Leveraging Large Models for Evaluating Novel Content A Case Study on Advertisement Creativity

Anonymous ACL submission

Abstract

Evaluating creativity is challenging, even for humans, not only because of its subjectivity but also because it involves complex cognitive processes. Inspired by work in marketing, we attempt to break down visual advertisement creativity into atypicality and originality. With fine-grained human annotations on these dimensions, we propose a suit of tasks specifically for such a subjective problem. We also evaluate the alignment between state-of-the-art (SoTA) vision language models (VLM) and humans on our proposed benchmark, demonstrating both the promises and challenges of using VLMs for automatic creativity assessment.

1 Introduction

001

005

011

017

024

027

Evaluation of creativity has been a lasting effort where researchers have taken diverse sets of perspectives (Said-Metwaly et al., 2017). Among those perspectives, evaluating the created product is the only one that focuses purely on the created artifact, no matter who authored that artifact, greatly reducing the complexity of problem framing. Examples from this perspective include evaluation of creative writing (Lu et al., 2024; Chakrabarty et al., 2024a; Porter and Machery, 2024; Marco et al., 2024), tool use (Tian et al., 2024), and advertisement (Smith et al., 2007).

In advertising, creativity plays a critical role in driving consumer behaviors, where multiple studies have shown positive effects of creative ads, including increased purchase intent and positive brand impression (Sharma, 2012; Terkan, 2014). Therefore, ad creators are motivated to consistently develop and evaluate creative ad content. Extensive research has been conducted to understand what the general public would consider creative (El-Murad and West, 2004; Rosengren et al., 2020; Swee Hoon Ang and Lou, 2014; Smith et al., 2007), many of which consider ads creativity to be a combination of "atypicality" (or, synthesis, abnormalness,



Figure 1: Top: two ads from dataset; Middle: human rating and VLM outputs (25 each); Bottom: average human rating, VLM outputs, and VLM pairwise predictions; Scores are 3-scale, 3 being the best. Ad **A** receives a higher average rating in all three categories compared to Ad **B**, also reflected in the model-predicted ratings and pairwise preference prediction.

etc.) and "originality" (or, novelty, uniqueness, etc.). Here is an example that is both atypical and original: in the advertisement (A) in Figure 1, the image of a cow typing on a computer is highly atypical, as cows do not do that; the text "*Eat chikin or I'll de-friend U*", which consists of multiple slangs, the cow, the computer, and the small logo of Chickfil-A is a rarely-seen combination given that it is a fast-food advertisement, making it very original. Decoding ads creativity under such a framework requires extensive reasoning, making the evaluation of creativity a challenging task. Unsurprisingly, previous work heavily relies on domain experts, who are expensive and inaccessible.

041

Recently, foundational models demonstrate impressive performances in other evaluation tasks, such as summarization (Zhong et al., 2022), Long-Form QA (Jiang et al., 2023), and commonsense text generation (Xu et al., 2023), many of which were previously dominant by human evaluation. For creativity evaluation, prior work (Chakrabarty et al., 2024a; Atmakuru et al., 2024; Ismayilzada et al., 2024) explored the ability of LLMs to access writing creativity. This poses the question of whether we can leverage foundational models to expand automatic evaluation of creativity to multimodal data such as visual advertisements, with the help of Vision Language Models (VLM).

061

063

064

067

077

089

096

100

101

102

104

To this end, we conduct several fine-grained, automatic creativity evaluations for visual advertisements, including creativity, originality, and atypicality. We decompose creativity into atypicality and originality and then collect high-quality human ratings of advertisements in those dimensions, as shown in Figure 1 as "human rating". We experiment with state-of-the-art (SoTA) VLMs to predict these ratings and examine the human-model alignment in both intrinsic (i.e., one image at a time) and pairwise fashion. In contrast to the traditional emphasis on prediction accuracy, we evaluate models' ability to capture the task's subjective nature and to gauge annotator disagreements. We also show that VLMs perform impressively in a pairwise setting, reasonably well in distribution modeling, and less effectively in disagreement. Our results highlight the importance of both high-quality annotations and distribution modeling as the objective, instead of majority label prediction, in future research on subjective tasks. We believe our benchmark and evaluation metrics provide a solid foundation for utilizing VLMs to assist visual content creators.

2 Related Work

Evaluation of Creativity Research in evaluation of creativity includes cognitive science (Said-Metwaly et al., 2017; Simonton, 2012; James Lloyd-Cox and Bhattacharya, 2022), marketing (El-Murad and West, 2004; Rosengren et al., 2020; Swee Hoon Ang and Lou, 2014; Smith et al., 2007), creative writing (Skalicky, 2022), human computer interaction (Chakrabarty et al., 2024b), and artificial intelligence (Chakrabarty et al., 2023, 2024a; Atmakuru et al., 2024; Ismayilzada et al., 2024). There are two consensus. First, creativity requires both divergence and effectiveness. We ensure effec-

Dimension	Questions	Answer
Atypicality	The ad connected usually unrelated objects The ad contained unusual connection The ad brought unusual items together	agree (1), neutral (0), disagree (-1)
Originality	The ad was out of the ordinary The ad broke away from habit-bound and stereotypical thinking The ad was unique	agree (1), neutral (0), disagree (-1)
Creativity	What is the overall level of creativity of this advertisement?	integer (1-5)

Table 1: Questions in Amazon Mechanical Turk; full annotation interface and instructions is in Figure 2.

106

107

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

tiveness via quality check questions in human annotation and mainly focus on modeling divergence with atypicality and originality. Second, evaluation of creativity is subjective. This motivates our distribution modeling task instead of a majority-label prediction task. Our work is closely related to two groups of works. Smith et al. (2007) focused on advertisement images and proposed five creativity dimensions, including atypicality and originality. We adopt their creativity decomposition. Chakrabarty et al. (2024a); Atmakuru et al. (2024); Ismayilzada et al. (2024) use LLMs to evaluate creative writing; in contrast, we instead study the alignment between VLM judgments and human ratings.

Automatic Evaluation with Foundation Models GPTScore (Fu et al., 2023) and UniEval (Zhong et al., 2022) decompose the evaluation of a complex task into ones that can be done by language models; whereas PandaLM (Wang et al., 2024) focuses on pairwise evaluation for free-form text quality. In vision domain, (Jayasumana et al., 2024; Otani et al., 2023) explore evaluating generated images with CLIP embeddings. These works focus on single modality, instead of our image-text pairs.

3 Dataset

We use the Pitt Ads Dataset (referred to as Pitt-Ads) as our basis, where each ad image is annotated with its topic, expected actions from viewers after seeing the ad, binary labels of atypical objects in it (when applicable), and the topic of the ad (10 topic groups in total) (Hussain et al., 2017; Ye et al., 2019). We sample 100 ads and collect fine-grained human creativity annotations (Creative-100 in 3.1); we also sample an additional 300 ads from the remaining data points for atypicality prediction(Atypical-300 in 3.2).

