

---

# Large Language Models are Fixated by Red Herrings: Exploring Creative Problem Solving and Einstellung Effect using the Only Connect Wall Dataset

---

Saeid Alavi Naeini<sup>1,3,4</sup> Raeid Saqur<sup>2,4</sup> Mozhgan Saeidi<sup>2,4,6</sup> John Giorgi<sup>2,4,5</sup> Babak Taati<sup>1,2,3,4</sup>

<sup>1</sup>Kite Research Institute, Toronto Rehabilitation Institute, University Health Network

<sup>2</sup>Department of Computer Science, University of Toronto

<sup>3</sup>Institute of Biomedical Engineering, University of Toronto <sup>4</sup>Vector Institute for AI

<sup>5</sup>Donnelly Centre for Cellular & Biomolecular Research, University of Toronto

<sup>6</sup>Department of Biomedical Data Science, Stanford University

{saeid.alavi, john.giorgi}@mail.utoronto.ca

raeidsaqur@cs.toronto.edu, mozhgans@stanford.edu, babak.taati@uhn.ca

## Abstract

The quest for human imitative AI has been an enduring topic in AI research since its inception. The technical evolution and emerging capabilities of the latest cohort of large language models (LLMs) have reinvigorated the subject beyond academia to the cultural zeitgeist. While recent NLP evaluation benchmark tasks test some aspects of human-imitative behavior (e.g., BIG-bench’s ‘human-like behavior’ tasks), few, if not none, examine *creative problem solving* abilities. Creative problem solving in humans is a well-studied topic in cognitive neuroscience with standardized tests that predominantly use the ability to associate (heterogeneous) connections among clue words as a metric for creativity. Exposure to misleading stimuli — distractors dubbed *red herrings* — impede human performance in such tasks via the *fixation effect* and *Einstellung* paradigm. In cognitive neuroscience studies, such fixations are experimentally induced by pre-exposing participants to orthographically similar incorrect words to subsequent word-fragments or clues. The popular British quiz show Only Connect’s *Connecting Wall* segment essentially mimics Mednick’s Remote Associates Test (RAT) formulation with built-in, deliberate red herrings, which makes it an ideal proxy task to explore and study the fixation effect and *Einstellung* paradigm from cognitive neuroscience in LLMs. In this paper, we present the novel Only Connect Wall (OCW) dataset and report results from our evaluation of selected pre-trained language models and LLMs on creative problem solving tasks like grouping clue words by heterogeneous connections and identifying correct open knowledge domain connections in respective groups. We synthetically generate two additional datasets: OCW-Randomized, OCW-WordNet to further analyze our red-herrings hypothesis in language models. The code and link to the dataset are available at <https://github.com/TaatiTeam/OCW>.

## 1 Introduction

The remarkable capabilities of state-of-the-art large language models (LLMs) [91], across a variety of domains and downstream tasks [78, 10], have spurred their comparisons with artificial general intelligence (AGI) [5, 14] and human-imitative AI [31] systems. The extraordinary leap in capabilities of these LLMs over a short span — from the advent of transformer-based [69] pre-trained, context-aware language models (PLMs) [52, 17, 40, 36, 53] circa 2018 to 2020, to the current and latest

Wall A: Season 11, Episode 23					Wall B: Season 12, Episode 27				
Gala	Twelfth	Bonfire	Hen	... night	Jazz	Gala	Honeygold	Jonathan	Apples
Oriov	Churchill	Digby	Tony	Advert Animals	Healing	Join	Greg	Show of	Can Precede Hands
Burns	Marx	Clarke	Bender	Cigar Smokers	Pippin	Merry	Gaffer	Sam	Hobbits
Canal Street	Castro	Chelsea	Darlinghurst	Gay Villages	Twill	Duct	Ticker	Cassette	Types of Tape
Wall C: Season 15, Episode 10					Wall D: Season 10, Episode 2				
Cameo	Fuji	Bramley	Jazz	Apples	Shrewsbury	Wellington	Ludlow	Madeley	Shropshire Towns
Amy	Lady Bird	Dakota	Dwayne	Johnsons	Bath	Boarding	Doge	Cathode	Begin with Animals
Thunder	Magic	Heat	Celtics	US Basketball Teams	Chelsea	Gum	Snow	Cowboy	Boots
Gala	Costume	Goggles	Pool	Swimming	Bolt	Bond	Churchill	Coward	English Playwrights

Figure 1: Examples of *Only Connect* walls with ground-truth groupings (rows) and connections (last column). *Red herrings* include orthographically identical words, e.g., **Gala**, **Churchill** and **Chelsea** in different connected groups — **Gala**: *Gala night*, *Apples*, *Swimming gala*, **Churchill**: *Advert Animals*, *English Playwrights* and **Chelsea**: *Gay Villages*, *Boots* — across walls. In Wall A (top left), the clues **Churchill**, **Marx**, **Castro** provide misleading stimuli inducing plausible fixation on historical figures within the wall.

cohort of increasingly larger (billions of parameters) LMs [59, 77, 57, 89, 16, 67, 18] spearheaded by the OpenAI’s GPT series [13], notably ChatGPT [49] and GPT-4 [48] — justifiably warrants such comparisons. Several natural language processing (NLP) benchmarks have been proposed to standardize the evaluation of these LLMs, including MMLU [27], BIG-bench [66], HELM [38], and Global-Bench [65]. The tasks inventory under these benchmarks are open (type of tasks) and dynamic (rolling additions). While a subset of these tasks aims to test for human imitative intelligence (e.g., nineteen tasks listed under the *human-like behaviour* category in BIG-bench), none tests for *creative problem solving* abilities [44] — a hallmark of human-like intelligence [31].

Creative problem-solving by humans is a well-studied topic in cognitive neuroscience and human behavioural sciences literature. These studies and methods use (word) associative fluency to model and test creativity objectively [44, 9]. Empirical research in this context commonly employs single or continuous word association tests that are variants of Mednick’s seminal Remote Associates Test (RAT) [45]. Such tests entail finding connections or links among a presented group of words using associations that can be heterogeneous (e.g., synonymy, semantic, compounding) [86, 43]. To exemplify, consider the cue words: *{Tennis, Same, Head}*. A correct connection in this triplet is: *Match*, which connects by semantic link (*tennis match*), synonymy (*same match*), and compounding (*match head*). Further, the word connections can also vary in degrees of figurativeness (e.g., *Star-Actress* vs. *Star-Planet*) and abstractness (e.g., *Humor-Sense* vs. *Apple-Tree*). In humans, such creative problem-solving abilities are impeded by exposure to wrong answers [61, 62, 85] — a finding referred to as the *fixation effect* [34, 82]. A closely related similar concept is the *Einstellung effect* [42], which postulates the negative effect of previous experience when solving new problems.

Studies examining the fixation effect induce fixations by presenting clue words intended as wrong answers (misleading stimuli) [61] dubbed “red herrings” or, by pre-exposing participants to red herrings before attempting creative problem-solving tasks like the RAT [45]. A slew of works in negative transfer learning in human cognition attempt to explain the RAT fixation phenomenon that involves pre-exposure to red herrings by the negative effects of prior learning on indirect or implicit measures of memory [63]. This negative transfer effect was demonstrated and studied using orthographically similar words to subsequent test word fragments as red herrings [63]. Intuitively, the red herrings lead participants away from the memory retrieval (or down incorrect neurological pathways by Hebbian terminology) required for correct responses and fixate on wrong connections [60]. Fixation in creative problem-solving can be increased by making red herrings more retrievable. Thus, creative problem-solving can be thought of as a type of indirect memory measure whose retrieval is degraded by red herrings due to the negative transfer effect. The *red herring retrieval hypothesis* states that factors that make red herrings more retrievable should reduce creative problem-solving performance,

as measured with RAT problems. Two such factors are repetition and context. A following corollary states that the memory strengths of red herrings determine the magnitude of a fixation effect [8].

In this work, we study the juxtaposition of these theories from human cognitive neuroscience (*fixations, negative transfer learning, red herring memory retrieval hypothesis*) from the context of LLMs and natural language processing. While negative transfer learning has been observed and studied in AI research [76, 22], the context of these studies is limited to strict machine learning sub-domains like statistical distribution measures and computer vision. There has not been any work that systematically examines these specific concepts’ relation in AI research. Our major contributions are as follow:

**1. Only Connect Wall (OCW) Dataset and creative problem solving tasks.** We introduce a novel dataset for evaluating *creative problem solving* tasks by curating the problems and human performance results from the popular British quiz show Only Connect [81, 3]. Specifically, the *Connecting Walls* segment of the show, where the tasks entail grouping sixteen (16) jumbled up clue words into four (4) connected groups, and naming the correct connections (Figure 1). The presented words have heterogeneous connections with open-domain knowledge retrieval, e.g., history, places, famous people, tools, and cultural references. These ‘walls’ contain red herrings or misleading stimuli by design, which makes this dataset an analogical proxy for RAT tests in evaluating LLMs for creative problem-solving. Section §2 provides a detailed description of the dataset.

**2. Experiments, results, and key findings of baseline LLMs evaluation.** We evaluate a suite of NLP models from static embeddings to PLMs to LLMs and demonstrate that none can solve the tasks of the OCW dataset. Our findings show that SOTA LLMs (e.g. GPT-4 [48]) perform significantly worse than the expert human baseline, and somewhat surprisingly, that increasing the number of in-context examples in few-shot in-context-learning is ineffective. Sections §3 and §4 provide details.

