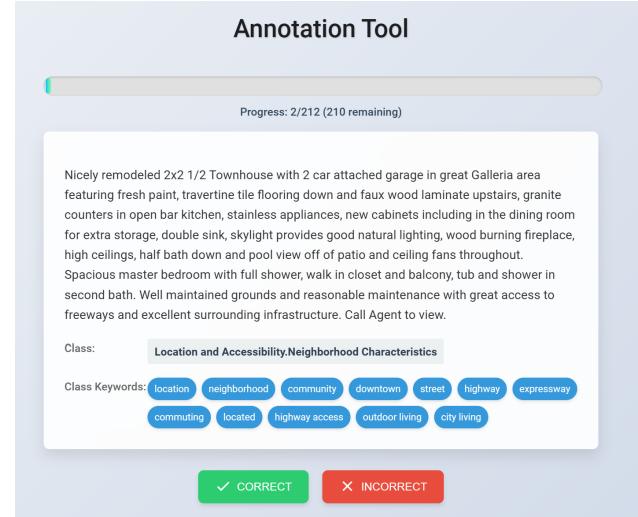


Figure 9: Annotation Interface

1242 **C IMPLEMENTATION DETAILS**  
 1243

1244 In this section, we provide a full description of the implementation detail of  AI Realtor.  
 1245

1246 **C.1 GROUNDING MODULE: PREDICTING MARKETABLE FEATURES**  
 1247

1249 Our model assumes the existence of attribute-feature mappings in different marketing problems,  
 1250 with which a seller can use to influence the buyer’s beliefs and behaviors. However, a key chal-  
 1251 lenge lies in determining how to accurately obtain such mappings. Specifically, we must identify  
 1252 which *signaling features* to include and under what conditions it is natural to market a product as  
 1253 possessing a particular feature. Traditionally, acquiring this knowledge from human experts is both  
 1254 labor-intensive and costly. Instead, we take a learning approach to uncover the mapping from our ex-  
 1255 periment dataset. While the raw dataset contains no annotation of any signaling feature, we employ  
 1256 LLMs to construct a high-quality feature schema and label the dataset accordingly in preparation  
 1257 for learning the attribute-feature mapping. This approach notably presents a novel unsupervised  
 1258 learning paradigm, harnessing the broad knowledge of LLMs to distill expert-level insights from  
 1259 unlabeled data with minimal human supervision.

1260 **Inductive Construction of Feature Schema** Our dataset only contains the raw attributes of each  
 1261 product. In order to learn a high-quality attribute-feature mapping, the first task is to obtain a good  
 1262 representation of feature schema  $S$ . On the one hand, if we miss some useful signaling features,  
 1263 it could significantly hinder the performance of subsequent marketing task. On the other hand,  
 1264 there are so many possible token that can serve as the signaling features in the natural language  
 1265 space, and many of these tokens might have duplicate or similar meaning. If there is no structured  
 1266 representation of the features, the resulting label classes could be too sparse to learn. Indeed, we  
 1267 discover that the feature schema obtained by directly prompting an LLM includes many similar  
 1268 features while miss some important ones. Based on this observation, we turn to a more sophisticated  
 1269 prompting strategy to inductively improve the quality and representation of the feature schema (see  
 1270 a high-level sketch of the construction pipeline in Figure 2).

1271 First, we construct a basis of feature schema, represented as a list of tokens used in the human-written  
 1272 marketing description to describe some house features. We begin with *Mixtral-8x7B-Instruct-*  
 1273 *v0.1* (Jiang et al., 2024) to extract keywords or phrases  $\{k_1, k_2, \dots\} = \text{LLM}_{\text{gen}}([\mathcal{I}_{\text{Keyword}}; D_{\text{human}}])$   
 1274 that summarize each human-written description  $D_{\text{human}}$  under a keyword-extraction prompt  $\mathcal{I}_{\text{Keyword}}$   
 1275 (Appendix F.1). We observed that, in some cases, the model output could not be directly parsed  
 1276 into a clean list of keywords, or it contained excessive quantifiers and modifiers. To address this,  
 1277 we re-prompted the model using  $\mathcal{I}_{\text{Norm}}$  (Appendix F.2) to normalize each keyword. Through this  
 1278 process, we initially extracted 112688 keywords—too many to handle effectively. We then ap-  
 1279 plied additional normalization steps, including lowercasing, lemmatization, and synset merging via  
 1280 NLTK (Bird et al., 2009). We also filtered the keywords, retaining only those that appeared in at  
 1281 least 50 descriptions. This reduced the final set to 1114 keywords as our *induction base*.

1282 Next, we organize the feature-related keywords into a structured feature schema. Since many key-  
 1283 words are related to each others and hard to distinguish, we use a hierarchical representation of  
 1284 feature schema to better capture the relations between different feature classes and to ease the sub-  
 1285 sequent labeling task. To achieve this goal, we prompted *Claude-3.5-Sonnet* (Anthropic, 2024) with  
 1286 a 100-keyword batch to iteratively generate a hierarchical schema that covers the majority of the  
 1287 keywords (an example run can be found in Appendix F.3). We temporarily switched to *Claude-3.5-*  
 1288 *Sonnet* because we found it particularly difficult for open-source models, even the state-of-the-art  
 1289 *GPT-4o* (OpenAI, 2024a), to induce such a schema without grouping most keywords into overly  
 1290 broad categories like “others” or “misc”, resulting in a shallow and uninformative schema. In  
 1291 contrast, when fed keywords in small batches, *Claude-3.5-Sonnet* followed our instructions more  
 1292 faithfully, organizing the keywords into a carefully structured hierarchy. Every leaf node in the  
 1293 schema was associated with a set of relevant keywords. From this process, we obtain a relatively  
 1294 well-structured and comprehensive feature schema.

1295 Finally, to evaluate the quality of the generated feature schema, monitor potential hallucination  
 1296 issues, and further refine the schema, we asked three human participants to conduct manual review.  
 1297 We prompt *Mixtral-8x7B-Instruct-v0.1* to determine whether a feature from the schema presents in  
 1298 each human-written description, and each participant is asked to independently verify this result (see

1296 our annotation interface in Appendix B.4). Based on the participants’ feedback on 636 samples, we  
 1297 found that features labeled by LLMs are mostly agreed across all human annotators, except for some  
 1298 ambiguous or subjective features (e.g., the aesthetic features of a house), where the agreement rates  
 1299 (around 60%) between models and human are about as good as that among human annotators. We  
 1300 refine the schema for two more iterations, where we prompt LLMs to merge some similar features  
 1301 and reduce the ambiguity of some features with more precise example keywords. We list our final  
 1302 feature schema in Appendix D.2 and it is used in the subsequent stages of our pipeline.