3.1 Creative-100

Creative-100 consists of 100 ads, with 10 from each topic group: food, pet, drinks, automobile, electronics, service, education, beauty, healthcare, clothing, home, leisure, shopping, and noncommercial. For fine-grained creativity evaluation, we include two dimensions of creativity: originality and atypicality, two most influential ones according to (Smith et al., 2007). Human annotations are then collected via Amazon Mechanical Turk (Mturk) to represent fine-grained ratings in all three dimensions: originality, atypicality, and creativity (see Figure 2 for annotation interface and instructions).

145

146

147

148

149

150

151

152

153

154

155

156

158

159

160

163

164

165

166

167

170

171

172

173

175

176

178

179

181

183

184

187

191

Due to the inherent subjectivity of creativity judgment, we formulate the measurement of creativity as several multiple-choice questions with possible answers as a categorical distribution of those choices. Hense, the predictive target is not a binary label (e.g., "creative") but a distribution of human ratings. This motivates us to collect 25 annotations per ad image to approximate the true rating distribution within certain error rate (McHugh, 2012). Refer to Appendix B.2 for more details.

For atypicality and originality, we follow Smith et al. (2007) and record responses about various statements (Table 1). We record a 5-scale overall score for creativity and convert it to a 3-scale, aligning with other dimensions. We also include a quality check question, asking annotators to choose the action after seeing a given ad (e.g., "I should go to Chick-fil-A" for ad A in Figure 1). Five actions are given, with one correct action and four randomly sampled from Pitt-Ads. Annotators get 96.88% accuracy in this question, highlighting their accurate understanding of visual advertisements. More dataset construction details are in Appendix B.

Atypical-300 3.2

We also randomly sampled 300 ads (Atypical-300) from Pitt-Ads, where 185(62%) include atypical object(s). Different from Creative-100, each ad here only has three binary annotations on atypicality. Both Smith et al. (2007) and Creative-100 (Appendix B.4) show that atypicality has a positive correlation with creativity. Thus, we include this dataset to further study VLM's ability to evaluate ads creativity.

4 **Experimental Setup**

4.1 Models

189 We experiment with open-sourced VLMs, InternVL2-8B (Chen et al., 2024), LLaVA-7B and 190 13B (Li et al., 2024), and close-sourced VLMs, GPT4-v (OpenAI et al., 2024). All experiments are done with zero-shot prompting with vllm (Kwon 193

et al., 2023). More details are in Appendix C.

4.2 Task Formulation

We define the following three tasks for each dimension (i.e., creativity, originality, and atypicality):

Distribution Modeling is designed to evaluate the model's ability to simulate human group behavior when it comes to creativity ratings. In practice, we prompt VLMs multiple times with high temperatures to get the same number of VLM outputs as the number of annotators. In this way, we simulate a "group behavior" instead of a single-point judgment about the level of creativity in the ad. To evaluate the quality of this simulation, we use two metrics: Spearman's correlation between the average rating from humans and that from VLMs and the average KL Divergence between human rating distribution and that from VLMs. These two results are in Rating Correlation and Distribution Divergence columns in Table 2.

Disagreement Prediction tries to capture the annotator's level of disagreement, which is important in domains like marketing. An ad with high creativity ratings and minimal disagreement is desired and could have a more positive impact on the product. In practice, we directly prompt VLMs to predict the level of disagreement (low, middle, or high) for each scoring dimension. We then compute Spearman's correlation between the prediction and standard deviation of human ratings. This metric studies the level of creativity ambiguity of the advertisements. A very creative ad will have a low disagreement rate with a high creativity score. The results are in Disagreement column in Table 2.

Pairwise Preference aims at evaluating the model's ability to correctly pick the more creative ad out of two ads, given that an absolute rating of creativity can be hard when there is no reference. For each scoring dimension, we include all ad pairs with average human ratings differences greater than 0.5. For Creative-100, we have 938, 2708, and 2631 pairs in creativity, originality, and atypicality; for Atypical-300, we sampled 1000 image pairs from 300 images due to constraints in computation resources. The results are evaluated by F1 score and are shown in Pairwise column in Table 2.

5 **Results**

Promising Results in Rating Correlation For all dimensions in Creative-100, the correlations between average human and VLM ratings are both

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

229

230

231

232

233

234

235

236

237

238

239

240

241

242

195

		Distribution Modeling		Distribution Modeling Disagreement		Pairwise	
Dimension	Model	Rating Correlation	Distribution Divergence		All	Easy	Hard
		$\uparrow R (p-value)$	\downarrow KL Divergence	$\uparrow R (p-value)$	$\uparrow Fl$	$\uparrow Fl$	$\uparrow Fl$
	LLaVA-7B	0.65 (.00*)	1.01	0.06 (.52)	0.78	0.80	0.75
Creativity	LLaVA-13B	0.65 (.00*)	0.37	nan	0.67	0.69	0.65
(Creative-100)	InternVL2-8B	0.72 (.00 [*])	1.45	0.20 (.05)	0.80	0.85	0.75
	GPT-4v	-	-	-0.04 (.72)	0.97	0.98	0.96
	LLaVA-7B	0.76 (.00*)	0.56	0.07 (.49)	0.73	0.80	0.67
Originality	LLaVA-13B	0.71 (.00*)	0.30	nan	0.67	0.68	0.67
(Creative-100)	InternVL2-8B	0.82 (.00*)	0.57	0.11 (.27)	0.69	0.71	0.68
	GPT-4v	-	-	0.15 (.13)	0.87	0.93	0.85
	LLaVA-7B	0.75 (.00*)	0.40	0.17 (.08)	0.76	0.81	0.72
Atypicality	LLaVA-13B	0.63 (.00*)	0.26	nan	0.69	0.68	0.69
(Creative-100)	InternVL2-8B	0.78 (.00*)	0.40	0.24 (.02*)	0.78	0.81	0.75
	GPT-4v	-	-	-0.07 (.47)	0.89	0.94	0.80
	LLaVA-7B	0.21 (.00*)	0.32	0.01 (.92)	0.79	-	-
Atypicality	LLaVA-13B	0.17 (.00*)	0.26	-0.05 (.43)	0.66	-	-
(Atypical-300)	InternVL2-8B	0.23 (.00*)	0.35	0.02 (.75)	0.77	-	-
	GPT-4v	-	-	-0.00 (.96)	0.90	-	-

Table 2: **Bold**: best-performing models; *: statically significant results ($\alpha = 0.05$); *nan*: disagreement predictions are uniform, making correlation test fail; "-" in GPT-4v rows: no distribution modeling task is done due to budget constraint; "-" in Pairwise columns: the classification of easy and hard is not available in Atypical-300.

high and statistically significant, with InternVL2-8B being the best-performing model in every dimension. However, the correlations are much lower in Atypical-300, although they are statistically significant. We believe this is due to the small annotation size (3 per ad) in Atypical-300, which can easily be biased by one annotation data point, leading to an overall more noisy distribution of ratings. Distribution Divergence also shows promising results, with the lowest divergence achieved by the LLaVA-13B model. Cross-dataset disparity is also much lower, where the KL divergence is similar for the atypicality in both datasets. Output examples and reasoning text word cloud in Appendix D.2.

