Nothing New Under the Sun?
The Study of Biblical Hebrew in the Era of Generative Pre-trained AI

A.G. Elrod
Vrije Universiteit, Amsterdam
elrodag@gmail.com

Abstract: This article investigates the potential impact of generative artificial intelligence, specifically OpenAI's GPT, on the field of biblical studies, particularly biblical Hebrew. The study is divided into three main categories: (1) knowledge retrieval or language understanding, (2) generative modeling or creative problem solving, and (3) command interpretation or query parsing. Experiments are conducted using OpenAI's GPT, the ETCBC's BHSA dataset, and Text-Fabric Python libraries. Results demonstrate GPT's limitations and proficiencies in biblical Hebrew and its capacity to employ its proficiencies creatively in problem-solving scenarios involving multifaceted forms of reasoning. The study concludes that understanding the capabilities and potential trajectories of these technologies is vital for biblical Hebrew scholarship, as they already possess the capacity to disrupt established scholarly norms and democratize access to advanced tools.

Keywords: Biblical Hebrew, generative artificial intelligence, large language models, GPT.

Introduction

It would be remiss not to acknowledge the staggering pace at which the technological landscape is evolving. The swift ascent of generative artificial intelligence in recent months has elicited reactions and prognostications of equal magnitude. As the rapidity of this technological revolution increases, forecasts suggest unprecedented levels of disruption in technology and economics. A report by Goldman Sachs, "The Potentially Large Effects of Artificial Intelligence on Economic Growth," anticipates AI automation could result in the loss of nearly 300 million jobs in the not-too-distant future.[1] Consequently, some industry leaders, disconcerted by the rapid advancement of Large Language Models (LLMs), are advocating for a temporary halt on what they term "Giant AI Experiments."[2] Yet, amidst this climate of uncertainty, opportunists eagerly envision a new horizon of possibilities. Indeed, such rapid transformation invites a reconsideration of the wisdom of Qohelet: "There is nothing new under the sun."

Regardless of the eventual ramifications of AI, it is undeniable that this technology holds significant potential across a broad range of fields. Of note is the fact that recent advancements in Natural Language Processing (NLP) models have paved the way for novel explorations in linguistics and statistical analysis of various corpora. This includes the study of ancient languages, given that LLMs such as OpenAI's GPT — the technology that powers ChatGPT — have been trained on these. This article poses the question: how might this swiftly evolving technology impact the field of biblical studies,

---

[1] "Generative AI Could Raise Global GDP by 7%," April 5, 2023, https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html.

[2] Future of Life Institute, "Pause Giant AI Experiments: An Open Letter," March 22, 2023, https://futureoflife.org/open-letter/pause-giant-ai-experiments/.

particularly biblical Hebrew? Our underlying experiments aim to shed light on the potential contributions of contemporary technological advancements to the study of biblical Hebrew.

Our experimentation aligns with three main categories, each reflecting different capabilities of OpenAI's Generative Pre-trained AI (GPT) relevant to the diverse field of biblical studies. Although there is a certain degree of overlap among these categories, they are designed to evaluate distinct aspects of NLP AI. Specifically, the areas examined are (1) knowledge retrieval or language understanding, (2) generative modeling or creative problem solving, and (3) command interpretation or query parsing.

1. Knowledge Retrieval or Language Understanding: This facet explores GPT's capacity to tap into its pre-trained knowledge reservoir pertinent to specific topics — in our case, those topics germane to biblical exegesis and hermeneutics. From an exegetical standpoint, we scrutinize the AI's competence in responding to inquiries regarding ancient Hebrew syntax, grammar, and the content of the Hebrew Bible. Hermeneutically, we challenge the model to deliver interpretations in line with prescribed frameworks.

2. Generative Modeling or Creative Problem Solving: This set of experiments delves into the AI's creative problem-solving and generative modeling capabilities through tasks such as filling simulated and actual lacunae in ancient documents and integrating disparate passages of ancient Hebrew literature into a single coherent narrative. Given that generative models in AI are designed to generate new content bearing a statistical structure akin to the data they were trained on, GPT is anticipated to exhibit proficiency in this realm.

3. Command Interpretation or Query Parsing: The final series of experiments assess GPT's ability to interact with Text-Fabric and the BHSA, examining its proficiency in command interpretation. Essentially, we evaluate the model's capability to understand and respond to high-technical plain language prompts by translating them into efficient Python queries. Considering GPT's well-documented proficiency in generating effective and accurate code across various domains, this is another sphere where GPT (specifically GPT-4) is expected to excel.

Use of Technology

Prior to delving into the experimentation, it is imperative to establish a foundational understanding of the tools utilized in the forthcoming analyses. The experimental exploration pertains to the employment of OpenAI's Generative Pre-trained AI, the ETCBC's BHSA dataset, and Text-Fabric Python libraries.