## 2 Only Connect Walls Dataset

Here we focus on the *Connecting Walls* segment (usually the third round) of the quiz-show. Each wall contains sixteen jumbled-up word clues that must be sorted into four groups, each with four connected words. Once the groups are formed, contestants must also identify the right connection or relationship among the items in each group. While there is only one correct solution to each wall, the puzzles are designed to include several red herring clues that can fit into another category and red herring categories fitting multiple clues. Figure 1 shows solved sample walls from the show highlighting a couple of typical red herrings.

### 2.1 Dataset Collection and Structure

The OCW dataset contains 618 connecting wall puzzles and solutions in total from 15 seasons of the show. Each show episode has two walls. The total number of walls per season varies based on the (varying) number of aired season episodes. The walls were scraped from fan websites<sup>1</sup>, and human performance results (for grouping and connection tasks) were manually curated by watching all the episodes. Figure 2 depicts the high-level structure of the dataset in JSON format with self-explanatory object keys and comments.

### 2.2 Tasks and Evaluation Metrics

The two dataset tasks: **Task 1 (Grouping)**, and **Task 2 (Connections)** are identical to the quiz-show’s human participant tasks. We evaluate Task 1 (Groupings) via six metrics: number of solved walls, number of correct groups (max. four per wall), Adjusted Mutual Information (AMI) [71], Adjusted Rand Index (ARI) [28], Fowlkes Mallows Score (FMS) [21], and Wasserstein Distance (WD) [54], normalized to (0, 1) range, between predicted and ground-truth labels [88, 70].

We similarly evaluate Task 2 (Connections) with three metrics: exact string matching, ROUGE-1 F1 [39], and BERTScore F1 [90]. Exact match is the most strict, assigning a score of 1 when the predicted connection is identical to the ground-truth and 0 otherwise. ROUGE-1 F1 relaxes this criterion; it is large when there is a high proportion of ground-truth tokens in the model’s predicted

---

<sup>1</sup>The primary source was the Only Connect fan website: <https://ocdb.cc> [6].

Predicted Connection	Ground-truth Connection	Exact Match	ROUGE-1 F1	BERTScore F1
Types of numbers	Types of numbers	1.00	1.00	1.00
Slang terms for money	Slang for money	0.00	0.86	0.79
Types of trees	Trees	0.00	0.50	0.63
Bridges in London	Thames bridges	0.00	0.40	0.31
Medieval occupations	Chaucer characters	0.00	0.00	0.15

Table 1: Examples of predicted and ground-truth connections and their performance according to the chosen metrics. Exact match is 0 for anything but identical strings. Empirically, we observe that a ROUGE-1/BERTScore F1 of  $\geq 0.5$  indicates that a predicted connection is likely *correct*.

connection *and* a low proportion of non-ground truth tokens. BERTScore F1 is similar but further relaxes this criterion, assigning a non-zero score for *semantically* similar (but non-identical) predicted tokens. Together these three metrics provide a more holistic view of model performance on Task 2 than any one metric alone. Empirically, we find that a ROUGE-1 or BERTScore F1 of  $\geq 0.5$  indicates that a predicted connection would likely be considered *correct* (Table 1). Note that BERTScore has many parameters affecting the final score; a hashcode is produced and reported for reproducibility.

Each of the evaluation metrics for Task 1 of Task 2 could be calculated per wall, per episode, per season, or for the entire test set. We present results on the entire test set in this paper (§4). We split the dataset into a train set (62 walls), validation set (62 walls), and test set (494 walls). The primary goal of our dataset is to evaluate the zero- and few-shot creative problem-solving abilities of LLMs; as such, we elect to set the size of the test set to be much greater than train or validation sets.

### 3 Experiments: Language Model Evaluations

This section describes methods and models used to provide baseline results for the dataset. For Task 1 (Grouping), we use clustering techniques on word-embeddings from classical and pre-trained language models (PLMs) (§3.1), and few-shot in-context learning (ICL) with LLMs (§3.2). For Task 2 (Connections), we only provide baseline results using few-shot ICL with LLMs (§3.2).

```
{// Contains a mapping of number of walls per season
"season_to_walls_map": {"1": 30, "2": 16, ...},
// Contains the actual dataset, one object per wall
"dataset": [{
  // Each wall has a unique identifier, season and episode number
  "wall_id": "3ca8", "season": 1, "episode": 8,
  // The list of 16 words or "clues" associated with the wall
  "words": ["Holmes", "Indiana Jones", "Bannister", ...],
  // Ground-truth connections for each of the four groups of this wall
  "gt_connections": ["Parts of a staircase", "___ cake", ...],
  // The four ground-truth groups for this wall
  "groups": {
    // Each group is an object with...
    "group_1": {
      // ...a unique ID
      "group_id": "3ca8_01",
      // ...ground-truth words
      "gt_words": ["Newel", "Bannister", "Tread", "Riser"],
      // ...and ground-truth connection
      "gt_connection": "Parts of a staircase",
      // Human performance is recorded as solved (1) or unsolved (0)
      "human_performance": {"grouping": 1, "connection": 1}
    }, ...
  }
  // Overall human performance for each group within the wall
  "overall_human_performance": {
    "grouping": [1, 1, 1, 1],
    "connections": [1, 1, 1, 0],
  }
} ...]}
```

Figure 2: JSON Structure of the OCW dataset. One truncated example is shown.

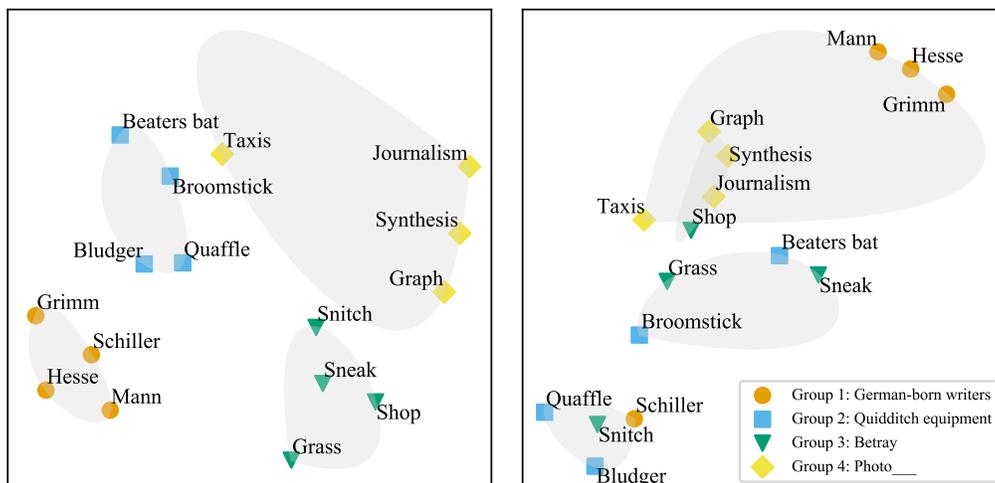


Figure 3: Solved wall (wall\_id="8cde") for Task 1 (Grouping) using best performing model ( $E5_{BASE}$ ) with both static and contextual embeddings. **Left:** solved wall using static embeddings. **Right:** unsolved wall using contextual embeddings. 2D projection of embeddings using t-SNE is shown. Colors and shapes correspond to true clusters, and grey convex regions correspond to predicted clusters. The legend shows the ground truth connection for each group.

### 3.1 Task 1: Grouping using Word Embeddings

For the *grouping task* evaluation (§2.2), we use clustering algorithms on word-embeddings of the sixteen clue words in each wall, to group them into four predicted groups that are subsequently evaluated against the four ground-truth groups for each wall. A vanilla  $k$ -means (with  $k = 4$ ) clustering algorithm [25] does not guarantee each predicted group to have four words, thus we use variants like constrained clustering.

**Clustering** Semi-supervised constrained clustering [72, 7] is used when the user has pre-existing knowledge about the desired partition (in our case, 4 groups). Here, we adopt a *minimum cost flow network* clustering approach [12] with a cluster size of four for grouping. Our preliminary analysis showed that clustering results exhibited slight variations across runs. This slight discrepancy could be attributed to the initializations of cluster centroids. To address this issue and ensure reliable results, we report the mean and variance of results (Table 3) across sixteen (16) runs, each with a unique seed and randomized order of sixteen-word clues. We tested two additional clustering approaches motivated by [47, 19]: (1) We constructed a self-similarity matrix containing pair-wise similar information about the words prior to applying constrained clustering; (2) We performed dimensionality reduction using Principal Component Analysis (PCA) [58] and t-distributed stochastic neighbor embedding (t-SNE) [68] before applying constrained clustering. Neither approach improved performance over raw embeddings’ clusters, and, for brevity, results are not included.

**Static word embedding** We used two well-known classic word embedding models, GloVe [51] and FastText [23], both of which are accessed through the Flair library (Table 2). We used two FastText models, one pre-trained on the Common Crawl corpus and another on Wikipedia. Approximately 10% of the total clues encountered in the dataset were out-of-vocabulary (OOV). A significant portion (~80%) of the OOV cases were addressed by mean pooling for clues comprised of multiple words to obtain one unified embedding. For the remaining OOV instances, we combined the static embeddings with BytePair encoded[26] sub-words.