243

244

245

246

247

253

256

257

259

260

262

263

264

267

268

269

270

Disagreement Prediction Remains Challenging For all scenarios in Disagreement Prediction (except for Atypicality in Creative-100, with InternVL2-8B), disagreement predictions have no statistical correlation with human rating standard deviations. For LLaVA-13B, all outputs are "middle", making the correlation result *nan*. This suggests that using VLM as a group-opinion synthesizer remains challenging ¹.

Great Performance in Pairwise Preference Results from the pairwise preference task are very impressive, with the best-performing GPT-4v achieving more than 0.9 F1 score. We also further analyze the performance by dividing image pairs into "easy" and "hard" subsets: compared to median human rating differences if a pair of ad images has a higher absolute difference, it is an easy image pair, and vice versa. As shown in Table 2, all VLMs perform better in easy tasks than hard ones, confirming the alignment with human preference differences.

271

272

273

274

275

276

277

278

279

280

281

284

285

287

289

290

291

293

294

295

296

297

298

299

Smaller Models' Superior Performance LLaVA-7B and InternVL2-8B consistently outperform LLaVA-13B in Rating Correlation and Pairwise Preference, both of which are ranking-based evaluations. Based on the error analysis (Appendix D.1), we believe this can be explained by ranking tasks requiring higher reasoning capability from the language part of the VLM. The language parts of those two smaller models are Mistral-7B and InternLM2.5-7B-Chat, both having a higher ranking on the HuggingFace Open LLM LeaderBoard ² in reasoning tasks compared to the language part (Vicuna-13B) of LLaVA-13B.

6 Conclusion

We present a case study of using SoTA VLMs to evaluate creativity in advertisements. Inspired by marketing research, we collect fine-grained human annotations. We show decent alignment between VLMs and humans in Distribution Modeling and Pairwise Preference, whereas Disagreement Prediction remains challenging. Our work opens the opportunity for automatic evaluation of advertisement creativity by providing a benchmark and metrics.

¹We also calculated the correlation between standard deviations of model predictions and that of human ratings, but the correlations are all near zero.

²HuggingFace LLM Leaderboard

300 301

302

310

311

313

314

315

316

319

320

321

323

325

326

327

330

332

341

342

343

345

347

7 Limitations

One obvious limitation is the size of our dataset. The fine-grained creativity annotation only consists of 100 ad images. Two bottlenecks that lead to such a limited number is budget and annotation quality. Since we want to explore distribution modeling, we need more annotation than typical machine learning tasks, leading to a huge budget requirement. Also because of the dataset size, we are not able to conduct fine-tuning experiments, as we only have 100 high-quality annotations. We will leave exploration of fine-tuning experiments to future work.

Another limitation is the natural biases contained in our annotation as a majority of our annotators are located in the U.S. We have plans to expand the annotation to other platforms (e.g., LabInTheWild) where a more diverse set of annotators is available. We would also suggest that researchers be cautious when applying our method to data from other countries or languages. Future work could also explore alternative prompting approaches to simulate group behavior or conduct a demographic analysis of human annotations, which could check whether VLM holds opinions comparable to those of particular groups.

Regarding the design of simulating "group behavior" by prompting a VLM 25 times with the same prompt, we recognize the simplicity of how we prompted the VLMs to make predictions. However, creating 25 different prompts for each ad could complicate the analysis and results. For example, certain prompts could disproportionately increase the likelihood of a "creative" label compared to others. One interesting approach to generating prompts more meaningfully could involve exploring persona prompts to simulate multiple annotators' behavior. However, we believe that starting with a straightforward approach using the sampling strategy is essential. Therefore, we leave the use of 25 distinct persona prompts for future work.

Also, due to hardware constraints, we only experiment with VLMs in the 7B to 13B range when much larger models, such as LLaVA-34B, are available. We will leave more extensive prompt tuning and model selections to future work.

References

Anirudh Atmakuru, Jatin Nainani, Rohith Siddhartha Reddy Bheemreddy, Anirudh Lakkaraju, Zonghai Yao, Hamed Zamani, and Haw-Shiuan Chang. 2024. CS4: Measuring the creativity of large language models automatically by controlling the number of story-writing constraints. *arXiv* [*cs.CL*].

350

351

352

353

354

358

359

360

361

362

363

364

365

366

367

368

369

370

371

372

373

374

375

376

377

378

379

381

382

383

384

385

389

390

392

393

394

395

396

397

398

400

401

402

403

404

405

- Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. 2024a. Art or artifice? large language models and the false promise of creativity. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.
- Tuhin Chakrabarty, Vishakh Padmakumar, Faeze Brahman, and Smaranda Muresan. 2024b. Creativity support in the age of large language models: An empirical study involving emerging writers. *Preprint*, arXiv:2309.12570.
- Tuhin Chakrabarty, Arkadiy Saakyan, Olivia Winn, Artemis Panagopoulou, Yue Yang, Marianna Apidianaki, and Smaranda Muresan. 2023. I spy a metaphor: Large language models and diffusion models co-create visual metaphors. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7370–7388, Toronto, Canada. Association for Computational Linguistics.
- Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, Ji Ma, Jiaqi Wang, Xiaoyi Dong, Hang Yan, Hewei Guo, Conghui He, Botian Shi, Zhenjiang Jin, Chao Xu, Bin Wang, Xingjian Wei, Wei Li, Wenjian Zhang, Bo Zhang, Pinlong Cai, Licheng Wen, Xiangchao Yan, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhai Wang. 2024. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Preprint*, arXiv:2404.16821.
- Qi Cheng, Michael Boratko, Pranay Kumar Yelugam, Tim O'Gorman, Nalini Singh, Andrew McCallum, and Xiang Lorraine Li. 2024. Every answer matters: Evaluating commonsense with probabilistic measures. *arXiv preprint arXiv:2406.04145*.
- Jaafar El-Murad and Douglas C. West. 2004. The definition and measurement of creativity: What do we know? *Journal of Advertising Research*, 44(2):188– 201.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *Preprint*, arXiv:2302.04166.
- Zaeem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. 2017. Automatic understanding of image and video advertisements. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1100–1110.
- Mete Ismayilzada, Claire Stevenson, and Lonneke van der Plas. 2024. Evaluating creative short story generation in humans and large language models. *arXiv* [cs.CL].