1303  
 1304 **Learning the Feature-Attribute Mapping** With the feature schema, we guide the LLM to an-  
 1305 notate for each product with attributes  $\mathbf{x}$  whether each feature  $s_i$  is described in the human-written  
 1306 marketing text (see the prompt in Appendix F.4). We perform a few additional pre-processing steps  
 1307 to this correspondence data to supervise the learning of the feature-attribute mapping.

1308 First, we found that some human-written marketing descriptions are of relatively low quality and  
 1309 these data points can negatively impact the learnt feature-attribute mapping. Hence, we only select  
 1310 marketing descriptions of products that are relatively popular, according to a simple heuristic ratio  
 1311 between the number of likes and views received by a listing recorded on the marketing platform.  
 1312 We expect the quality of feature-attribute mapping uncovered from this filtered set of human-written  
 1313 descriptions would be higher than average.

1314 Next, we normalize the attributes of each listing  $\mathbf{x}$  and embed existing knowledge of these attributes  
 1315 into their representation. Since the raw attributes of each listing  $\mathbf{x}$  have different value types (cat-  
 1316 egorical, integer, float, etc.), we convert each attribute  $x_i$  into a natural language statement using  
 1317 the template, “The attribute *attribute\_name* is *attribute\_value*.”, and then use an embedding model,  
 1318 *SFR-Embedding-Mistral* (Meng et al., 2024), to convert each natural language statement into a fixed-  
 1319 dimensional vector  $e_i = \text{LLM}_{\text{embed}}(x_i) \in \mathcal{R}^d$ . We also perform some standardized normalization  
 1320 techniques such as removing irrelevant attributes and dropping attributes with missing values. Fi-  
 1321 nally, we use a simple multi-layer perceptron (MLP) to learn the attribute-feature mapping as,

$$\pi(s_i \mid \mathbf{x}) = \sigma(O_i^T \text{ReLU}(W\bar{e}(\mathbf{x}))),$$

1322 where  $\bar{e}(X)$  is the mean-pooled attribute embedding, and  $O_i \in \mathcal{R}^{d/2}$ ,  $W \in \mathcal{R}^{d \times d/2}$  are the model’s  
 1323 weights. The function  $\sigma$  represents the sigmoid activation function. Here, we assume conditional in-  
 1324 dependence between highlights given the raw features  $X$ . We use the standard logistic loss function  
 1325 to training the neural network. We apply a random train-test split of 4 : 1 ratio in our dataset and  
 1326 achieve testing accuracy 69.39% and F1 score 67.43%. We find the accuracy to be reasonably high,  
 1327 given the stochastic nature of signaling process. That is, the features deterministically predicted  
 1328 based on our mapping cannot exactly match with the features used in the human written description  
 1329 with some degree of randomness — just as the accuracy of predicting a fair coin toss is at most 50%.

1330 The typical implementation of a signaling scheme is to follow the attribute-feature mapping  $\pi$  to  
 1331 randomly draw a signal  $S_j$  with probability  $s_j(\mathbf{x})$ . This is necessary in theory to maintain the partial  
 1332 information carried by each signal. However, we implement a deterministic feature selection strat-  
 1333 egy to only use feature  $S_j$  with probability above some threshold  $\alpha$ . This is because our generated  
 1334 marketing content only accounts for a tiny portion of the corpus so that it should have almost no  
 1335 influence on people’s perception of a feature (e.g., the partial knowledge inferred upon observing  
 1336 each feature). This also ensures that the product would have the feature with high probability, as  
 1337 our objective prioritizes the rigorousness of our marketing content. As a simple heuristics in our  
 1338 implementation, we set the threshold  $\alpha = 1/2$  and we will refer to this set of features as,

$$\text{Marketable Features: } \mathcal{S}_1(\mathbf{x}) = \{S_j : s_j(\mathbf{x}) \geq \alpha\}. \quad (3)$$

## 1342 C.2 PERSONALIZATION MODULE: ALIGNING WITH PREFERENCES

1343 This stage seeks to steer the persuasive language generation toward the buyer’s preference, which  
 1344 is another crucial objective of grounded persuasion. In particular, with the advent of LLM, there  
 1345 is an unprecedented opportunity for our data-driven approach could achieve much higher degree of  
 1346 personalization with significantly lower cost than the conventional marketing designed for a larger  
 1347 population. Our solution has two parts: the first part is to properly elicit the useful information about  
 1348 a user’s preference and structure it in a good representation; the second part is to select a subset of  
 1349 features based on the user preference in order to maximize the influence to the user’s belief.

1350     **Structured Preference Representation** As mentioned previously, our evaluation environment is  
 1351     built to have an information elicitation process from each buyer. However, such information cannot  
 1352     directly describe the user’s preference. So, we ask the LLM to act like a human realtor to determine  
 1353     the features that the users might be interested in based on their initial selection. To do this, we  
 1354     prompt the language model to convert the user preference into information structured according to  
 1355     the feature schema. We then ask the user to give a rating  $r_j$  on a scale of 1-5 on how important  
 1356     each feature  $S_j$  is. We also elicit the user’s rationale behind this rating to nudge users to give more  
 1357     thoughts on their selection and thereby improve the credibility of their rating responses. While our  
 1358     implementation mostly relies on user surveys and the information processing power of LLMs, this  
 1359     design is a reasonable simulation of digital marketing in real-world applications, where  $r_j$  can be  
 1360     learned through the standard industrial techniques of cookie analysis.

1361  
 1362     **Personalized Feature Selection** While the marketable features in Equation (3) are predicted at a  
 1363     population level, it is also useful to select features that are tailored to the user’s special interests.  
 1364     However, because real-world marketing descriptions are not optimized for individual users, we can-  
 1365     not simply rely on a data-driven machine learning approach for personalization. Instead, we leverage  
 1366     the innate capability of LLMs to understand and analyze human preference. In our implementation,  
 1367     we select a set of features that are marketable and preferred by the buyer and let the LLMs to decide  
 1368     which personalized features to emphasize on in the marketing content. Our heuristic method for  
 1369     personalized feature selection is to adjust the population-level feature scores  $s(\mathbf{x})$  with the user’s  
 1370     rating over each feature  $\mathbf{r}$  as follows,

$$1371 \quad \text{Personalized Features: } \mathcal{S}_2(\mathbf{x}) = \{s_j | s_j(\mathbf{x}) + c(r_j - r_0) \geq \alpha\}, \quad (4)$$

1372  
 1373     where the constant  $c$  reflects the intensity of personal preference,  $r_0$  is the basis rating of each  
 1374     attribute. In our human-subject experiment, we choose  $c = 0.01$ ,  $r_0 = 2$  and set the threshold value  
 1375      $\alpha$  such as to select features of the top 10 highest scores. We list these features in the prompt to  
 1376     generate persuasive marketing description (see a full specification in Appendix F.5).