**PLMs** We explored general-purpose PLMs (BERT [17], RoBERTa [40], DistilBERT [56], ELMo [52]) as well as Sentence Transformers (MPNet [64], E5 [75]; see Table 2). We evaluated performance with and without contextual embeddings.<sup>2</sup> Depending on the context, some clues in the dataset may appear across different walls with different meanings. As an example, the word

<sup>2</sup>Static embeddings are obtained from the PLMs by passing clues through the model *independently*.

Model	# Parameters		Version	Accessed via
<i>Word Embeddings</i>				
BPEmb [26]	En	–	en	Flair [4]
GloVe [51]	6B	–	glove	Flair
FastText [23]	Crawl	–	crawl	Flair
	News	–	news	Flair
<i>Pre-trained Language Models (PLMs)</i>				
ELMo <sub>LARGE</sub> [52]		94M	large	Flair [4]
DistilBERT <sub>BASE</sub> [56]	uncased	66M	distilbert-base-uncased	HuggingFace [83]
BERT <sub>BASE</sub> [17]	uncased	110M	bert-base-uncased	HuggingFace
BERT <sub>LARGE</sub>	uncased	340M	bert-large-uncased	HuggingFace
RoBERTa <sub>LARGE</sub> [40]		355M	roberta-large	HuggingFace
<i>Sentence Transformers</i>				
all-mpnet <sub>BASE</sub> [64]	V2	110M	sentence-transformers/all-mpnet-base-v2	HuggingFace
E5 <sub>BASE</sub> [75]	V2	110M	intfloat/e5-base-v2	HuggingFace
E5 <sub>LARGE</sub>	V2	340M	intfloat/e5-large-v2	HuggingFace
<i>Large Language Models (LLM)</i>				
GPT-3.5-turbo		–	gpt-3.5-turbo-0301	OpenAI API
GPT-4		–	gpt-4-0314	OpenAI API

Table 2: Details about the baselines and models used in our experiments.

“Gala” was found in three distinct walls, each associated with a different meaning: *apples*, *swimming* \_\_\_\_, and \_\_\_\_ *night* (Figure 1). The contextual embeddings were aimed to capture contextual semantic similarity among the clues (if any). They were generated by joining the 16 clues in the wall as a pseudo-sentence. We randomly shuffle the word order across sixteen different runs for each wall to account for the positional ordering. We note that such faux sentences (for inducing context) are not valid English syntactic sentences. We used mean pooling to generate embeddings for clues comprised of multiple words to capture the collective meaning of the entire clue.

### 3.2 Task 2: Connections using Few-shot In-context Learning (ICL) with LLMs

Few-shot ICL with LLMs has emerged as a performant and broadly applicable paradigm in NLP [13]. To evaluate the performance of this approach on our proposed dataset, we designed a few-shot prompt for GPT-3.5-turbo and GPT-4 [48], which are amongst the strongest performing LLMs currently available.<sup>3</sup> For Task 1 (Grouping, §2.2), the prompt consists of some natural language instructions, several examples of solved walls from the training set, and the current example’s 16 clues, randomly sorted. For Task 2 (Connections), in place of the 16 clues, the prompt contains a solved wall *without* the connections (Figure 4).

We developed our prompts on the validation set and reported the final performance on the test set. In-context examples are randomly selected from the train set; the same examples are used across all test inputs. We experiment with 0, 1, 3, 5, and 10 in-context examples. When necessary, we apply simple post-processing to the LLMs output. For example, in both Task 1 and Task 2, we take a maximum of 4 predictions for the groups and connections, respectively, and pad up to 4 with the empty string in cases where the model outputs fewer than 4.<sup>4</sup> To make results as reproducible as possible, we set the temperature=0 and used the 03/01/2023 GPT-3.5-turbo snapshot and the 03/14/2023 GPT-4 snapshot. The max output length is set to 144 tokens. All other hyperparameters of the OpenAI API are left at their defaults [2]. Prompts were designed as per the Guidance library [1].

## 4 Results and Discussions

### 4.1 Task 1: Grouping Results

**Embedding Clustering Techniques** In Table 3 we report the performance of several static embedding baselines on Task 1 (Grouping). E5<sub>BASE</sub> was the most performant model and, on average, solved

<sup>3</sup>In preliminary experiments, we found that open-source LLMs like LLaMA [67] perform poorly and typically do not follow the task instructions.

<sup>4</sup>Please see our codebase for all post-processing steps: <https://github.com/TaatiTeam/OCW>

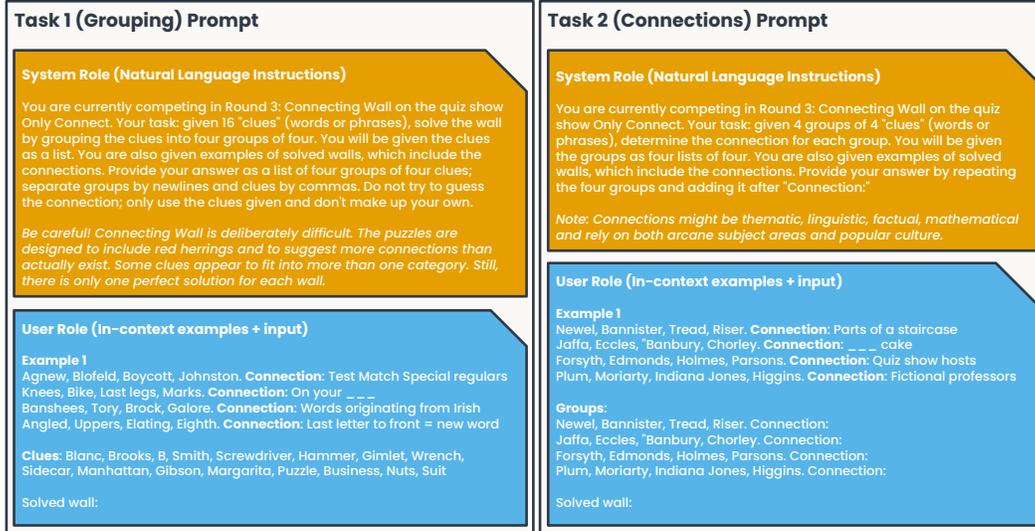


Figure 4: Example prompts for Task 1 (Grouping) and Task 2 (Connections) used with GPT-3.5-turbo and GPT-4. The system’s role includes natural language instructions. The user role includes  $n$  in-context examples and the current examples 16 clues (Task 1) or the solved wall without connections (Task 2). For Task 1, the model is instructed to output the solved wall as four lines of four clues separated by commas. For Task 2, the model is instructed to copy the solved wall and fill in the connections. *Emphasis* and **bold text** are for visualization purposes only.

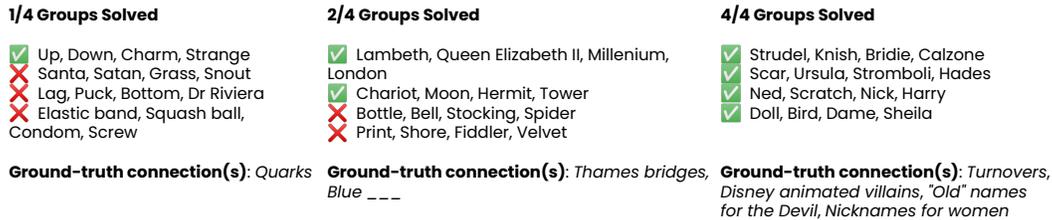


Figure 5: Examples of partially and fully solved walls predicted by GPT-4.

1 wall and correctly clustered 89 groups. Contextual embeddings had the lowest overall performance amongst all methods (Table 6 in Appendix §A). One explanation is that by concatenating all the clues in each wall, the resulting input may not adhere to the sentence structure that PLMs are accustomed to during training. This may disrupt the natural flow of information and, in turn, lead to less meaningful contextual embeddings. Moreover, the context may change abruptly when combining clues from different parts of the wall. This can introduce ambiguity and contextual shifts that the model may struggle to interpret accurately. Another possible explanation is the effect of positional encoding in the underlying models. Unlike other main components of PLMs, positional encoding is variant to sequence order [33]. Even though the addition of positional to word embeddings helps with learning the contextual representation of words at different positions, intrinsic similarities may be more important than contextual usage. The embedding dependence on neighboring clues may have hindered the clustering process by introducing noise and capturing irrelevant information that is specific to a particular context. For instance, in Figure 3, the contextual embedding model erroneously associated the clue “Shop” with the connection “Photo \_\_\_”, resulting in the formation of the word “Photoshop”; however, this association is incorrect as it is an example of a red herring in the wall. In contrast, the static embedding model correctly mapped “Shop” to its British slang meaning connection “Betray”. Please refer to Appendix §B for more examples.