- 407 408
- 411
- 409 410

412

413

414

415

416

417

418

419

420

421

422

423

424

425

426

497

428

429

430

431

432

433

434

435

436

437

438

439

440

441

442

443

444

445

446

447

448

449

450

451

452 453

454

455 456

457

458

459

460

461

462

463

- Alan Pickering James Lloyd-Cox and Joydeep Bhattacharya. 2022. Evaluating creativity: How idea context and rater personality affect considerations of novelty and usefulness. Creativity Research Journal, 34(4):373–390.
- Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. 2024. Rethinking fid: Towards a better evaluation metric for image generation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 9307–9315.
- Dongfu Jiang, Yishan Li, Ge Zhang, Wenhao Huang, Bill Yuchen Lin, and Wenhu Chen. 2023. Tigerscore: Towards building explainable metric for all text generation tasks. ArXiv, abs/2310.00752.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles.
- Bo Li, Kaichen Zhang, Hao Zhang, Dong Guo, Renrui Zhang, Feng Li, Yuanhan Zhang, Ziwei Liu, and Chunyuan Li. 2024. Llava-next: Stronger llms supercharge multimodal capabilities in the wild.
- Ximing Lu, Melanie Sclar, Skyler Hallinan, Niloofar Mireshghallah, Jiacheng Liu, Seungju Han, Allyson Ettinger, Liwei Jiang, Khyathi Chandu, Nouha Dziri, and Yejin Choi. 2024. Ai as humanity's salieri: Quantifying linguistic creativity of language models via systematic attribution of machine text against web text. Preprint, arXiv:2410.04265.
- Guillermo Marco, Julio Gonzalo, M Teresa Mateo-Girona, and Ramón Del Castillo Santos. 2024. Pron vs prompt: Can large language models already challenge a world-class fiction author at creative text writing? In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, pages 19654-19670, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Mary L McHugh. 2012. Interrater reliability: the kappa statistic. Biochemia medica, 22(3):276-282.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben

Chess, Chester Cho, Casey Chu, Hyung Won Chung, 464 Dave Cummings, Jeremiah Currier, Yunxing Dai, 465 Cory Decareaux, Thomas Degry, Noah Deutsch, 466 Damien Deville, Arka Dhar, David Dohan, Steve 467 Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, 468 Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, 469 Simón Posada Fishman, Juston Forte, Isabella Ful-470 ford, Leo Gao, Elie Georges, Christian Gibson, Vik 471 Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-472 Lopes, Jonathan Gordon, Morgan Grafstein, Scott 473 Gray, Ryan Greene, Joshua Gross, Shixiang Shane 474 Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, 475 Yuchen He, Mike Heaton, Johannes Heidecke, Chris 476 Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, 477 Brandon Houghton, Kenny Hsu, Shengli Hu, Xin 478 Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, 479 Joanne Jang, Angela Jiang, Roger Jiang, Haozhun 480 Jin, Denny Jin, Shino Jomoto, Billie Jonn, Hee-481 woo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Ka-482 mali, Ingmar Kanitscheider, Nitish Shirish Keskar, 483 Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, 484 Christina Kim, Yongjik Kim, Jan Hendrik Kirch-485 ner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, 486 Łukasz Kondraciuk, Andrew Kondrich, Aris Kon-487 stantinidis, Kyle Kosic, Gretchen Krueger, Vishal 488 Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan 489 Leike, Jade Leung, Daniel Levy, Chak Ming Li, 490 Rachel Lim, Molly Lin, Stephanie Lin, Mateusz 491 Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, 492 Anna Makanju, Kim Malfacini, Sam Manning, Todor 493 Markov, Yaniv Markovski, Bianca Martin, Katie 494 Mayer, Andrew Mayne, Bob McGrew, Scott Mayer 495 McKinney, Christine McLeavey, Paul McMillan, 496 Jake McNeil, David Medina, Aalok Mehta, Jacob 497 Menick, Luke Metz, Andrey Mishchenko, Pamela 498 Mishkin, Vinnie Monaco, Evan Morikawa, Daniel 499 Mossing, Tong Mu, Mira Murati, Oleg Murk, David 500 Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, 501 Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, 502 Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex 503 Paino, Joe Palermo, Ashley Pantuliano, Giambat-504 tista Parascandolo, Joel Parish, Emy Parparita, Alex 505 Passos, Mikhail Pavlov, Andrew Peng, Adam Perel-506 man, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Poko-508 rny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, 510 Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, 511 Cameron Raymond, Francis Real, Kendra Rimbach, 512 Carl Ross, Bob Rotsted, Henri Roussez, Nick Ry-513 der, Mario Saltarelli, Ted Sanders, Shibani Santurkar, 514 Girish Sastry, Heather Schmidt, David Schnurr, John 515 Schulman, Daniel Selsam, Kyla Sheppard, Toki 516 Sherbakov, Jessica Shieh, Sarah Shoker, Pranav 517 Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, 518 Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin 519 Sokolowsky, Yang Song, Natalie Staudacher, Fe-520 lipe Petroski Such, Natalie Summers, Ilya Sutskever, 521 Jie Tang, Nikolas Tezak, Madeleine B. Thompson, 522 Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, 523 Preston Tuggle, Nick Turley, Jerry Tworek, Juan Fe-524 lipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, 525 Chelsea Voss, Carroll Wainwright, Justin Jay Wang, 526 Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, 527

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

582

583

CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. Gpt-4 technical report. *Preprint*, arXiv:2303.08774.

528

529

535

542

543

547

548

549

553

556

557

558

559

561

562

563

564

566

571

573

574

575

576

577 578

579

581

- Mayu Otani, Riku Togashi, Yu Sawai, Ryosuke Ishigami, Yuta Nakashima, Esa Rahtu, Janne Heikkilä, and Shin'ichi Satoh. 2023. Toward verifiable and reproducible human evaluation for text-toimage generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14277–14286.
 - Brian Porter and Edouard Machery. 2024. AI-generated poetry is indistinguishable from human-written poetry and is rated more favorably. *Sci. Rep.*, 14(1):26133.
 - Sara Rosengren, Martin Eisend, Scott Koslow, and Micael Dahlen. 2020. A meta-analysis of when and how advertising creativity works. *Journal of Marketing*, 84(6):39–56.
 - Sameh Said-Metwaly, Wim Van den Noortgate, and Eva Kyndt. 2017. Approaches to measuring creativity: A systematic literature review. *Creativity. Theories – Research - Applications*, 4(2):238–275.
 - Pooja Sharma. 2012. Advertising effectiveness:" understanding the value of creativity in advertising", a review study in india. *Online Journal of Communication and Media Technologies*, 2(3):1.
 - Dean Keith Simonton. 2012. Quantifying creativity: can measures span the spectrum? *Dialogues in Clinical Neuroscience*, 14(1):100–104. PMID: 22577309.
 - Stephen Skalicky. 2022. Liquid gold down the drain: Measuring perceptions of creativity associated with figurative language and play. *Cognitive Semantics*, 8(1):79 – 108.
 - Robert E. Smith, Scott B. MacKenzie, Xiaojing Yang, Laura M. Buchholz, and William K. Darley. 2007.
 Modeling the determinants and effects of creativity in advertising. *Marketing Science*, 26(6):819–833.
 - Yih Hwai Lee Swee Hoon Ang, Siew Meng Leong and Seng Lee Lou. 2014. Necessary but not sufficient: Beyond novelty in advertising creativity. *Journal of Marketing Communications*, 20(3):214–230.
 - Remziye Terkan. 2014. Importance of creative advertising and marketing according to university students' perspective. *International Review of Management and Marketing*, 4(3):239–246.
 - Yufei Tian, Abhilasha Ravichander, Lianhui Qin, Ronan Le Bras, Raja Marjieh, Nanyun Peng, Yejin Choi,