### 1378     C.3 MARKETING MODULE: CAPTURING SURPRISING VIA RAG

1380     The last stage is designed to better ground the persuasive language generation on factual evidences,  
 1381     problem contexts and localized information in automated marketing. There are many ways to im-  
 1382     prove the grounding for different settings of automated marketing. As a case study, we choose to  
 1383     focus on the surprising effect, a common marketing strategy studied by many work (Lindgreen &  
 1384     Vanhamme, 2005; Ludden et al., 2008; Ely et al., 2015), under which the buyers would derive enter-  
 1385     tainment utility and have a deeper impression. In our setting of real estate marketing, we consider  
 1386     the type of features that are relatively rare in its surrounding area. That is, we say a marketable  
 1387     feature  $S_j$  is *surprising* if it is among the top  $\beta$ -quantile of the distribution of  $S_j$  values under the  
 1388     prior distribution  $s_j(\mu)$ , or formally,

$$1389 \quad \text{Surprising Features: } \mathcal{S}_3(\mathbf{x}) = \{S_j \subset \mathcal{S}_1 : \\ 1390 \quad s_j(\mathbf{x}) \text{ is within } \beta\text{-quantile of distribution } s_j(\mu)\}. \quad (5)$$

1391  
 1392     In our implementation, we determine a set of features for each listing that have its comparative ad-  
 1393     vantage among different groups of similar listings. We consider two kinds of retrieval criteria: (1)  
 1394     select all listings within the proximal location at different levels of granularity (e.g., neighbourhood,  
 1395     zipcode or city); (2) select the 10 listings with the most similar features via an information retrieval  
 1396     system (implemented by the ElasticSearch framework<sup>6</sup>) — the search engine implementation de-  
 1397     tails can be found in Appendix F.8. For each group of similar listings, we determine an empirical  
 1398     distribution function on each attribute score  $\tilde{F}_i$ . We then set  $1 - \tilde{F}_i(p_i)$  as the percentile ranking of  
 1399     the listing’s attribute  $i$  among this group. We then select all attributes that are among the top 30%  
 1400     percentile ranking for some group and provide the information in the prompt to generate persuasive  
 1401     marketing language (see a full specification in Appendix F.6). This gives the LLMs localized feature  
 1402     information at different granularity level.

1403  
 6 $\text{https://www.elastic.co/elasticsearch}$

1404

## D DATA CURATION

1405

1406

## D.1 DATASET RAW ATTRIBUTE SCHEMA

1407

1408

To ensure both quality and fidelity of our evaluation, we collect the real data of real estate listings on the market. The dataset for this experiment was sourced primarily from Zillow and includes around 50000 listings collected in the month of April in 2024. We follow the Zillow terms of services<sup>7</sup> to avoid any commercial use of their data. Each of these listings is from one of the top 30 most populous cities in the United States as described by the U.S. Census Bureau. Listings that were not residential in nature or were missing crucial data to this experiment were excluded from this dataset. This dataset is composed of 95 columns, with features ranging from number of bedrooms, price, views, and more (see Table 1). These many features associated with each listing provide us sufficient space to develop and test improved models for grounded persuasion.

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

## D.2 FINAL FEATURE SCHEMA

1448

Here is the condensed version of the final feature schema to save pages:

1449

1450

1451

## Interior Features:

1452

## Rooms:

1453

```
[bath,bathroom,bedroom,kitchen,living room,secondary
→ bedrooms,patio,backyard,closet,room,living,dining
→ room,pantry,space,office,laundry room,dining,living
→ space,living area,primary suite,master suite,family
→ room,cellar,foyer,game room,great room,den,master
```

1454

1455

1456

1457

<sup>7</sup><https://www.zillow.com/z/corp/terms/>

```

1458      ↳ bedroom,utility room,sunroom,bedroom suite,living
1459      ↳ areas,primary bedroom,office space,kitchenette,owner
1460      ↳ 's suite,playroom,storage room,living rooms,ensuite,
1461      ↳ wet bar,loft area,sitting room,mud room,exercise
1462      ↳ room,clothes closets,walk-in closet,mudroom,
1463      ↳ conference room]
1464  Flooring:
1465      [flooring,stories,carpeting,hardwood floors,tile,tile
1466      ↳ floors,hardwood flooring,wood flooring,hardwood
1467      ↳ floors]
1468  Furniture:
1469      [desk,table,chair,bed,dressers,cupboards,sofa,bench,
1470      ↳ seating]
1471  Additional Spaces and Versatility:
1472      [bonus room,flex space,flex room,den]
1473  Kitchen Features:
1474      [countertop,granite countertops,marble countertops,island,
1475      ↳ cabinetry,kitchen island,kitchen cabinets,waterfall,
1476      ↳ dining space,cooktop]
1477  Architectural Elements:
1478      [roof,window,floor plan,cabinet,molding,staircase,brick,
1479      ↳ paneling,siding,beam,ceiling fans,stair,chandelier,
1480      ↳ finishing trim,baseboard,trim]
1481  Bathroom Features:
1482      [shower,vanity,powder room,jacuzzi,ensuite,half bath,water
1483      ↳ closet,mirror,faucet]
1484  Storage:
1485      [storage,closet space,cabinet space,shelving,storage space
1486      ↳ ,mudroom,drawer,bookshelf,storage unit,clothes
1487      ↳ storage,bike storage]
1488  Comfort and Ambiance:
1489      Lighting:
1490      [lighting,natural light,light fixtures,skylight,
1491      ↳ lighting fixtures]
1492      Temperature Control:
1493      [fireplace,hvac,fan,ac,a/c,central air conditioning]
1494  Exterior Features:
1495      Outdoor Spaces:
1496      [patio,backyard,yard,pool,spa,balcony,porch,deck,roof deck
1497      ↳ ,outdoor space,rv parking,outdoor spaces,outdoor
1498      ↳ living space,fenced yard,pavers,garden,outdoor
1499      ↳ living,backyard oasis,pergola,gazebo,cabana,
1500      ↳ landscaping,shade,lawn,fountain,sod,outdoor bench]
1501      Outdoor Activities:
1502      [gardening,outdoor cooking,barbecue,bbq]
1503  Location and Accessibility:
1504      Neighborhood Characteristics:
1505      [location,neighborhood,community,downtown,street,highway,
1506      ↳ expressway,commuting,located,highway access,outdoor
1507      ↳ living,city living]
1508      Nearby Amenities:
1509      [shopping,restaurant,park,school,grocery,cafe,hospital,
1510      ↳ food,stadium,museum,boutique,shopping centers,
1511      ↳ station,elementary,bus,trader joe's,golf,brewery,
1512      ↳ elementary school,school district,recreation
1513      ↳ facility]
1514  Cities/Regions:

```

```

1512 [Austin, Denver, Charlotte, Houston, Dallas, San Antonio,
1513   ↪ Nashville, Phoenix, Los Angeles, LA, Manhattan, Detroit,
1514   ↪ Philadelphia, Portland]
1515 Access and Transportation:
1516   [access to amenities, proximity to schools, proximity to
1517     ↪ restaurants, proximity to shops, access to shopping,
1518     ↪ bus stop, walking distance, proximity to shopping,
1519     ↪ freeway access, public transit nearby, public
1520     ↪ transportation, road]
1521 Walkability and Bikeability:
1522   [walkability, bike score, walk score]
1523 Housing Types:
1524   [studio, cottage, ranch, duplex, townhome, brownstone, row home,
1525     ↪ bungalow]
1526 Building Features:
1527   Structure:
1528     [condo, loft, unit, townhouse, estate, square feet, duplex,
1529       ↪ garage, carport, story, penthouse, sf, triplex, colonial]
1530   Parking:
1531     [garage, parking, parking space, parking spaces, garage door,
1532       ↪ parking spot]
1533 Appliances:
1534   [appliance, refrigerator, dishwasher, washer/dryer, range, fridge,
1535     ↪ microwave, washer, ac unit, dryer, hood, laundry facilities,
1536     ↪ washer and dryer, oven, garbage disposal, wolf appliances,
1537     ↪ thermador appliances]
1538 Amenities:
1539   [community center, community pool, spa, firepit, fire pit,
1540     ↪ outbuilding, tennis courts, club house, rooftop, rooftop
1541     ↪ deck, rooftop terrace, dog park, lounge, elevator, recreation
1542     ↪ room, gym, fitness center, clubhouse, swimming pool, pool,
1543     ↪ spa, sauna, hot tub, putting green, tennis courts, basketball
1544     ↪ , pickleball, tennis court, golf, management, booking,
1545     ↪ concierge, trash, maintenance, doorman, superintendent,
1546     ↪ nightlife, brewery]
1547 Utilities and Systems:
1548   [plumbing, water heater, heater, hot water heater, water, water
1549     ↪ filtration system, gas, sprinkler system, hvac, ac, a/c,
1550     ↪ wiring, solar panels, solar, electrical panel, electricity,
1551     ↪ generator, security, security system, camera, internet, wifi,
1552     ↪ cable, phone, satellite, fiber, internet access, satellite TV
1553     ↪ , internet service, irrigation system, ac unit, hvac unit,
1554     ↪ central air conditioning]
1555 Design and Style:
1556   Interior Design:
1557     [paint, style, home style, architecture, woodwork, ensemble,
1558       ↪ accent, open floor plan, drawing]
1559   Aesthetics:
1560     [elegance, sophistication]
1561   Architectural Styles:
1562     [tudor, colonial, craftsman, farmhouse]
1563 Smart Home Features:
1564   [smart home technology, surround sound, home technology, camera]
1565 Lifestyle Features:
1566   Work from Home:
1567     [workspace, home office]
1568 Entertainment:
1569   [entertaining space, party, entertainment options, wet bar,
1570     ↪ entertainment]

```

```

1566 Sustainability Features:
1567   [solar system, sustainability, solar, heated floors, solar panels,
1568     ↪ tankless water heater]
1569 Real Estate Financial and Legal Aspects:
1570   [condo fee, hoa fee, hoa fees, equity, hoa dues, condo fees, cdd
1571     ↪ fees, occupied, rental potential, income potential,
1572     ↪ appreciation, airbnb, investment opportunity, investor
1573     ↪ opportunity, warranty, pricing, rental income, income,
1574     ↪ financing, utility, sale, closing, furnished, slip, tax, flip
1575     ↪ tax, abatement, zoning, hoa, rental cap, option]
1576 Water Features:
1577   [soaking tub, softener]
1578 Views and Scenery:
1579   [mountain views, lake views, ocean views, sunset, city views,
1580     ↪ skyline, skyline views]
1581 Property Characteristics:
1582   Specialty Rooms:
1583     [wine cellar, media room, suite]
1584   Distinctive Interior Elements:
1585     [exposed brick, high ceilings]
1586   Exterior Appearance:
1587     [curb appeal, facade, exterior paint]
1588   Atmosphere:
1589     [oasis, retreat, sanctuary, flow]
1590   Environment:
1591     [surroundings]
1592   Property Metrics:
1593     [lot, corner lot, sqft, br, walk score, foot, inch]
1594   Property Condition:
1595     Improvements:
1596       [improvement, tlc, fixer, flooded]
1597     Age and Status:
1598       [new, renovated, remodeled, renovated, rehabbed, home age,
1599         ↪ upgrade, update, built, finish, updated, move,
1600         ↪ readiness, move-in ready, maintained]
1601
1602 Real Estate Industry:
1603   [builder, agent]
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

```

### D.3 DIVERSITY OF THE REAL ESTATE MARKET IN CHICAGO

In this section, we analyze the diversity of the real estate market in Chicago compared to other major US cities. We use two quantitative signals: (1) the diversity of home types measured by entropy, and (2) the dispersion of prices measured by percentile ratios (p90/p10 and p75/p25). Higher values in either metric indicate a more heterogeneous market. The home type entropy for each city is summarized in Table 2, and the cross-city price dispersion is reported in Table 3.

Overall, Chicago emerges as the most diverse market among the major cities examined. As shown in Table 2, it has the highest home type entropy, indicating a well-balanced mix of condos, single-family homes, multi-family units, and townhouses. Meanwhile, Table 3 shows that Chicago also exhibits the strongest price dispersion, reflecting a wide range of housing options across different price tiers. Together, these signals highlight Chicago as a particularly heterogeneous and versatile real estate market.