**Few-shot ICL with LLMs** Performance of GPT-4 far surpassed the static (Table 3) and contextual embedding baselines (Table 6), particularly in terms of the number of solved walls and correct groups (>2X the next most performant model, E5), but was still far below human performance (§4; see §5 for example predictions). Examining the predictions of the best-performing model (GPT-4,

	WD ↓	FMS ↑	ARI ↑	AMI ↑	# Solved Walls	# Correct Groups
<i>Classic Word Embeddings</i>						
GloVe	84.9 ± .4	31.5 ± .3	14.4 ± .3	17.6 ± .4	0 ± 0	68 ± 4
FastText (Crawl)	84.2 ± .5	32.1 ± .3	15.2 ± .3	18.4 ± .4	0 ± 0	80 ± 4
FastText (News)	85.5 ± .5	30.4 ± .2	13.0 ± .2	15.8 ± .3	0 ± 0	62 ± 3
<i>Pre-trained Language Models (PLMs)</i>						
ELMo <sub>LARGE</sub>	86.3 ± .6	29.5 ± .3	11.8 ± .4	14.5 ± .4	0 ± 0	55 ± 4
DistilBERT <sub>BASE</sub>	86.7 ± .6	29.1 ± .2	11.3 ± .3	14.0 ± .3	0 ± 0	49 ± 4
BERT <sub>LARGE</sub>	88.3 ± .5	26.5 ± .2	8.2 ± .3	10.3 ± .3	0 ± 0	33 ± 2
BERT <sub>BASE</sub>	89.5 ± .4	25.1 ± .2	6.4 ± .3	8.1 ± .4	0 ± 0	22 ± 2
RoBERTa <sub>LARGE</sub>	88.4 ± .4	26.7 ± .2	8.4 ± .3	9.4 ± .4	0 ± 0	29 ± 3
<i>Sentence Transformers</i>						
all-mpnet <sub>BASE</sub>	86.3 ± .4	29.4 ± .3	11.7 ± .4	14.3 ± .5	0 ± 0	50 ± 4
E5 <sub>LARGE</sub>	84.4 ± .7	32.3 ± .4	15.4 ± .5	18.5 ± .6	0 ± 0	76 ± 5
E5 <sub>BASE</sub>	<b>83.8 ± .6</b>	<b>33.1 ± .3</b>	<b>16.3 ± .4</b>	<b>19.5 ± .4</b>	<b>1 ± 0</b>	<b>89 ± 6</b>
Human Performance	–	–	–	–	285 / 494	1405 / 1976

Table 3: Results of selected models on Task 1 (Grouping) using static embeddings. WD: Wasserstein Distance. FMS: Fowlkes Mallows Score. ARI: Adjusted Rand Index. NMI: Normalized Mutual Information. Mean ± standard deviation over 16 random seeds is shown. **Bold**: best scores.

	# In-context Examples	WD ↓	FMS ↑	ARI ↑	AMI ↑	# Solved Walls	# Correct Groups
GPT-3.5-turbo	0-shot	82.5	34.0	18.4	21.6	0	114
	1-shot	82.3	34.4	18.2	21.2	0	123
	3-shot	80.9	36.8	21.3	24.7	0	140
	5-shot	80.6	37.3	22.0	25.4	2	149
	10-shot	81.2	36.1	20.4	24.0	2	137
GPT-4	0-shot	75.8	41.5	27.2	30.7	6	239
	1-shot	73.4	43.7	29.7	33.5	4	262
	3-shot	73.7	<b>43.9</b>	<b>29.9</b>	<b>33.6</b>	5	<b>272</b>
	5-shot	<b>72.9</b>	43.4	29.1	32.8	<b>7</b>	269
	10-shot	73.6	42.8	28.5	32.3	3	249
Human Performance		–	–	–	–	285 / 494	1405 / 1976

Table 4: Results on Task 1 (Grouping) using Large Language Models. WD: Wasserstein Distance. FMS: Fowlkes Mallows Score. ARI: Adjusted Rand Index. NMI: Normalized Mutual Information. **Bold**: best scores.

3-shot), we found common sources of error to include misformatted outputs (4.4% of all predicted groups) and hallucinated clues (6.6%).

Surprisingly, more in-context examples (from 1 to 10 shot) did not improve performance. One possible explanation for this observation is that, due to the huge variety of possible connection types, the in-context examples’ primary benefit is demonstrating the expected output format – as opposed to demonstrating how to perform the task – which likely requires only a single example. This is related to the concepts of *task learning* versus *task recognition*, which are thought to be the two distinct mechanisms through which ICL leverages demonstrations [50, 32]. Many clues require open-domain, arcane, cultural and intimate knowledge of niche subject areas (e.g., “*Professional snooker players*”, “*Female Radio 1 DJs*”) that, without prior memorization, are unlikely to help. The presence of orthographically similar clue words in the in-context examples could themselves act as red herrings and plausibly induce negative transfer learning. An interesting future direction would be the evaluation of retrieval augmented models [24, 37, 11, 29], which may be capable of solving groups about highly specific subject areas.

## 4.2 Task 2: Connections Results

In Figure 6, we present the results for Task 2 (Connections). In general, GPT-4 outperforms GPT-3.5-turbo, especially in the 0-shot regime. Performance for GPT-4 improves monotonically with an increasing number of in-context examples, although improvements are sometimes small (e.g.,

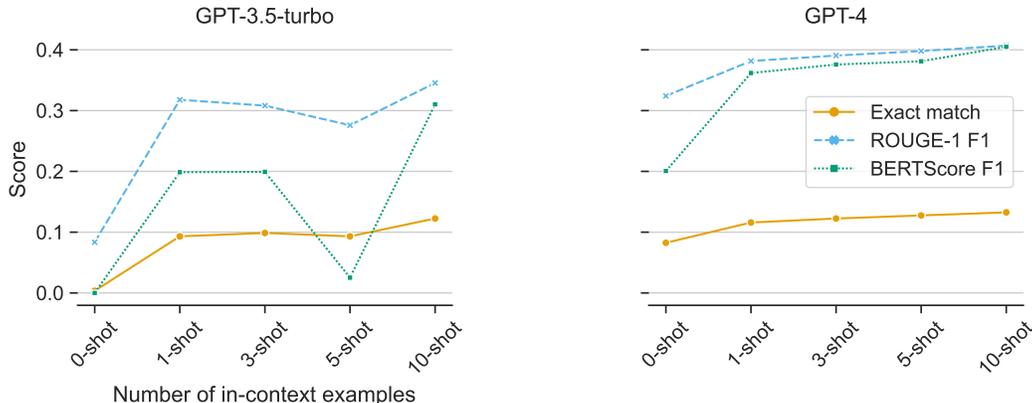


Figure 6: Results for Task 2 (Connections) with GPT-3.5-turbo and GPT-4. For reference, human performance is approximately 80% (fraction of correctly answered connections). We report  $\max(\text{BERTScore}, 0)$  in the case of GPT-3.5-turbo for readability.

$< 0.01$ ). As expected, the exact match score for both models is low ( $< 15\%$ ). This is explained by the fact that even insignificant differences between the model’s predictions and the ground truth will result in a score of 0 (e.g., “Made *of* rubber” vs. “Made *from* rubber”). For this reason, we also report ROUGE-1 and BERTScore F1 scores (§2.2). Although not a perfect comparison, we can contextualize these results with human performance, which we recorded as the fraction of correctly guessed connections:  $\sim 80\%$  on the test set. The quiz show Only Connect allows for some small deviations in guessed connections that will be accepted as correct, making the comparison to ROUGE and BERTScore more suitable than to exact match. Our results suggest that at 41-45% F1, the best performance achieved with few-shot ICL (GPT-4, 10-shot) is far below human performance. Lastly, we note that a common source of model error was the inclusion of clues in the predicted connection (occurring in 8.2% of all predicted connections for the best performing model), e.g., “Fireplace tools (Spade, Brush, Poker, Tongs)”, even though (1) the model was not instructed to do so, and (2) the in-context examples were not formatted like this.

More complicated post-processing or prompting strategies (e.g., “Chain of Thought” [79], “Tree of Thoughts” [87]) could mitigate these issues and improve performance. However, applying these more complicated prompting strategies to the OCW dataset is non-trivial, as they require breaking down the problem into intermediate steps, and the number or nature these intermediate reasoning steps should take is unclear. We leave their application to the OCW dataset for future work.

### 4.3 Effects of Red-Herrings: Additional Datasets, Experiments and Analyses

To analyze our *red-herring hypothesis* on language models, we designed and performed additional ablative experiments. The original OCW dataset contains red-herrings as distractors *by design*. We generate two additional datasets from OCW to decrease the presence of red-herrings: OCW-Randomized and OCW-WordNet. The goals, construction and other details are presented in Appendix §C.1.

In OCW-Randomized, we diluted the presence of red herrings by randomly swapping groups among the walls in the test set – thus negating the inherent deliberate distractor groups in each wall. We further simplify the grouping task in OCW-WordNet by removing red herrings altogether. This is achieved by using subordinate-superlative (or hyponym-hypernym) word hierarchy and synonyms in the English lexical database WordNet [46, 20]. Thus the results in Table 5 present results on datasets with a decreasing proportion of red herrings from left to right, and by our hypothesis, increasing task simplicity for LLMs. The results are aligned with our expectations, with GPT-3.5-turbo and GPT-4 performance increasing significantly with the reduction of red herrings from the test set.