Thomas Griffiths, and Faeze Brahman. 2024. Mac-Gyver: Are large language models creative problem solvers? In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5303–5324, Mexico City, Mexico. Association for Computational Linguistics.

- Yidong Wang, Zhuohao Yu, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, Wei Ye, Shikun Zhang, and Yue Zhang. 2024. Pandalm: An automatic evaluation benchmark for llm instruction tuning optimization. *Preprint*, arXiv:2306.05087.
- Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Wang, and Lei Li. 2023. Instructscore: Towards explainable text generation evaluation with automatic feedback. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pages 5967–5994.
- Keren Ye, Narges Honarvar Nazari, James Hahn, Zaeem Hussain, Mingda Zhang, and Adriana Kovashka. 2019. Interpreting the rhetoric of visual advertisements. *IEEE transactions on pattern analysis and machine intelligence*, 43(4):1308–1323.
- Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. Towards a unified multidimensional evaluator for text generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2023–2038, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A Ads Dataset

A.1 Terminology

In the original study (Smith et al., 2007) that proposed the breakdown of creativity, they use the term "synthesis" instead of atypicality. However, they defined "synthesis" as "...*combine, connect, or blend normally unrelated objects or ideas*" which is very similiar to "atypicality". We use the term "atypicality" because that is the term used in the Pitt-Ads Dataset and want to keep the terminology aligned.

A.2 Atypicality and Originality as Selected Dimensions of Creativity

The original marketing research (Smith et al., 2007) conducted an exploratory factor analysis (EFA) on five dimensions of divergence in ads creativity (numbers in parenthesis are loading factors of EFA, higher means more correlated with creativity): originality (0.23), synthesis (0.23), artistic value (0.20), flexibility (0.18), and elaboration (0.15). As mentioned previously, their definition of synthesis is
very similar to "atypicality" as defined in Pitt-Ads
Dataset. Therefore, we keep the highest-scoring
two aspects (originality and synthesis/atypicality)
in our study.

A.3 Sampling Process

641

647

654

657

665

The original Pitt Ads dataset consists of 38 topics in total (Hussain et al., 2017). We group those topics into 10 "topic groups": clothing, food, automobile, beauty, leisure, electronics, drinks, service, non-commercial, and healthcare. From each topic group, we sampled 10 ads, which makes up the 100 images in Creativity-100.

Details about "topic groups"

- food: "Restaurants, cafe, fast food", "Chocolate, cookies, candy, ice cream", "Chips, snacks, nuts, fruit, gum, cereal, yogurt, soups", "Seasoning, condiments, ketchup"
- pet: "Pet food"
- drinks: "Alcohol"", "Coffee, tea"", "Soda, juice, milk, energy drinks, water""
- automobile: "Cars, automobiles"
 - electronics: "Electronics"
 - service: "Phone, TV and internet service providers", "Financial services, "Security and safety services", "Software", "Other services"
 - education: "Education"
 - beauty: "Beauty products and cosmetics"
- healthcare: "Healthcare and medications"
 - clothing: "Clothing and accessories"
 - home: "Baby products", "Cleaning products", "Home improvements and repairs", "Home appliances"]
 - leisure: "Games and toys", "Vacation and travel", "Media and arts", "Sports equipment and activities", "Gambling"
- shopping: "Shopping"

non-commercial: "Environment, nature, pollution, wildlife", "Animal rights, animal abuse",
"Human rights", "Safety, safe driving, fire safety", "Smoking, alcohol abuse", "Domestic violence", "Self esteem, bullying, cyber bullying", "Political candidates", "Charities"

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

704

705

706

707

708

709

710

B Data Collection and Post Processing

B.1 Amazon Mechanical Turk Details

Payment for worker Each HIT receives \$0.5 compensation (estimated \$15/hour).

Annotation interface See Figure 2 for the annotation interface. Note that there is a section "artistic values". We dropped that section in the later parts of the experiment because 1) it is very subjective and could be further broken down into more fine-grained subcategories, and 2) to keep our focus on atypicality and originality.

In total, 43 workers contributed to our task and finished 2500 HITs. Their background can be found in Figure 3, 4, 5, 6, and 7. As we can see, the annotators are strongly skewed towards the US-based, white, female, middle age group, which should be kept in mind when applying our methodology when it comes to people from another background.

B.2 Number of Samples for Distribution Task

Following previous works (McHugh, 2012; Cheng et al., 2024), the number of samples required to approximate the real distribution can be calculated as follows:

$$P(D_{KL}(g_{n,k}||f) > \epsilon) \le e^{-n\epsilon} \left[\frac{3c_1}{c_2} \sum_{i=0}^{k-2} k_{i-1} (\frac{e\sqrt{n}}{2\pi})^i \right]$$
 703

 c_1 and c_2 are constant values (based on (McHugh, 2012) $c_1 = 2, c_2 = \frac{\pi}{2}$), k is the number of categories in the categorical distribution (in our case, k = 3), and n is the number of samples. If we fix the left-hand side to be less than 0.1, we would get n has to be 25 (see Figure 8).

B.3 Label Processing

We process the annotation by first converting the
categorical data to numerical values. For atypi-
cality and originality, we code agree, neutral, and
disagreement choices as 1, 0, and -1. As there are
three subquestions for both atypicality and origi-
nality, we simply add up the three scores from each711713
714
715712

Overview

Given an advertisement, provide your opinion on the statements below.

Atypicality: There are uncommon entities (objects, humans, animals, etc) or interactions of entities in the ad.
 Originality: The ad is distinctive to other ads in the same topic.
 Artistic Value: The ad is visually impressive or memorable.
 Effectiveness: The ad promots a strong message about the intended action from viewers. Choose the right action from five choices that viewers would take after seeing this ad
 Overall: The overall creativy of the advertisement is based on your on beliefs

Atypicality

The ad connected objects that are usually unrelated. ○ agree ○ neutral ○ disagree

The ad contained unusual connections. ○ agree ○ neutral ○ disagree

The ad brought unusual items together.