### E HALLUCINATION EXPERIMENT DETAILS

In this section, we introduce implementation details for hallucination verification experiments. We will introduce both automatic evaluation and human evaluation.

1620 Table 2: Home type entropy across major US cities. Higher entropy indicates a more balanced  
 1621 home-type distribution. Chicago exhibits the highest diversity.

1623	City	Home Type Entropy
1624	<b>Chicago, IL</b>	<b>0.8613</b>
1625	Seattle, WA	0.8415
1626	San Jose, CA	0.8399
1627	Los Angeles, CA	0.8018
1628	San Francisco, CA	0.7851
1629	Washington, DC	0.7796
1630	Portland, OR	0.7434
1631	Denver, CO	0.7064
1632	San Diego, CA	0.6849
1633	Philadelphia, PA	0.6375

1634 Table 3: Price dispersion across cities. Higher percentile ratios indicate larger heterogeneity in  
 1635 listing prices. Chicago shows the strongest price dispersion.

1636	City	Price p90/p10	Price p75/p25
1637	<b>Chicago, IL</b>	<b>10.09</b>	<b>3.36</b>
1638	Seattle, WA	5.44	2.07
1639	San Jose, CA	4.81	2.32
1640	Los Angeles, CA	5.48	2.37
1641	San Francisco, CA	5.30	2.28
1642	Washington, DC	7.03	2.50
1643	Portland, OR	5.09	2.24
1644	Denver, CO	5.85	2.47
1645	San Diego, CA	5.44	2.32
1646	Philadelphia, PA	5.97	2.38

### 1649 E.1 AUTOMATIC EVALUATION

1650 We adopt fine-grained fact-checking based on GPT-4o for automatic evaluation, similar to the  
 1651 pipeline introduced in FActScore(Min et al., 2023). Specifically, we select *price*, *living area* (in  
 1652 sqft), *#bedrooms* and *#bathroom* as  $X_{\text{hard}}$  and *home insights*, *address* as  $X_{\text{soft}}$  according to a prior  
 1653 survey of user preference.

1654 We use structured output API<sup>8</sup> on OpenAI to setup  $\text{eval}_{\text{soft}}(L, x)$  and  $\text{eval}_{\text{hard}}(L, x)$ . This means  
 1655 in both cases, we need to first define the structured output class specification and then prompt the  
 1656 model with it.

1657 For  $\text{Faithful}_{\text{hard}}$ , our structured output class specification is:

```
1659     class MainInfo(BaseModel):
1660         price_mentioned: bool
1661         price: float
1662         living_area_mentioned: bool
1663         living_area: str
1664         bedrooms_mentioned: bool
1665         bedrooms: float
1666         bathrooms_mentioned: bool
1667         bathrooms: float
1668         address_mentioned: bool
1669         address: str
```

1670 and our prompt for  $\text{eval}_{\text{hard}}(L, x)$  is:

1671 <sup>8</sup><https://platform.openai.com/docs/guides/structured-outputs/introduction>

```

1674 messages=[  

1675     {"role": "system", "content": "Extract Real Estate Information  

1676         ↪ . Find the price (e.g, 290000.0), living area (e.g.,  

1677         ↪ '990.0 sqft'), bedrooms (e.g., 2) and bathrooms (e.g.,  

1678         ↪ 3) from the description. Not all information may be  

1679         ↪ present, so you also have to determine whether each  

1680         ↪ field is mentioned or not."},  

1681     {"role": "user", "content": {description}}  

1682 ]  


```

1683 We then compare the extracted information with  $\text{supp}(L, X_{\text{hard}})$  to compute  $\text{Faithful}_{\text{hard}}$ . If certain  
1684 attributes are mentioned (i.e.,  $\text{xx\_mentioned}=\text{True}$ ) and the corresponding extracted values matched  
1685 the listing info  $\text{supp}(L, X_{\text{hard}})$ , then we will give one score, otherwise zero.  
1686

1687 For  $\text{Faithful}_{\text{soft}}$ , we will compute it in two stages. First, we will conduct attribute extraction as in  
1688  $\text{Faithful}_{\text{hard}}$ , but with a different set of attributes  $X_{\text{soft}}$ . Our structured output class specification is:

```

1689     class MainInfo(BaseModel):  

1690         home_insights_mentioned: bool  

1691         home_insights: list[str]  

1692         address_mentioned: bool  

1693         address: str  

1694

```

1695 and our prompt is:

```

1696 example_home_insights =["Large island", "Oversized bathroom", "  

1697     ↪ Open floor plan", "Lake views", "Orange l lines", "Newer  

1698     ↪ stainless steel appliances", "Gorgeous hardwood floors", "  

1699     ↪ Tons of cabinet space", "In-unit washer and dryer", "Skyline  

1700     ↪ view", "Private balcony", "Beautiful city"]  

1701 example_addr = "1255 S State St UNIT 703 Chicago IL 60601"  

1702 messages=[  

1703     {"role": "system", "content": "Extract Real Estate Information  

1704         ↪ . Find the home insights (e.g., {example_home_insights})  

1705         ↪ , and address (e.g., {example_addr}) from the  

1706         ↪ description. Not all information may be present, so you  

1707         ↪ also have to determine whether each field is mentioned  

1708         ↪ or not."},  

1709     {"role": "user", "content": {description}}  

1710 ]

```

1710 In the second stage, we will use JSON mode API<sup>9</sup> to check whether the extracted attributes match  
1711  $\text{supp}(L, X_{\text{soft}})$ . Our matching prompt is:

1713 Given the following information:

```

1715 1. Description: {description}  

1716 2. True value for {attribute_name}: {json.dumps(true_value)}  

1717 3. Extracted value for {attribute_name}: {json.dumps(  

1718     ↪ extracted_value)}

```

1719 Please analyze how well the extracted value matches the true value  
1720 ↪ , considering the context provided in the description.  
1721

1722 For 'home\_insights', consider it a good match if a significant  
1723 ↪ subset of the true insights is correctly identified.  
1724 For 'address', consider it a good match if at least a subset (e.g  
1725 ↪ .., city/state) is correctly identified, given it was  
1726 ↪ mentioned in the description.