## 5 Related Work

Various datasets and tasks have been proposed for evaluating language models against human-like linguistic capabilities. Earlier examples of such tasks include *word sense disambiguation* (WSD) [55],

		OCW		OCW-Randomized		OCW-WordNet	
		# Solved Walls	# Correct Groups	# Solved Walls	# Correct Groups	# Solved Walls	# Correct Groups
GPT-3.5-turbo	0-shot	0	114	5	274	337	1522
	1-shot	0	123	12	315	320	1400
	3-shot	0	140	10	306	415	1748
	5-shot	<b>2</b>	<b>149</b>	16	<b>337</b>	415	1759
	10-shot	2	137	<b>17</b>	333	<b>428</b>	<b>1800</b>
GPT-4	0-shot	6	239	59	595	<b>471</b>	<b>1926</b>
	1-shot	4	262	57	644	304	1581
	3-shot	5	<b>272</b>	62	649	279	1537
	5-shot	<b>7</b>	269	<b>68</b>	<b>655</b>	298	1584
	10-shot	3	249	55	614	378	1742
Human Performance		285 / 494	1405 / 1976	–	–	–	–

Table 5: Coalesced results of LLMs performance on Task 1 (grouping) using two additional test datasets OCW-Randomized and OCW-WordNet with decreasing presence of red-herrings from *left* to *right* in the walls, juxtaposed against the original OCW test set (left-most column). Only the main metrics are shown (details and full results in Appendix §C). **Bold**: best scores.

Winograd schema challenge [35] and *word sense induction* (WSI) [80]. WSD aims to determine a word’s correct meaning or sense within a specific context. WSI focuses on automatically clustering words into different senses or semantic categories based on their contextual usage patterns. Benchmarks like GLUE [74] and SuperGLUE [73] are aimed at aggregating and standardizing these classical NLP tasks to evaluate language models. The PLMs (e.g., BERT variants) and the first generation of LLMs, mostly solved or attained human-level performance on these tasks by 2020s [41].

In order to evaluate the human-imitative capabilities of modern LLMs, more challenging tasks have been proposed in recent benchmarks like BIG-bench [66] and HumanEval [15]. **BIG-bench** aims to address the limitations of existing benchmarks by providing a more comprehensive, open, and dynamic (tasks added on a rolling basis) evaluation benchmark. It covers a wide range of tasks, including a suite of tasks targeted specifically for *human-like behavior*. **HumanEval** is an evaluation set to measure the functional correctness of code synthesis from docstrings [15]. This benchmark includes 164 original programming problems that assess language comprehension, algorithms, and simple mathematics comparable to simple software interview questions. While these recent benchmarks include a wide net of complex tasks, evaluating a broad range of LLM capabilities, our work here is orthogonal to these since none of them aims to specifically measure creative problem-solving or creativity and their impediments in LLMs.

## 6 Limitations & Future Work

As with any machine learning dataset, especially one designed to evaluate the performance of LLMs, the OCW dataset has several limitations. First, we noticed that the performance of contextual approaches can vary significantly depending on the order that clues are provided to the model. To alleviate this (and where feasible), we evaluate models across 16 random sortings of the clues. Due to cost, we did not evaluate GPT-3.5-turbo and GPT-4’s sensitivity to this ordering; future work should report performance across multiple random sorts. Second, due to the nature of the quiz show *Only Connect*, the clues, groups, and connections in the dataset tend to be Western- (and specifically UK-) centric (e.g. “*Doctor Who companions*”, “*English cricket captains*”, “*Irish counties*”). Therefore, performance on the OCW dataset may not extrapolate to languages or cultures outside of Western English. In fact, the *US*-centric bias of LLMs like GPT-3.5/4 [84] might partially explain their poor performance on the *UK*-centric OCW dataset. We hope to add additional *Only Connect* inspired walls in multiple languages and with clues derived from various cultures & subcultures in future work. Finally, given that the walls are publicly available as text on fan sites like ocdb.cc, there is always the possibility that they are included in the training sets of LLMs like GPT. However, we think this is unlikely, given the low performance on the grouping and connection tasks. Preventing the test sets of publicly available datasets like our OCW from “leaking” into the training sets of LLMs remains an interesting and open problem. We have taken basic steps against this leakage by distributing our dataset in a compressed format [30].

## References

- [1] Microsoft guidance: A guidance language for controlling large language models. <https://github.com/microsoft/guidance>. 2023.
- [2] OpenAI API completions reference. <https://platform.openai.com/docs/api-reference/completions>. 2023.
- [3] Only Connect. Television show, 2008–2020. Created by Presentable, RDF Television and Parasol, Presented by Victoria Coren Mitchell.
- [4] A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter, and R. Vollgraf. FLAIR: An easy-to-use framework for state-of-the-art NLP. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-4010.
- [5] S. Altman. Planning for AGI and beyond. *OpenAI Blog*, February, 2023.
- [6] C. Anotado, T. Ruddle, and J. Halbur. The Only Connect Database.
- [7] S. Basu, I. Davidson, and K. Wagstaff. *Constrained clustering: Advances in algorithms, theory, and applications*. CRC Press, 2008.
- [8] Z. Beda and S. M. Smith. Chasing red herrings: Memory of distractors causes fixation in creative problem solving. *Memory & Cognition*, 46:671–684, 2018.
- [9] M. Benedek and A. C. Neubauer. Revisiting Mednick’s model on creativity-related differences in associative hierarchies. Evidence for a common path to uncommon thought. *The Journal of creative behavior*, 47(4):273–289, 2013.
- [10] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, et al. On the opportunities and risks of foundation models. *ArXiv preprint*, abs/2108.07258, 2021.
- [11] S. Borgeaud, A. Mensch, J. Hoffmann, T. Cai, E. Rutherford, K. Millican, G. van den Driessche, J. Lespiau, B. Damoc, A. Clark, D. de Las Casas, A. Guy, J. Menick, R. Ring, T. Hennigan, S. Huang, L. Maggiore, C. Jones, A. Cassirer, A. Brock, M. Paganini, G. Irving, O. Vinyals, S. Osindero, K. Simonyan, J. W. Rae, E. Elsen, and L. Sifre. Improving language models by retrieving from trillions of tokens. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvári, G. Niu, and S. Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 2206–2240. PMLR, 2022.
- [12] P. S. Bradley, K. P. Bennett, and A. Demiriz. Constrained k-means clustering. *Microsoft Research, Redmond*, 20(0):0, 2000.
- [13] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [14] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, et al. Sparks of artificial general intelligence: Early experiments with GPT-4. *ArXiv preprint*, abs/2303.12712, 2023.
- [15] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, et al. Evaluating large language models trained on code. *ArXiv preprint*, abs/2107.03374, 2021.

- [16] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al. PaLM: Scaling language modeling with pathways. 2022. *ArXiv preprint*, abs/2204.02311, 2022.
- [17] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423.
- [18] N. Du, Y. Huang, A. M. Dai, S. Tong, D. Lepikhin, Y. Xu, M. Krikun, Y. Zhou, A. W. Yu, O. Firat, B. Zoph, L. Fedus, M. P. Bosma, Z. Zhou, T. Wang, Y. E. Wang, K. Webster, M. Pellat, K. Robinson, K. S. Meier-Hellstern, T. Duke, L. Dixon, K. Zhang, Q. V. Le, Y. Wu, Z. Chen, and C. Cui. Glam: Efficient scaling of language models with mixture-of-experts. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvári, G. Niu, and S. Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 5547–5569. PMLR, 2022.
- [19] K. Ethayarajh. How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China, 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1006.
- [20] C. Fellbaum. Wordnet. In *Theory and applications of ontology: computer applications*, pages 231–243. Springer, 2010.
- [21] E. B. Fowlkes and C. L. Mallows. A method for comparing two hierarchical clusterings. *Journal of the American statistical association*, 78(383):553–569, 1983.
- [22] J. Gao, L. Ge, K. Li, H. Q. Ngo, and A. Zhang. On handling negative transfer and imbalanced distributions in multiple source transfer learning. In *Proceedings of the 13th SIAM International Conference on Data Mining, May 2-4, 2013. Austin, Texas, USA*, pages 261–269. SIAM, 2013. doi: 10.1137/1.9781611972832.29.
- [23] E. Grave, P. Bojanowski, P. Gupta, A. Joulin, and T. Mikolov. Learning word vectors for 157 languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, 2018. European Language Resources Association (ELRA).
- [24] K. Guu, K. Lee, Z. Tung, P. Pasupat, and M. Chang. Retrieval augmented language model pre-training. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 3929–3938. PMLR, 2020.
- [25] J. A. Hartigan and M. A. Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108, 1979.
- [26] B. Heinzerling and M. Strube. BPEmb: Tokenization-free pre-trained subword embeddings in 275 languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, 2018. European Language Resources Association (ELRA).
- [27] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt. Measuring massive multitask language understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [28] L. Hubert and P. Arabie. Comparing partitions. *Journal of classification*, 2:193–218, 1985.
- [29] G. Izacard, P. Lewis, M. Lomeli, L. Hosseini, F. Petroni, T. Schick, J. A. Yu, A. Joulin, S. Riedel, and E. Grave. Few-shot learning with retrieval augmented language models. *ArXiv preprint*, abs/2208.03299, 2022.