○ agree ○ neutral ○ disagree

Originality

The ad was out of the ordinary. \odot agree \odot neutral \odot disagree

The ad broke away from habit-bound and stereotypical thinking. \odot agree \odot neutral \odot disagree

The ad was unique.

 \odot agree \odot neutral \odot disagree

Artistic Value The ad was visually/verbally distinctive.

 \bigcirc agree \bigcirc neutral \bigcirc disagree

The ad made ideas come to life graphically/verbally. \bigcirc agree \bigcirc neutral \bigcirc disagree

The ad was artistically produced.

 \bigcirc agree \bigcirc neutral \bigcirc disagree

Effectiveness

Given this advertisement, out of these five possible actions, which one is the most likely one?

- \bigcirc a. I should get a porsche O b. I should get some tap shoes.
- O c. i should try this product
- O d. I should eat kfc
- O e. i should want to go here

Overall

What is the overall level of creativity of this advertisement? (1: NOT creative; 5: creative) 0102030405

Background (You only need to answer these questions once!)

- Race Identification (select one or more): American Indian or Alaska Native
 Asian
 Black or African American
 Native Hawaiian or Other Pacific Islander
- Native Hawaiian or Other Page
 White
 Another race not listed here
 I prefer not to say

Are you of Hispanic or Latino/a origin? ○ Yes ○ No ○ I prefer not to say

Which age group do you belong to? ○ Below 18 ○ 18~24 ○ 25~34 ○ 35~44 ○ 45~54 ○ 55~64 ○ 65 and above

What is your gender? ○ Male ○ Female ○ Non Binary ○ I prefer not to say

Which country did you live the longest growing up?

Feedbacks/Questions

If any part of this HIT is confusing or if you have any feedbacks or question for us, please let us know below.

Submit







Figure 3: Distribution of workers' response to "In which country did you live the longest time so far?"



Figure 4: Distribution of workers' response to "What is your age?"



Figure 5: Distribution of workers' response to "Race Identification (select one or more)"



Figure 6: Distribution of workers' response to "Are you of Hispanic or Latino/a origin?"



Figure 7: Distribution of workers' response to "What is your gender?"



Figure 8: Upper-bound of the error based on calculation.

717dimension and get one accumulated score for each.718For overall creativity, we keep the raw score (an719integer number between 1 and 5). Thus each an-720notation data point consists of three integer scores,721corresponding to atypicality, originality, and over-722all creativity.

723

725

726

727

728

730

732

733

734

737

740

741

742

743

745

746

747

748

751

753

755

756

758

We then normalize the score by individual annotators to mitigate the differences in people's rating preferences. In particular, for each score dimension, we group the scores provided by each annotator and standardize them (subtract mean and divide by standard deviation). We then map the standardized score to an integer (1, 2, or 3) by dividing the standardized score interval into three bins.

B.4 Connection between atypicality and creativity

After analyzing the fine-grained creativity data we collected (Sec. 3.1), we find out that the Pearson R correlation between the normalized atypicality and overall creativity score is 0.4017 (p < 0.01), a positive correlation (the sample size is 2500: 100 ads with 25 annotations each). Therefore, it makes sense to evaluate the same methodology on data with only atypicality annotation to prove its effectiveness at a larger scale.

C Experiment Details

C.1 Configurations

- Temperature: 0.75 (for distribution prediction) and 0.01 (for disagreement prediction)
- Max New Token: 256
 - Model Checkpoint
 - GPT-4: gpt-4-vision-preview
 - LLaVa-7B:
 - llava-v1.6-mistral-7b-hf
 - LLaVa-13B: llava-v1.6-vicuna-13b-hf
 - InternVL2-8B: InternVL2-8B-MPO
- Number of pairwise samples (% of label "1")
- creativity: 938 (46%)
 - atypicality: 2631 (51%)
 - originality: 2708 (50%)
- Device: single card NVIDIA L40S GPU

C.2 Running Time

• Creative-100

(Roughly, all tasks combined) 760

759

– GPT4-v: 1 hr	762
– LLaVA-7B: 4 hr	763
– LLaVA-13B: 5 hr	764
– InternVL: 4.5 hr	765
	100
• Atypical-300 (atypical data only)	766
– GPT4-v: 0.2 hr	767
– LLaVA7B: 0.5 hr	768
– LLaVA13B: 0.5 hr	769
– InternVL: 0.5 hr	770
C.3 Multiple Run Results	771
For the distribution modeling task, since there is	779
inherent randomness in the modeling process (we	773
set the temperature for LLM generation to 0.75)	774
we run the same setting for 3 times to see if the	775
randomness will have a huge impact. As shown	776
in Table 3, for the same task, same model, the	777
performance is still within reasonable margin. We	778
believe it is safe to believe that the results is not	779
affected by randomness in LLM.	780
C 4 VI M Prompts - Creativity	781
Distribution Madeline II	701
Distribution Modeling How creative is this da-	782
auidelines:	783
Choose your answer from [not creative neutral]	704
very creativel answer 1 for not creative 2 for neu-	786
tral and 3 for very creative	787
- Show your reasons and give a final answer (in a	788
single number), in the following format: "reason-	789
ing:{reasoning}, answer: {score}"	790
Disagreement You are a marketing analyst	791
tasked with predicting the audience's perception	792
of an advertisement's creativity. In next week, 25	793
individuals will rate this ad's creativity on a scale	794
of 1 to 3. Your goal is to predict the level of dis-	795
agreement among these ratings. As a reference, ads	796
that are more challenging to understand generally	797
result in higher disagreement, while simpler ads	798
lead to greater consensus.	799
Provide your assessment using a single score:	800
1: Low disagreement (high agreement),	801
2: Neutral disagreement,	802
3: High disagreement (low agreement).	803
Your response should follow this format:	804
'answer: {score}: explanation: {reasoning}"	805

Dimension	Model	Distribution Modeling		
Dimension	Model	R (p-value)	KL-Divergence	
	LLaVA-7B	0.6363 (<.01)	1.0721	
	LLaVA-7B	0.6548 (<.01)	1.0127	
	LLaVA-7B	0.6392 (<.01)	1.0464	
	LLaVA-13B	0.6548 (<.01)	0.3734	
Creativity	LLaVA-13B	0.6289 (<.01)	0.3933	
	LLaVA-13B	0.6257 (<.01)	0.3706	
	InternVL2-8B	0.6363 (<.01)	1.4903	
	InternVL2-8B	0.6548 (<.01)	1.4542	
	InternVL2-8B	0.7153 (<.01)	1.4658	
	LLaVA-7B	0.7573 (<.01)	0.5649	
	LLaVA-7B	0.7623 (<.01)	0.5946	
	LLaVA-7B	0.7781 (<.01)	0.5942	
	LLaVA-13B	0.7108 (<.01)	0.2983	
Originality	LLaVA-13B	0.6757 (<.01)	0.3297	
- 8 - 1	LLaVA-13B	0.7262 (<.01)	0.3448	
	InternVI 2-8B	0.8169 (< 01)	0 5707	
	InternVL 2-8B	0.0109((<.01))	0.6203	
	InternVL2-8B	0.8124 (<.01)	0.5935	
	LLaVA-7B	0.7282 (<.01)	0.4226	
	LLaVA-7B	0.7470 (<.01)	0.4043	
	LLaVA-7B	0.7061 (<.01)	0.4301	
	LLaVA-13B	0.6332 (<.01)	0.2617	
Atypicality	LLaVA-13B	0.5732 (<.01)	0.2692	
	LLaVA-13B	0.7126 (<.01)	0.2367	
	InternVL2-8B	0.7838 (< 01)	0 4041	
	InternVL2-8B	0.7821 (<.01)	0.4083	
	InternVL2-8B	0.7838 (<.01)	0.4041	