---

<sup>9</sup><https://platform.openai.com/docs/guides/structured-outputs/json-mode>

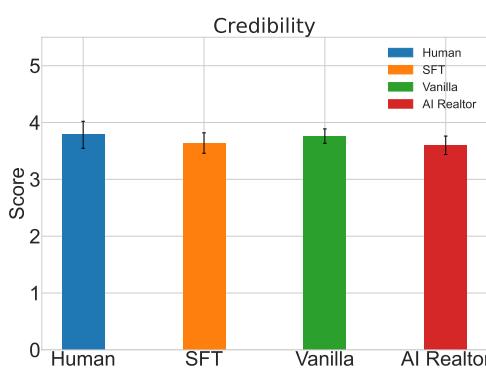


Figure 10: Credibility Scores for Hallucination Checks.

1728  
1729  
1730  
1731  
1732  
1733  
1734  
1735  
1736  
1737  
1738  
1739  
1740  
1741  
1742  
1743  
1744  
1745 Provide a score between 0 and 10, where:  
1746 0 = Completely incorrect or irrelevant  
1747 5 = Partially correct or relevant  
1748 10 = Perfect match

1749 Respond with a JSON object in the following format:  
1750 {  
1751 "score": int  
1752 }  
1753

1754 Where 'score' is an integer between 0 and 10.  
1755

1756 Finally we sum up all scores to compute  $\text{Faithful}_{\text{soft}}$ .  
1757

## E.2 HUMAN EVALUATION

1760 We recruit human annotators to replicate GPT-4o’s hallucination checks and assess the reliability  
1761 of its automatic evaluations. To ensure consistency with the LLM judge, we define factuality iden-  
1762 tically for human raters: verifying that claims made in the description are strictly grounded in the  
1763 provided attribute set  $X$ . In addition to the two factual attributes evaluated by GPT-4o— $X_{\text{hard}}$  and  
1764  $X_{\text{soft}}$ —we include an additional stylistic check: **credibility**, which captures users’ emotional judg-  
1765 ment of whether the persuasive description feels trustworthy.

1766 Given an attribute set  $X$  and a description  $L$ , either sampled from model- or human-generated out-  
1767 puts, we ask users to (1) rate the credibility of  $L$  on a 1–5 scale (Figure 11a), (2) evaluate how  
1768 well each hard attribute  $x_{\text{hard}} \in X_{\text{hard}}$  is reflected in  $L$ , if it is mentioned ( $X_{\text{hard}} \in \text{supp}(L, X_{\text{hard}})$ )  
1769 (Figure 11b), and (3) assess how well each soft attribute  $x_{\text{soft}} \in X_{\text{soft}}$  is reflected, if it is mentioned  
1770 ( $x_{\text{soft}} \in \text{supp}(L, X_{\text{soft}})$ ) (Figure 11c). The instruction files provided to human annotators will be  
1771 submitted in a separate supplementary file.

1772 As shown in Figure 5, and consistent with findings in § 5.3, AI Realtor achieves the highest  
1773 faithfulness on  $X_{\text{hard}}$ , while human-written descriptions score lowest in credibility. For evaluations  
1774 on  $X_{\text{soft}}$  (Figure 5) and credibility (Figure 10), which requires more subjective judgment, the per-  
1775 formance of AI Realtor is comparable to that of humans, suggesting AI Realtor does not  
1776 rely on hallucination or deception to persuade users.  
1777

## F PROMPTS

### F.1 KEYWORD EXTRACTION PROMPT

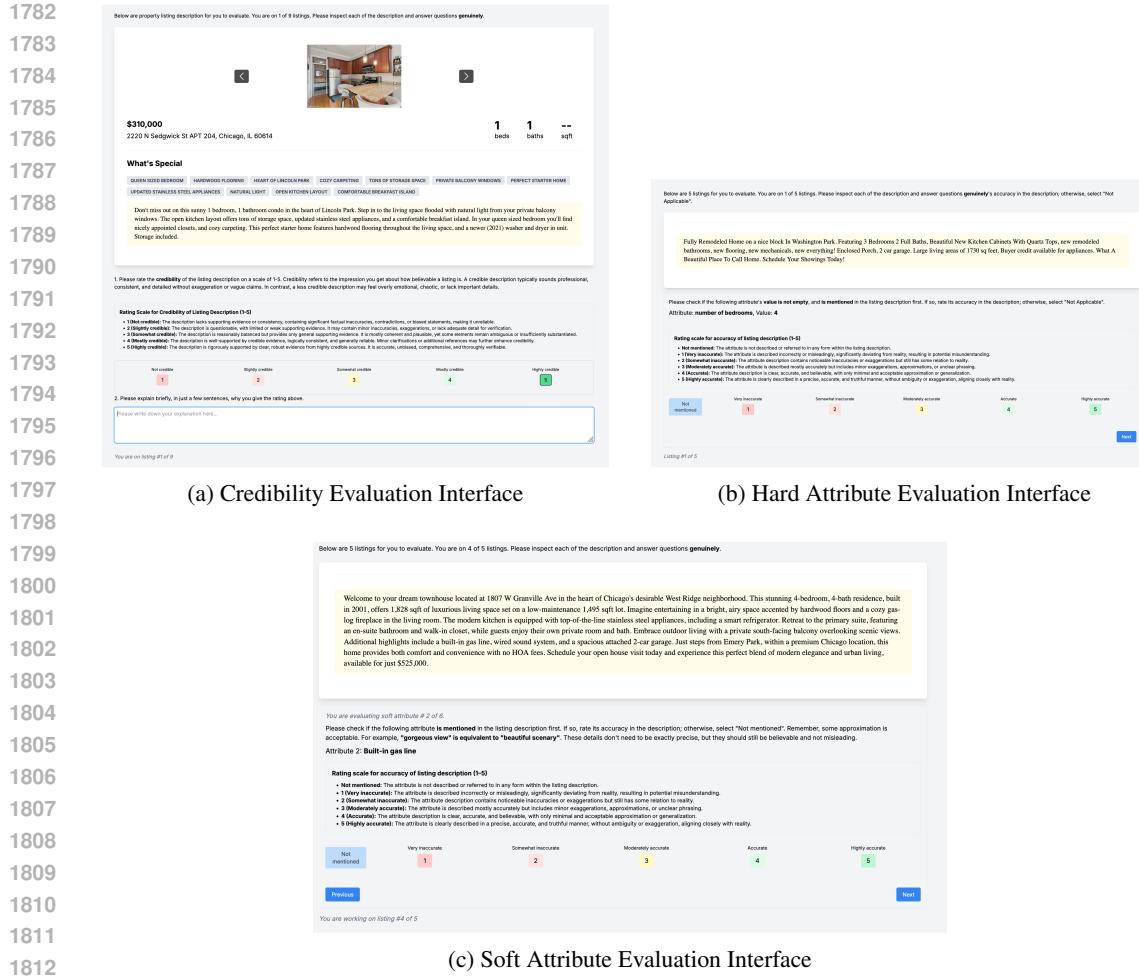


Figure 11: Interfaces used in the hallucination checks.