- [30] A. Jacovi, A. Caciularu, O. Goldman, and Y. Goldberg. Stop uploading test data in plain text: Practical strategies for mitigating data contamination by evaluation benchmarks. *ArXiv preprint*, abs/2305.10160, 2023.
- [31] M. I. Jordan. Artificial intelligence—the revolution hasn’t happened yet. *Harvard Data Science Review*, 1(1):1–9, 2019.
- [32] J. Kaddour, J. Harris, M. Mozes, H. Bradley, R. Raileanu, and R. McHardy. Challenges and applications of large language models. *ArXiv preprint*, abs/2307.10169, 2023.
- [33] G. Ke, D. He, and T. Liu. Rethinking positional encoding in language pre-training. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [34] N. Kohn and S. M. Smith. Partly versus completely out of your mind: Effects of incubation and distraction on resolving fixation. *The Journal of Creative Behavior*, 43(2):102–118, 2009.
- [35] H. Levesque, E. Davis, and L. Morgenstern. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*, 2012.
- [36] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.703.
- [37] P. S. H. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W. Yih, T. Rocktäschel, S. Riedel, and D. Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [38] P. Liang, R. Bommasani, T. Lee, D. Tsipras, D. Soylu, M. Yasunaga, Y. Zhang, D. Narayanan, Y. Wu, A. Kumar, et al. Holistic evaluation of language models. *ArXiv preprint*, abs/2211.09110, 2022.
- [39] C.-Y. Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, 2004. Association for Computational Linguistics.
- [40] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. *ArXiv preprint*, abs/1907.11692, 2019.
- [41] N. Lu, S. Liu, R. He, and K. Tang. Large language models can be guided to evade AI-generated text detection. *ArXiv preprint*, abs/2305.10847, 2023.
- [42] A. S. Luchins and E. H. Luchins. Rigidity of behavior: A variational approach to the effect of Einstellung. 1959.
- [43] M. Marko, D. Michalko, and I. Riečanský. Remote associates test: An empirical proof of concept. *Behavior research methods*, 51:2700–2711, 2019.
- [44] S. Mednick. The associative basis of the creative process. *Psychological review*, 69(3):220, 1962.
- [45] S. A. Mednick. The remote associates test. *The Journal of Creative Behavior*, 1968.
- [46] G. A. Miller. WordNet: A lexical database for English. In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*, 1992.
- [47] R. Navigli. Word sense disambiguation: A survey. *ACM computing surveys (CSUR)*, 41(2): 1–69, 2009.
- [48] OpenAI. GPT-4 technical report. *ArXiv preprint*, abs/2303.08774, 2023.

- [49] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [50] J. Pan, T. Gao, H. Chen, and D. Chen. What in-context learning "learns" in-context: Disentangling task recognition and task learning. In *Annual Meeting of the Association for Computational Linguistics*, 2023.
- [51] J. Pennington, R. Socher, and C. Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162.
- [52] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana, 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1202.
- [53] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- [54] A. Ramdas, N. García Trillos, and M. Cuturi. On wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017.
- [55] O. Sainz, O. L. de Lacalle, E. Agirre, and G. Rigau. What do language models know about word senses? zero-shot wsd with language models and domain inventories. *ArXiv preprint*, abs/2302.03353, 2023.
- [56] V. Sanh, L. Debut, J. Chaumond, and T. Wolf. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *ArXiv preprint*, abs/1910.01108, 2019.
- [57] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. *ArXiv preprint*, abs/2211.05100, 2022.
- [58] J. Shlens. A tutorial on principal component analysis. *arXiv:1404.1100*, 2014.
- [59] M. Shoenberger, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro. Megatron-LM: Training multi-billion parameter language models using model parallelism. *ArXiv preprint*, abs/1909.08053, 2019.
- [60] S. M. Smith. The constraining effects of initial ideas. *Group creativity: Innovation through collaboration*, pages 15–31, 2003.
- [61] S. M. Smith and S. E. Blankenship. Incubation effects. *Bulletin of the Psychonomic Society*, 27(4):311–314, 1989.
- [62] S. M. Smith and S. E. Blankenship. Incubation and the persistence of fixation in problem solving. *The American journal of psychology*, pages 61–87, 1991.
- [63] S. M. Smith and D. R. Tindell. Memory blocks in word fragment completion caused by involuntary retrieval of orthographically related primes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(2):355, 1997.
- [64] K. Song, X. Tan, T. Qin, J. Lu, and T. Liu. MpNet: Masked and permuted pre-training for language understanding. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [65] Y. Song, C. Cui, S. Khanuja, P. Liu, F. Faisal, A. Ostapenko, G. I. Winata, A. F. Aji, S. Cahyawijaya, Y. Tsvetkov, et al. GlobalBench: A benchmark for global progress in natural language processing. *ArXiv preprint*, abs/2305.14716, 2023.

- [66] A. Srivastava, A. Rastogi, A. Rao, A. A. M. Shob, A. Abid, A. Fisch, A. R. Brown, A. Santoro, A. Gupta, A. Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *ArXiv preprint*, abs/2206.04615, 2022.
- [67] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. LLaMA: Open and efficient foundation language models. *ArXiv preprint*, abs/2302.13971, 2023.
- [68] L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008.
- [69] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
- [70] C. Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.
- [71] N. X. Vinh, J. Epps, and B. J. *Information theoretic measures for clusterings comparison*, volume 09. Proceedings of the 26th Annual International Conference on Machine Learning - ICML, 2009.
- [72] K. Wagstaff, C. Cardie, S. Rogers, and S. Schrödl. Constrained k-means clustering with background knowledge. In C. E. Brodley and A. P. Danyluk, editors, *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001), Williams College, Williamstown, MA, USA, June 28 - July 1, 2001*, pages 577–584. Morgan Kaufmann, 2001.
- [73] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3261–3275, 2019.
- [74] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [75] L. Wang, N. Yang, X. Huang, B. Jiao, L. Yang, D. Jiang, R. Majumder, and F. Wei. Text embeddings by weakly-supervised contrastive pre-training. *ArXiv preprint*, abs/2212.03533, 2022.
- [76] Z. Wang, Z. Dai, B. Póczos, and J. G. Carbonell. Characterizing and avoiding negative transfer. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 11293–11302. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.01155.
- [77] J. Wei, M. Bosma, V. Y. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le. Finetuned language models are zero-shot learners. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [78] J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, et al. Emergent abilities of large language models. *ArXiv preprint*, abs/2206.07682, 2022.
- [79] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.

- [80] L. White, R. Togneri, W. Liu, and M. Bennamoun. Finding word sense embeddings of known meaning. In *Computational Linguistics and Intelligent Text Processing: 19th International Conference, CICLing 2018, Hanoi, Vietnam, March 18–24, 2018, Revised Selected Papers, Part II*, pages 3–16. Springer, 2023.
- [81] Wikipedia contributors. Only connect — Wikipedia, the free encyclopedia. [https://en.wikipedia.org/w/index.php?title=Only\\_Connect&oldid=1157929067](https://en.wikipedia.org/w/index.php?title=Only_Connect&oldid=1157929067), 2023. [Online; accessed 7-June-2023].
- [82] J. Wiley. Expertise as mental set: The effects of domain knowledge in creative problem solving. *Memory & cognition*, 26:716–730, 1998.
- [83] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, and A. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6.
- [84] R. Wolfe and A. Caliskan. American== white in multimodal language-and-image ai. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 800–812, 2022.
- [85] G. Wood and J. Pennington. Encoding and retrieval from long-term storage. *Journal of Experimental Psychology*, 99(2):243, 1973.
- [86] C.-L. Wu, S.-Y. Huang, P.-Z. Chen, and H.-C. Chen. A systematic review of creativity-related studies applying the remote associates test from 2000 to 2019. *Frontiers in psychology*, 11: 573432, 2020.
- [87] S. Yao, D. Yu, J. Zhao, I. Shafran, T. L. Griffiths, Y. Cao, and K. Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *ArXiv preprint*, abs/2305.10601, 2023.
- [88] B. Yin, M. Zhao, L. Guo, and L. Qiao. Sentence-BERT and k-means based clustering technology for scientific and technical literature. In *2023 15th International Conference on Computer Research and Development (ICCRD)*, pages 15–20. IEEE, 2023.
- [89] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, et al. OPT: Open pre-trained transformer language models. *ArXiv preprint*, abs/2205.01068, 2022.
- [90] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [91] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, et al. A survey of large language models. *ArXiv preprint*, abs/2303.18223, 2023.

# Appendices

## A Additional Experiments

**Task 1 – Grouping** In addition to grouping clue words using token embeddings (discussed in the main paper §4), we also ran grouping the words by clustering on ‘contextual’ embeddings. We experimentally induce ‘context’ by joining the sixteen (16) word tokens (in a random order) into a single pseudo-sentence. The embeddings for each token were different based on the ordering of the tokens. We repeat the random ordering sixteen times and report the mean and variance of the results obtained in Table 6.