Table 3: Results for multiple runs

Pairwise Here are two images of advertisement. Which one is more creative? Answer 1 for the one on the left and 2 for the one on the right. Give your answer in the following format: "explanation: {reasoning}; answer: {choice}"

C.5 VLM Prompts - Atypicality

809

810

811

812

813

814

815

816

817

818

819

- **Distribution Modeling** How atypical is this advertisement? Give your answer that follows these guidelines:
- Atypical advertisement either connected objects that are usually unrelated, contained unusual connections, or brought unusual items together;
- Choose your answer from [not atypical, neutral, very atypical], answer 1 for not atypical, 2 for neutral, and 3 for very atypical
- Show your reasons and give a final answer (in a single number), in the following format: "reason-ing: {reasoning}, answer: {score}; "

Disagreement You are a marketing expert task
with predicting how audiences perceive the atypi-

cality of an advertisement. Atypicality refers to ads 826 that connect objects typically unrelated, feature 827 unusual connections, or bring unexpected items 828 together. In next week, 25 individuals will rate this 829 ad's atypicality on a scale of 1 to 3. Your task is 830 to predict the level of disagreement among their 831 ratings. As a reference, Ads that are more difficult 832 to interpret tend to generate higher disagreement, 833 whereas more straightforward ads lead to greater 834 agreement. 835 Provide your answer using a single score: 836 1: Low disagreement (high agreement), 837 2: Neutral disagreement, 838 3: High disagreement (low agreement). 839 Format your response as follows: 840 "answer: {score}; explanation: {reasoning}" 841

Pairwise Here are two images of advertisement. Which one is more abnormal and unusual? Answer 1 for the one on the left and 2 for the one on the right. Give your answer in the following format: "explanation: {reasoning}; answer: {choice}" 842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

865

866

867

868

869

870

871

872

873

874

C.6 VLM Prompts - Originality

Distribution Modeling How novel is this advertisement? Give your answer that follows these guidelines:

- Novel advertisement either is out of the ordinary, break away from habit-bound and stereotypical thinking, or is unique;

- Choose your answer from [not novel, neutral, very novel], answer 1 for not novel, 2 for neutral, and 3 for very novel

- Show your reasons and give a final answer (in a single number), in the following format: "reasoning: {reasoning}, answer: {score};"

Disagreement You are a marketing analyst tasked with predicting how audiences perceive the novelty of an advertisement. Novelty refers to ads that are out of the ordinary, break free from stereotypical thinking, or exhibit uniqueness. In next week, 25 individuals will rate this ad's novelty on a scale of 1 to 3. Your goal is to predict the level of disagreement among their ratings. Ads that are harder to interpret typically result in higher disagreement, while clearer ads lead to more agreement.

Provide your assessmen	t using	a single	e score:
------------------------	---------	----------	----------

- 1: Low disagreement (high agreement),
- 2: Neutral disagreement,
- 3: High disagreement (low agreement).

Format your response as follows:
"answer: {score}; explanation: {reasoning}"

Pairwise Here are two images of advertisement.
Which one is more unique compared with other ads
in the same product category? Answer 1 for the one
on the left and 2 for the one on the right. Give your
answer in the following format: : "explanation:
{reasoning}; answer: {choice}"

883

887

891

910

917

3 C.7 Atypical-300 Prompts (atypicality only)

Distribution Modeling How atypical is this advertisement? Give your answer that follows these guidelines:

- Atypical advertisement either connected objects that are usually unrelated, contained unusual connections, or brought unusual items together;

- Choose your answer from [not atypical, neutral, very atypical], answer 0 for not atypical and 1 for very atypical

Show your reasons and give a final answer (in a single number), in the following format: "reason-ing: {reasoning}, answer: {score};"

Disagreement You are a marketing expert analyzing how audiences perceive the atypicality of an advertisement. Atypicality refers to ads that con-898 nect objects typically unrelated, feature unusual connections, or bring unexpected items together. 900 Imagine 25 individuals have rated the ad's atypical-901 ity either 0 or 1. Your task is to determine the level 902 of disagreement among their ratings. Ads that are 903 904 more difficult to interpret tend to generate higher disagreement, whereas more straightforward ads 905 lead to greater agreement. 906

- 907 *Provide your answer using a single score:*
- 908 0: Low disagreement (high agreement),
- 909 1: High disagreement (low agreement).
 - Format your response as follows:
- 911 "answer: {score}; explanation: {reasoning}"

912PairwiseHere are two images of advertisement.913Which one is more abnormal and unusual? 1 for914the left image and 2 for the right image. Give915your answer in the following format: "answer:916{number}; explanation: {reasoning}"

D Output Analysis

918 D.1 Error Analysis on Pairwise Outputs

Here, we present two error analysis examples on
the pairwise task. Detailed analysis are in the caption of each image pair (Figure 9, 10).

D.2 Distribution Modeling Examples

We have three examples with all the scoring met-923 rics; see Figure 11, 12, 13. We have also plotted 924 WordClouds for the *reasoning* part of the output 925 (Figure 16, 19, 22). As shown in those word clouds, 926 the commonly used phrases generally closely corre-927 spond to the task definition: 'creative', 'message', 928 'imagery', 'effective' for creativity task, 'unique', 929 'novel', 'habit' for originality task, and 'atypical', 930 'connection', 'unusual' for atypicality task). 931