1782 'Your task is to extract attractive keywords. (e.g., 'modern  
1783 → amenities', 'great views', 'lush landscaping', 'bamboo  
1784 → flooring'). Please express these keywords as phrases or  
1785 → single word from the following house description. Each  
1786 → keyword should be separated by a comma. \n\nDescription: {  
1787 → desc}\n\nKeywords:

## F.2 KEYWORD EXTRACTION NORMALIZATION PROMPT

1782 "Please remove the quantifiers, numbers, adjectives or any  
1783 → modifiers in the provided input. "  
1784 "Uppercase or lowercase doesn't matter. "  
1785 "If the given input is already precise enough, please provide  
1786 → the same input."  
1787 "If you are not sure what to do, please also provide the input  
1788 → as it is. "  
1789 "Do not explain or provide additional information."  
1790 "Here are a few examples:  
1791 "\n\nInput: Two Bedrooms.\n\nOutput: Bedrooms."  
1792 "\n\nInput: Newly Renovated Kitchen.\n\nOutput: Kitchen."  
1793 "\n\nInput: landscape. \n\nOutput: landscape."  
1794 "[Example Ends]"

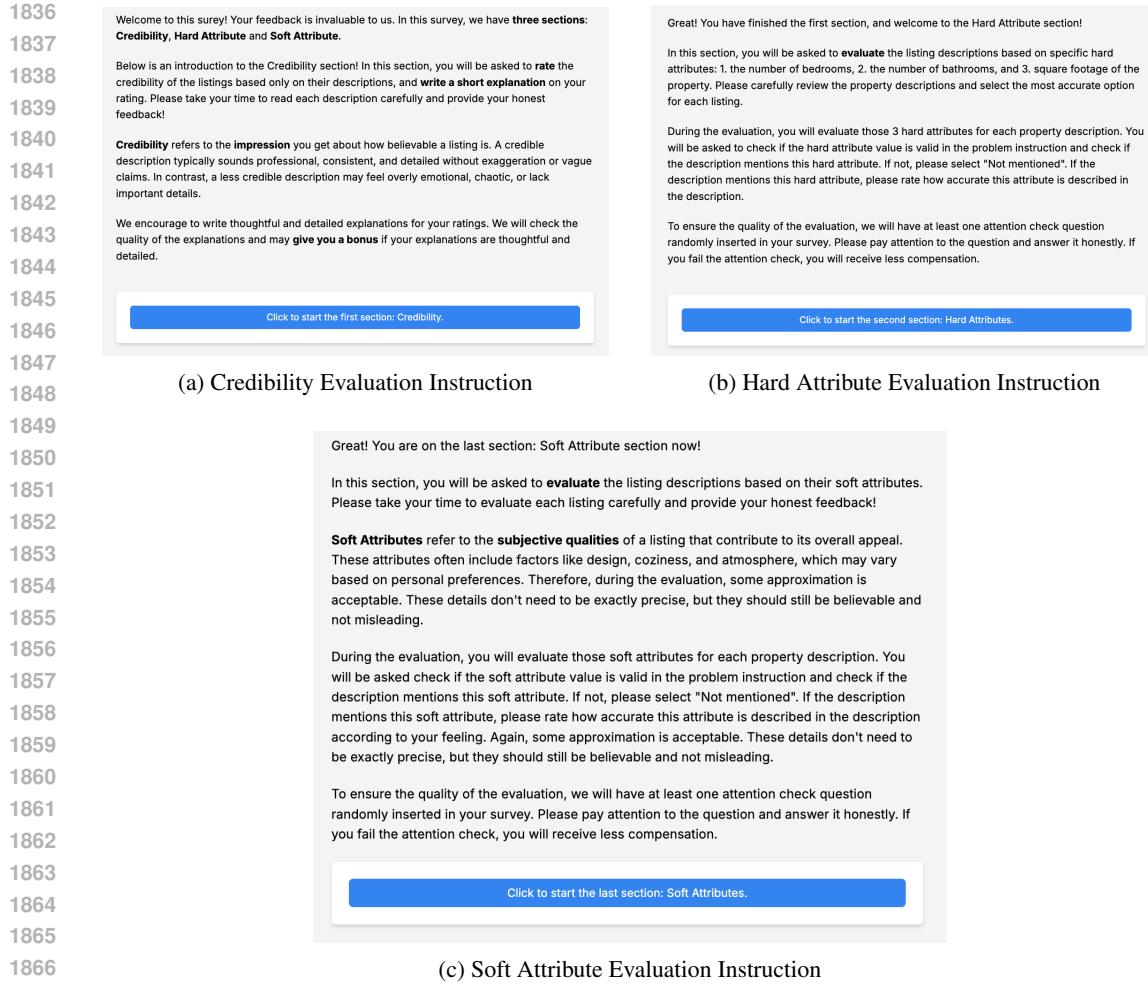


Figure 12: Interfaces used in the hallucination checks.

1870 "Now, given the Input, please precisely provide the Output."  
1871 "\n\nInput: {} \n\nOutput (should only be a noun phrase or  
1872 ↪ keyword) : "  
1873