	WD ↓	FMS ↑	ARI ↑	AMI ↑	# Solved Walls	# Correct Groups
ELMo <sub>LARGE</sub>	90.0 ± .3	23.6 ± .4	4.5 ± .5	5.6 ± .7	0 ± 0	19 ± 3
DistilBERT <sub>BASE</sub>	88.4 ± .7	26.7 ± .3	8.3 ± .4	10.4 ± .5	0 ± 0	30 ± 4
BERT <sub>LARGE</sub>	<b>87.2 ± .6</b>	<b>28.3 ± .5</b>	<b>10.4 ± .6</b>	<b>12.8 ± .7</b>	0 ± 0	<b>46 ± 5</b>
BERT <sub>BASE</sub>	87.7 ± .5	28.0 ± .2	10.0 ± .3	12.4 ± .4	0 ± 0	39 ± 2
RoBERTa <sub>LARGE</sub>	88.4 ± .5	25.9 ± .2	7.4 ± .3	9.3 ± .4	0 ± 0	30 ± 4
all-mpnet <sub>BASE</sub>	87.6 ± .5	28.0 ± .3	10.0 ± .4	12.4 ± .5	0 ± 0	38 ± 3
E5 <sub>LARGE</sub>	87.7 ± .5	28.1 ± .3	10.2 ± .4	12.7 ± .5	0 ± 0	37 ± 4
E5 <sub>BASE</sub>	<b>87.2 ± .3</b>	28.2 ± .2	10.2 ± .3	12.5 ± .4	0 ± 0	<b>46 ± 5</b>
Human Performace	–	–	–	–	285 / 494	1405 / 1976

Table 6: Results of selected models on Task 1 (Grouping) using contextual embeddings. WD: Wasserstein Distance. FMS: Fowlkes Mallows Score. ARI: Adjusted Rand Index. NMI: Normalized Mutual Information. Mean ± standard deviation over 16 random seeds is shown. **Bold**: best scores.

**Task 2 – Connections** In addition to prompting based results on GPT-4 (discussed in §4), we ran experiments on additional LLMs like LLaMa [67] (7B, 13B) using pre-trained configuration weights obtained by permission from Meta AI. However, without additional fine-tuning on the specific task, these LLMs were unable to solve the task in a meaningful manner. To elucidate, LLaMa generated a bunch of hallucinated words with unequal group sizes. We omit these unintelligible results for brevity.

## B Additional Figures

In this section, we provide additional t-SNE projections of embeddings from various methods used.

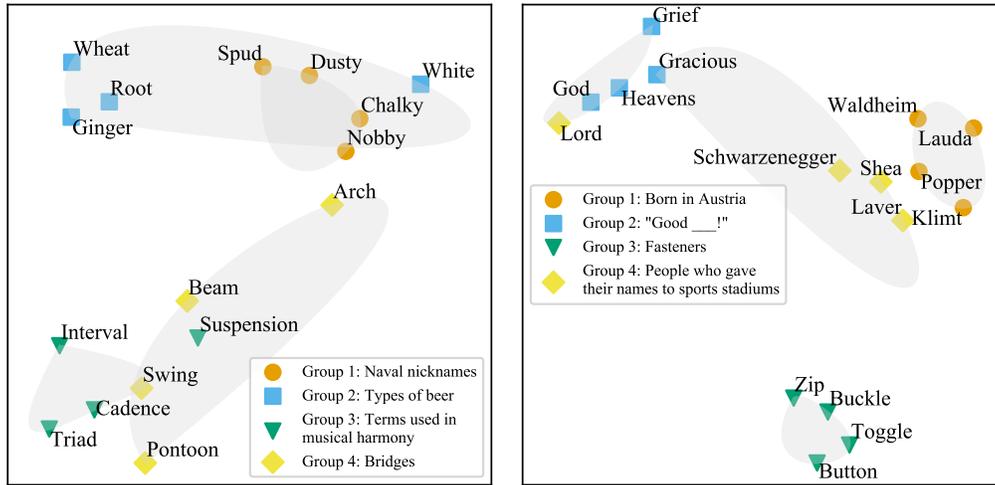


Figure 7: Solved wall for Task 1 (Grouping) using GloVe. **Left:** (`wall_id="7ed3"`), the embedding model erroneously associated the clue “*Suspension*” with the connection “*Bridges*”; however, this association is an example of a red herring. “*Suspension*” is “*a term used in musical harmony*” in this context. **Right:** (`wall_id="5e3c"`), shows that clue “*Lord*” is close to “*God, Heavens, and Grief*” in the embedding space, which matches the “*Good \_\_\_!*” connection. However, this is another example of a red herring as, in this context, “*Lord*” refers to “*Lord’s cricket Ground*”, a cricket stadium named after “*Thomas Lord*”.

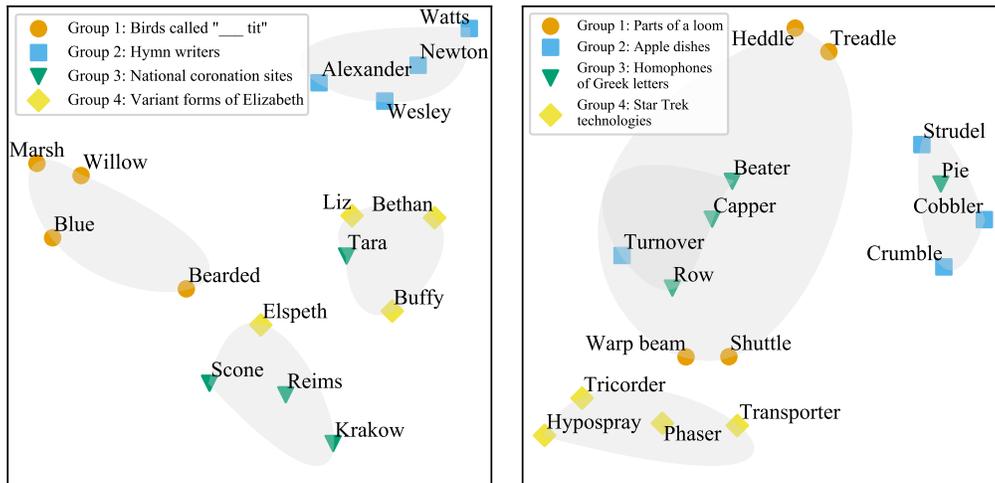


Figure 8: Solved wall for Task 1 (Grouping) using FastText (Crawl). **Left:** (`wall_id="d5e6"`), the embedding model erroneously associated the clue “*Tara*” other girls’ names; but here, “*Tara*” is short for “*Hill of Tara*” and belongs to the “*national coronation sites*” group. **Right:** (`wall_id="4c22"`), shows that clue “*Pie*” associated with the connection “*Apple*”. Even though it is acceptable in general context, here it represents a homophone for the Greek letter “ $\pi$ ”.

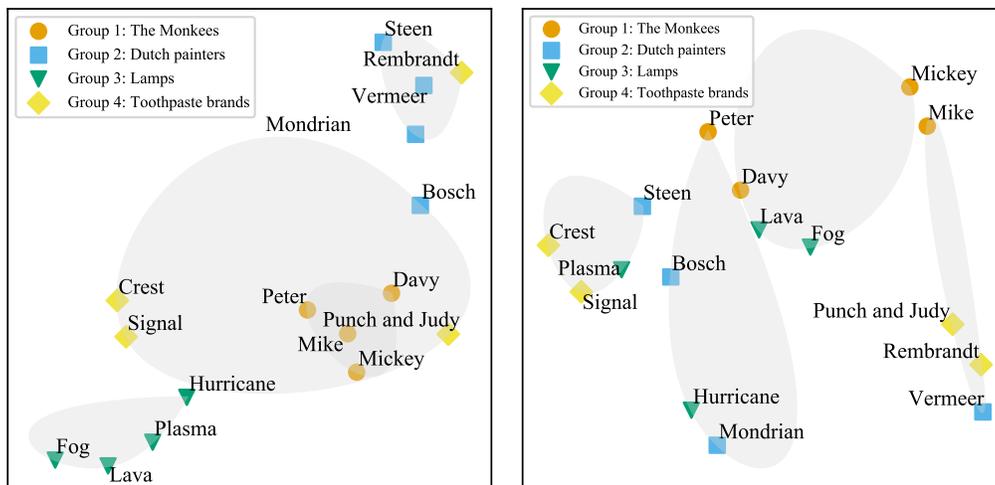


Figure 9: Solved wall (wall\_id="2d8f") for Task 1 (Grouping) using BERT<sub>LARGE</sub> with both static and contextual embeddings. **Left:** contextual embedding solved 3/4 groups. Here the clue “*Rambrandt*” is placed near other Dutch painters. The correct grouping for this clue in this wall is “*Toothpaste Brands*”. **Right:** static embedding solved 0/4 groups.

## C Effects of Red-Herrings: Additional Experiments, Analysis and Results

### C.1 Additional Datasets

Both of the additional datasets described in this section for ablation experiments have been made available via our code repositories on Github and HuggingFace.

#### C.1.1 OCW-Randomized Dataset

This test dataset generates a version of the test set where red herrings are removed or largely reduced in frequency. This is achieved by rebuilding every wall using a randomly selected group from different walls. We only applied the process to the (original OCW) test set, the train and validation sets are left untouched.

**Method** For each wall in the existing test set, we leave the first group untouched, and sample three new groups, each from a different wall, such that none of the groups share a word in common. The connections for each group are unmodified. The result is a new version of the test set where every wall is composed of 4 random groups from 4 different walls.

#### C.1.2 OCW-WordNet Dataset

WordNet [46, 20] is a large lexical database of English. Nouns, verbs, adjectives, and adverbs are grouped into sets of cognitive synonyms (synsets), each expressing a distinct concept. We use the hypernym/hyponym (or superlative/subordinative) hierarchical lexical structure aggregated in WordNet to generate an easy test set to further analyze the effects of red-herring in OCW.