Figure 9: Task: atypicality; model outputs are in Table 4

Model	Model Output	Prediction
LLaVA-7B	explanation: The image on the left features a man sitting on a chair with a cup of soda in front of him. The text on the ad is about the risks of consuming sugary drinks and the importance of cutting one's portion sizes. This is a typical advertisement that uses a real-life scenario to convey a message. On the other hand, the image on the right is more unusual and abnormal. It features a large, brown, cylindrical structure that resembles a bottle or a vase. The text on this ad is about the risks of consuming sugary drinks and the importance of cutting one's portion sizes. The presence of the large structure in the ad is unexpected and out of place, making it more unusual and abnormal compared to the first image. answer: 2	right (2)
LLaVA-13B	The image on the left is more abnormal and unusual because it features a person with a significantly larger waistline than the rest of their body, which is not a typical human proportion. The image on the right shows a large, ancient-looking structure, which is not an unusual sight in itself, but it is presented in a way that suggests it is a part of the advertisement, which is not a common practice for advertising.	left (1)
InternVL2-8B	explanation: The left advertisement features a person with a large belly, which is a common depiction in media to represent obesity. The use of a person in the image is typical in health campaigns to visually convey the message about portion sizes and health risks. The right advertisement, on the other hand, uses the Parthenon in Athens, an ancient monument, to convey the message about portion sizes. Using a historical or cultural landmark in an advertisement is less common and more unusual compared to using a person. Therefore, the right advertisement is more abnormal and unusual. answer: 2 <i>{repeating answer 4 times}</i>	right (2)
GPT-4v	explanation: The advertisement on the right is more abnormal and unusual because it portrays a well-known architectural column, which is typically straight, as a bottle of Absolut Vodka. This is a surreal and artistic reinterpretation of a cultural landmark, intended to resemble the shape of the iconic Absolut bottle, which is not a normal sight in real-world settings. On the other hand, the advertisement on the left, while addressing a serious health issue regarding portion sizes and diabetes, doesn't alter common objects in an unexpected way but uses real-life elements (a person and soft drink cups) to convey its message. Therefore, the right image is more surreal and unconventional in its approach. answer: 2	right (2)

Table 4: Model output for ads pair in Figure 9. The correct answer is right (2). We can see LLaVA-13B is the only model that gets the answer incorrectly and its explanation is also self-contradicting (the bold part); for InternVL2, the problem is that it always repeat the answer for multiple times, even when it answers correctly.



Figure 10: Task: originality; model outputs are in Table 5

Model	Model Output	Prediction
LLaVA-7B	explanation: The left image features a woman with a caption that reads "When I sneeze, I wet my pants. Normal?" This ad is unique because it addresses a common but often embarrassing issue, using humor to connect with the audience . The right image, on the other hand, shows a group of women in a bathing suit, with a caption that reads: Seh die Schönheit! Seh die Schönhe	left (1)
LLaVA-13B	The left image is more unique compared to the right one. The reason is that the left image features a modern woman with a contemporary hairstyle and clothing, while the right image is an old-fashioned advertisement with a vintage design and illustration style. The juxtaposition of the modern woman with the vintage advertisement creates a contrast that makes the left image more unique.	left (1)
InternVL2-8B	explanation: The left advertisement is more unique compared to the right one because it uses a modern approach with a relatable and humorous question about a common pregnancy symptom, which can engage and resonate with potential customers. The use of a smiling woman and a direct question makes it more approachable and relevant to current audiences. In contrast, the right advertisement uses a vintage style and focuses on a specific product, which might not immediately connect with modern viewers as well. The humor and relatability in the left ad make it stand out in the same product category . Answer: 1 { <i>repeating answer 4 times</i> }	left (1)
GPT-4v	explanation: The advertisement on the left provides a candid, contemporary approach to discussing health issues, specifically addressing pregnancy-related concerns in a way that openly discusses a commonly experienced, yet often unspoken issue. This upfront and direct style can be considered unique in comparison to typical pregnancy services ads which may traditionally use softer language and imagery. On the right, the Nivea advertisement takes a more classic approach, using vintage-style artwork to represent the product in a historical context, which is creative but not necessarily unique as it aligns with a familiar tradition of period-specific advertising. This style, though distinctive, may not stand out as much in the skincare product category where vintage and nostalgic themes are frequently employed. answer: 1	left (1)

Table 5: Model output for ads pair in Figure 10. The correct answer is right (2). ALLaVA-13B is the only one that incorrectly combined the visual element in the left image ("modern woman") and the right image ("vintage advertisement") together as if they are in the same ad ("the juxtaposition of..."); LLaVA-7B and InternVL2-8B are able to capture the humor in the left ad but still make the mistake. In particular, LLaVA-7B is stuck in consistently generating "Seh die Schönheit!", which is a description of the left image, without even get a chance to get to the second image. This might be the reason for its failure. Note that GPT-4v is the only model that recognizes the brand name ("Nivea"), highlighting its strong capability in recognizing details.



Figure 11: Example (A) and creativity predictions by models; complete output in Table 6

Aspect	Human	LLaVA-7B (KL)	LLaVA-13B (KL)	InternVL (KL)
Creativity	2.60	3.00 (1.0054)	2.76 (0.3986)	3.00 (1.0054)
Originality	2.92	2.92 (0.9643)	2.88 (0.3144)	3.00 (1.0054)
Atypicality	2.92	2.88 (0.6030)	2.64 (0.1191)	2.84 (0.2223)

Table 6: Model output (average across 25 runs) and human ratings for Example (A), see ad image and distribution modeling result in Figure 11; *KL* refers to KL(Human||Model)



Figure 12: Example (B) and creativity predictions by models; complete output in Table 7

Aspect	Human	LLaVA-7B (KL)	LLaVA-13B (KL)	InternVL (KL)
Creativity	2.60	2.52 (0.7701)	2.60 (0.8803)	2.88 (1.6395)
Originality	2.92	2.28 (0.2762)	1.60 (0.0963)	2.12 (0.1791)
Atypicality	2.92	1.76 (0.2560)	1.68 (0.070)	1.36 (0.5113)

Table 7: Model output and human ratings for Example (B), see ad image and distribution modeling result in Figure 12; *KL* refers to KL(Human||Model)



Figure 13: Example (C) and creativity predictions by models; complete output in Table 8

Aspect	Human	LLaVA-7B (KL)	LLaVA-13B (KL)	InternVL (KL)
Creativity	2.60	2.64 (1.3395)	2.20 (0.4060)	2.00 (0.3070)
Originality	2.92	2.36 (0.7893)	1.56 (0.0643)	1.60 (0.3263)
Atypicality	2.92	1.80 (0.6703)	1.60 (0.5048)	1.16 (0.8899)

Table 8: Model output and human ratings for Example (C), see ad image and distribution modeling result in Figure 13; *KL* refers to KL(Human||Model)



Figure 14: Task: Creativity, Model: LLaVA-7B



Figure 15: Task: Creativity, Model: LLaVA-13B



Figure 16: Task: Creativity, Model: InternVL2-8B



Figure 17: Task: Originality, Model: LLaVA-7B



Figure 18: Task: Originality, Model: LLaVA-13B



Figure 19: Task: Originality, Model: InternVL2-8B



Figure 20: Task: Atypicality, Model: LLaVA-7B



Figure 21: Task: Atypicality, Model: LLaVA-13B



Figure 22: Task: Atypicality, Model: InternVL2-8B