### 1874 F.3 SCHEMA INDUCTION PROMPT

1875 Here is an initial listing keyword schema that I have, but it may  
1876 ↪ not be comprehensive. I have a manually extracted  
1877 ↪ comprehensive keyword list, but there are many duplicated  
1878 ↪ words (e.g., different keywords may bear similar semantic  
1879 ↪ meanings) and some of them may inspire new categories in  
1880 ↪ this schema. I will give you that 1k+ keyword list in a file  
1881 ↪ and the schema below. Can you do it this way: for every 100  
1882 ↪ keywords in the file, either try to assign it to one of the  
1883 ↪ categories below, or create a new (sub)category and assign  
1884 ↪ the keyword to this new (sub)category. You CANNOT use too  
1885 ↪ broad categories like "others" "misc" and "uncategorized".  
1886 ↪ Only create informative categories if necessary. Give me the  
1887 ↪ final zip files containing all 100-ish intermediate  
1888 ↪ assignment results. Each result should be represented as a  
1889 ↪ JSON-like file with key=subcategory, value=[

```

1890     ↪ list_of_original_keywords_in_file], or key=category, value=
1891     ↪ subcategory (in other words, I want a rich hierarchical
1892     ↪ structure with the leaf nodes as a list of original keywords
1893     ↪ in the file).
1894
1895     #####schema#####
1896     Appliances:
1897         Refrigerator
1898         Oven
1899         Dishwasher
1900         Washer/Dryer
1901         Microwave
1902         Garbage Disposal
1903
1904     Transportation:
1905         Garage
1906         Carport
1907         Parking Space
1908         Public Transit Nearby
1909
1910     Interior Features:
1911         Hardwood Floors
1912         Fireplace
1913         Central Air Conditioning
1914         Walk-in Closet
1915         Open Floor Plan
1916         High Ceilings
1917
1918     Exterior Features:
1919         Balcony
1920         Patio
1921         Deck
1922         Fenced Yard
1923         Garden
1924         Pool
1925
1926     Building Features:
1927         Elevator
1928         Fitness Center
1929         Laundry Room
1930         Security System
1931         Concierge
1932
1933     Utilities:
1934         Water
1935         Gas
1936         Electricity
1937         Cable/Satellite TV
1938         Internet
1939
1940     Neighborhood Features:
1941         Nearby Schools
1942         Parks
1943         Shopping Centers
1944         Restaurants
1945         Hospitals
1946         Recreation Facilities

```

1944     **F.4 FEATURE EXTRACTION BASED ON DESCRIPTION PROMPT**  
 1945  
 1946     "Your task is to determine whether the given feature is mentioned  
 1947        ↳ in the description. The meaning of the feature will be  
 1948        ↳ explained by example keywords. Only respond with 'YES' or '  
 1949        ↳ NO'. "  
 1950     "Feature: {feature\_name}. \n\nExample Keywords for explaining this  
 1951        ↳ feature: {keywords}\n\n"  
 1952     "\n\nDescription: {human\_description}\n\nResponse (Yes/No) : "  
 1953  
 1954  
 1955     **F.5 PERSUASIVE LANGUAGE GENERATION WITH PERSONALIZED FEATURES**  
 1956  
 1957     "Your task is to generate a marketing description for a real  
 1958        ↳ estate listing with the provided features to highlight, and  
 1959        ↳ the client's preferences.  
 1960        - The listing has the following attributes:\n{attributes}  
 1961        - The listing has the following features (accounted for the  
 1962            ↳ client's preference) that are worth highlighting:\n{  
 1963            ↳ highlight\_features\_reweighted }  
 1964        - The client has the following general preferences:\n{  
 1965            ↳ user\_preference}  
 1966        - The client has the following specific preferences over  
 1967            ↳ features:\n{feature\_preference}  
 1968        - You should emphasize the feature or attributes that matches  
 1969            ↳ with the user's preference.  
 1970        Make sure the description is persuasive while concise under  
 1971            ↳ one paragraph."  
 1972  
 1973  
 1974     **F.6 PERSUASIVE LANGUAGE GENERATION WITH LOCALIZED FEATURE PROMPT**  
 1975  
 1976     "Your task is to generate a marketing description for a real  
 1977        ↳ estate listing with the provided features to highlight and a  
 1978        ↳ list of attributes that are competitive among similar  
 1979        ↳ listings."  
 1980        - The listing has the following attributes:\n{attributes}  
 1981        - Compared with {K} similar listings, the listing stands out  
 1982            ↳ in the following features that you want to emphasize:  
 1983            {surprise\_features}  
 1984        - Compared with listings in Chicago, the following features of  
 1985            ↳ this listing are competitive:\n{city\_rankings}  
 1986        - Compared with listings in this neighborhood {neighbourhood},  
 1987            ↳ the following features of this listing are competitive  
 1988            ↳ :\n{neighbourhood\_rankings}  
 1989        - Compared with listings in this zipcode {zipcode}, the  
 1990            ↳ following features of this listing are competitive:\n{  
 1991            {zipcode\_rankings}  
 1992        - Finally, You should explicitly highlight the listing  
 1993            ↳ features or attributes that stands out above or those  
 1994            ↳ ones that exactly matches with the user's preferences as  
 1995            ↳ a surprise factor.  
 1996        Make sure the description is persuasive while concise under  
 1997            ↳ one paragraph."

1998  
1999

## F.7 USER SIMULATION PROMPT

2000 To avoid positional bias as demonstrated in (Zheng et al., 2023), for each pairwise comparisons  
 2001 of descriptions generated by different models, we will prompt the GPT-4o-mini twice to generate  
 2002 separate scores as integers within [0, 100], and compare the final scores to decide which model  
 2003 wins. The prompt below shows an example of this prompt to obtain GPT-4o-mini judgement for the  
 2004 first description presented. “Description 0” and “Description 1” refers to descriptions generated by  
 2005 different models and are randomly shuffled.

2006 You will be given a user profile, a listing and two descriptions  
 2007 ↪ of this listing. Optionally, you may also be given the user'  
 2008 ↪ s history of preferences. Your task is to predict which  
 2009 ↪ description the user would prefer. \n\n  
 2010 User Profile: {user\_profile}  
 2011 Listing: {listing}\n\n  
 2012 Description 0: {description\_0}\n\n  
 2013 Description 1: {description\_1}\n\n  
 2014 Please first generate an analysis of the user's profile and  
 2015 ↪ history (if available), and then analyze why the user might  
 2016 ↪ prefer the first description. You can use the following  
 2017 ↪ format: 'The user might prefer the first description because  
 2018 ↪ ...'  
 2019 The score for the first description (an integer within [0, 100]):

2020  
2021

## F.8 RETRIEVER CONFIGURATION

2022 "mappings": {  
 2023 "properties": {  
 2024 "bedrooms": {"type": "float"},  
 2025 "bathrooms": {"type": "float"},  
 2026 "price": {"type": "float"},  
 2027 "description": {"type": "text"},  
 2028 "area": {"type": "float"},  
 2029 "street\_address": {"type": "text"},  
 2030 "home\_type": {"type": "keyword"},  
 2031 "state": {"type": "keyword"},  
 2032 "city": {"type": "keyword"},  
 2033 "page\_view\_count": {"type": "float"},  
 2034 "favorite\_count": {"type": "float"},  
 2035 "home\_insights": {"type": "keyword"},  
 2036 "neighborhood\_region": {"type": "keyword"},  
 2037 "id": {"type": "keyword"}  
 2038 }  
 2039 }  
 2040  
 2041  
 2042  
 2043  
 2044  
 2045  
 2046  
 2047  
 2048  
 2049  
 2050  
 2051