**Method** We use the existing words in a wall to select synonyms from the word’s synsets. We only consider synsets that have at least five synonymous lexical names, then randomly sample four words. The original test set word and its definition (`ss.definition()`) subsequently becomes the connection phrase for the group. Four groups were generated for each wall, and the easy wall generation process was repeated for the total number of walls (494) in the original test data set.

For the group connections, we concatenate the superlative parent word with a synset definition giving a description of the word. This allows for an ideal semantic similarity score to be calculated using BERTScore. For a few cases (approx. 70/494 walls in the test set), the number of generated groups per wall is less than four, due to the unavailability of direct synonyms from word synsets. In those edge cases, we generate and append groups using common hypernym words like animal, mammal, furniture, etc. to ensure a wall is valid with four groups.

A sample generated easy group is shown below, where we prefix the `group_id` from the original OCW dataset with ‘easy’ to aid with mapping or identification.

```
{
  ...
  "group_3": {
    "group_id": "easy_691a_3",
    "gt_words": ["gibe", "shaft", "jibe", "barb"],
    "gt_connection": "Shaft: an aggressive remark directed at a person
    like a missile and intended to have a telling effect"
  }
  ...
}
```

Further, we generate easy to train and validation sets mimicking the original dataset, package and release these three additional easy sets, as **OCW-WordNet** as added contributions.

## C.2 Results of Ablation Experiments

### C.2.1 PLMs: Performance on Task 1 (Grouping)

We perform and present the results using ‘static’ embeddings due to the noted superior results and the word order related deficiency already shown by using contextual embeddings pertinent to our task setup.

	WD ↓	FMS ↑	ARI ↑	AMI ↑	# Solved Walls	# Correct Groups
<i>Classic Word Embeddings</i>						
GloVe	76.8 ± .7	39.2 ± .3	24.0 ± .4	27.7 ± .4	7 ± 1	213 ± 8
FastText (Crawl)	76.1 ± .5	40.5 ± .3	25.0 ± .6	28.6 ± .7	<b>13 ± 1</b>	236 ± 7
FastText (News)	79.3 ± .5	36.8 ± .3	21.0 ± .3	24.5 ± .4	5 ± 1	176 ± 6
<i>Pre-trained Language Models (PLMs)</i>						
ELMo <sub>LARGE</sub>	80.9 ± .4	35.2 ± .3	18.9 ± .3	22.2 ± .4	3 ± 1	154 ± 6
DistilBERT <sub>BASE</sub>	82.3 ± .6	34.2 ± .4	17.7 ± .5	21.1 ± .5	1 ± 1	124 ± 8
BERT <sub>LARGE</sub>	86.2 ± .4	29.2 ± .3	11.5 ± .3	14.2 ± .4	0 ± 0	66 ± 4
BERT <sub>BASE</sub>	87.5 ± .4	27.7 ± .3	9.6 ± .6	11.8 ± .5	0 ± 0	48 ± 4
RoBERTa <sub>LARGE</sub>	86.7 ± .5	28.6 ± .2	10.8 ± .3	13.4 ± .3	1 ± 0	56 ± 4
<i>Sentence Transformers</i>						
all-mpnet <sub>BASE</sub>	81.4 ± .4	35.1 ± .4	18.9 ± .5	22.0 ± .6	8 ± 1	154 ± 7
E5 <sub>LARGE</sub>	76.0 ± .5	40.7 ± .3	25.9 ± .4	29.7 ± .4	8 ± 1	230 ± 5
E5 <sub>BASE</sub>	<b>75.1 ± .8</b>	<b>41.8 ± .3</b>	<b>27.2 ± .3</b>	<b>31.1 ± .3</b>	8 ± 1	<b>249 ± 8</b>
Human Performance	-	-	-	-	-	-

Table 7: Results of **OCW-Randomized** using static embeddings. WD: Wasserstein Distance. FMS: Fowlkes Mallows Score. ARI: Adjusted Rand Index. NMI: Normalized Mutual Information. Mean ± standard deviation over 16 random seeds is shown. **Bold**: best scores.

	WD ↓	FMS ↑	ARI ↑	AMI ↑	# Solved Walls	# Correct Groups
<i>Classic Word Embeddings</i>						
GloVe	43.0 ± 1.0	66.1 ± .4	57.4 ± .5	60.9 ± .5	118 ± 3	886 ± 1
FastText (Crawl)	30.6 ± 1.0	75.8 ± .6	69.6 ± .7	72.4 ± .7	195 ± 6	1173 ± 18
FastText (News)	44.9 ± 1.2	64.9 ± .5	55.9 ± .6	59.5 ± .6	105 ± 3	844 ± 12
<i>Pre-trained Language Models (PLMs)</i>						
ELMo <sub>LARGE</sub>	52.5 ± 1.1	58.9 ± .3	48.2 ± .4	52.5 ± .4	67 ± 3	682 ± 9
DistilBERT <sub>BASE</sub>	45.5 ± 1.0	64.1 ± .4	55.0 ± .5	58.7 ± .5	105 ± 3	835 ± 13
BERT <sub>LARGE</sub>	76.9 ± 1.0	38.9 ± .2	23.4 ± .3	27.5 ± .3	7 ± 0	197 ± 6
BERT <sub>BASE</sub>	73.0 ± 1.3	42.5 ± .5	27.9 ± .6	32.5 ± .6	8 ± 2	268 ± 12
RoBERTa <sub>LARGE</sub>	57.4 ± 1.3	54.8 ± .3	43.3 ± .3	47.5 ± .3	48 ± 2	573 ± 8
<i>Sentence Transformers</i>						
all-mpnet <sub>BASE</sub>	<b>22.6 ± .7</b>	<b>81.9 ± .4</b>	<b>77.1 ± .5</b>	<b>79.4 ± .4</b>	<b>256 ± 4</b>	<b>1365 ± 12</b>
E5 <sub>LARGE</sub>	23.6 ± .8	80.9 ± .4	75.9 ± .5	78.3 ± .4	250 ± 4	1347 ± 12
E5 <sub>BASE</sub>	26.9 ± .9	78.0 ± .4	72.3 ± .5	75.0 ± .5	224 ± 4	1259 ± 10
Human Performance	-	-	-	-	-	-

Table 8: Results of **OCW-WordNet** using static embeddings. WD: Wasserstein Distance. FMS: Fowlkes Mallows Score. ARI: Adjusted Rand Index. NMI: Normalized Mutual Information. Mean ± standard deviation over 16 random seeds is shown. **Bold**: best scores.

### C.2.2 LLMs: Performance on Task 1 (Grouping) using GPT3.5/4

Here we present the results of repeating Task 1 (grouping) on the ablation datasets OCW-Randomized (C.1.1) and OCW-Wordnet (C.1.2) to analyze the effects of red-herrings in walls on LLM performance.

The results adhere to the expected results of superior performance with the dilution/removal of red-herrings from the walls.

	# In-context Examples	WD ↓	FMS ↑	ARI ↑	AMI ↑	# Solved Walls	# Correct Groups
GPT-3.5-turbo	0-shot	74.3	40.4	26.4	29.8	5	274
	1-shot	72.0	43.1	29.0	32.3	12	315
	3-shot	72.7	43.4	29.4	32.9	10	306
	5-shot	70.7	44.6	30.9	34.4	16	337
	10-shot	70.5	43.8	30.0	33.5	17	333
GPT-4	0-shot	58.2	56.2	45.4	48.8	59	595
	1-shot	55.1	<b>58.0</b>	<b>47.5</b>	<b>51.0</b>	57	644
	3-shot	55.0	57.5	46.9	50.3	62	649
	5-shot	<b>54.1</b>	<b>58.0</b>	<b>47.5</b>	50.9	<b>68</b>	<b>655</b>
	10-shot	56.6	56.1	45.1	48.5	55	614
Human Performance		–	–	–	–	–	–

Table 9: Results of **OCW-Randomized** using Large Language Models. WD: Wasserstein Distance. FMS: Fowlkes Mallows Score. ARI: Adjusted Rand Index. NMI: Normalized Mutual Information. **Bold**: best scores.

	# In-context Examples	WD ↓	FMS ↑	ARI ↑	AMI ↑	# Solved Walls	# Correct Groups
GPT-3.5-turbo	0-shot	15.9	86.3	83.4	84.9	337	1522
	1-shot	24.8	76.4	74.4	75.4	320	1400
	3-shot	8.65	92.7	91.2	91.8	415	1748
	5-shot	8.09	94.0	92.4	93.1	415	1759
	10-shot	6.55	95.3	94.0	94.7	428	1800
GPT-4	0-shot	<b>1.51</b>	<b>98.5</b>	<b>98.0</b>	<b>98.2</b>	<b>471</b>	<b>1926</b>
	1-shot	19.2	87.9	84.3	83.7	304	1581
	3-shot	21.5	86.6	82.5	81.8	279	1537
	5-shot	19.1	88.1	84.5	83.8	298	1584
	10-shot	11.2	92.9	90.7	90.4	378	1742
Human Performance		–	–	–	–	–	–

Table 10: Results of **OCW-WordNet** using Large Language Models. WD: Wasserstein Distance. FMS: Fowlkes Mallows Score. ARI: Adjusted Rand Index. NMI: Normalized Mutual Information. **Bold**: best scores.