1. GPT: OpenAI's Generative Pretrained Transformer (GPT) is a technology in the realm of natural language processing and artificial intelligence. Built on the principle of transformers, GPT models process words in relation to all other words in a sentence, regardless of their respective positions. This diverges from traditional language models that process words sequentially. Through an approach called unsupervised learning, these models are trained on vast amounts of data, enabling them to generate human-like text (and other reasoning abilities) by predicting what word should logically follow a given input. GPT's prowess lies not only in its language generation capabilities but also in its ability to understand context, nuance, and

    even perform tasks such as coding, translation, summarization, and (increasingly) creative problem-solving.[3]

2. BHSA: The Eep Talstra Centre for Bible and Computer (ETCBC) Biblical Hebrew Syntax (BHSA) dataset is a comprehensive resource for scholars studying the Hebrew Bible. This dataset is notable for its extensive linguistic annotations and syntactic structuring of the entire Hebrew Bible, allowing for in-depth computational research. Each word in the BHSA dataset is encoded with an array of linguistic features, including morphological, lexical, syntactical, and semantic attributes, to name a few. It also includes clause and phrase segmentation and constituent analysis, offering researchers a multi-faceted resource for exploring biblical Hebrew on a granular level.[4]

3. The Text-Fabric (TF) Python library is a powerful tool for navigating, querying, and analyzing text datasets like the BHSA. Designed with the challenges of handling large, complex datasets in mind, TF provides an efficient and flexible way to interact with structured data. One of the key features of Text-Fabric is its ability to query text using syntax, thereby offering a comprehensive and nuanced approach to text analysis. It also permits the addition of new data features, allowing scholars to annotate and customize datasets according to their research needs. By creating a platform that is both robust and adaptable, Text-Fabric contributes to more dynamic, innovative study in fields such as biblical studies, linguistics, and digital humanities.[5]

Experiments

1. Knowledge Retrieval or Language Understanding

1.1 Appearance of the conjunction וְ in Genesis 1?

As a preliminary test, we assessed GPT's pre-trained capabilities to parse the Hebrew Bible with a straightforward prompt, "How many times is the conjunction וְ used in the first chapter of Genesis?" We leveraged GPT's API framework to address this prompt with various models and configurations. Models spanning from versions 3.5 to 4 (as of May 24) were evaluated using temperature settings of 0 and 16. We excluded temperature settings greater than 1 from this analysis, as they frequently resulted in hallucinatory or nonsensical responses.

Previous research has consistently demonstrated that GPT's performance is enhanced through extended collaboration and techniques such as few-shot and follow-up prompting [source: OpenAI Cookbook]. In alignment with this knowledge, we incorporated straightforward follow-up prompting

---

[3] OpenAI, "GPT-4 Is OpenAI's Most Advanced System, Producing Safer and More Useful Responses," n.d., accessed May 28, 2023.

[4] ETCBC, "BHSA," accessed May 28, 2023, https://etcbc.github.io/bhsa/.

[5] Dirk Roorda, "Text-Fabric," accessed May 28, 2023, https://annotation.github.io/text-fabric/tf/.

[6] Changing the temperature setting affects the model such that lower temperature produce more predictable or "accurate" results. Higher temperatures produce results that are more random or "creative." For more information, this setting is explained in OpenAI's GPT API documentation; cf. OpenAI, "API Reference: Create Completion," *OpenAI API Reference*, accessed May 28, 2023, https://platform.openai.com/docs/api-reference/models.

techniques in our experiments. The initial follow-up prompt seeks to verify the response with a simple question, "Are you certain?" Subsequently, an additional prompt is introduced with the intent of stimulating more systematic thought in the AI, potentially improving its response, "How many times does the conjunction "וְ" appear in each verse of Genesis 1?" Lastly, to further probe GPT-4's capabilities, we provided the Hebrew passage (in UTF-8 encoding) as a prompt for direct analysis.

Results:

Table 1. Responses and performance on the prompt: How many times is the conjunction ו used in the first chapter of Genesis? (Correct: 101)

GPT Model	Temp	Prompt Response	Certainty Response	Per Verse Response	Given Text Response	Accuracy
3.5	0	29	29	35		1 NC, 2 NC, 3 NC
3.5	1	7	10	11		1 NC, 2 NC, 3 NC
3.5 (0301)	0	7	10	29		1 NC, 2 NC, 3 NC
3.5 (0301)	1	19	19	25		1 NC, 2 NC, 3 NC
4 (0314)	0	102	106	89		1 NC, 2 NC, 3 NC
4 (0314)	1	102	74	55		1 NC, 2 NC, 3 NC
4 (current)	0	102	106	62	30	1 NC, 2 NC, 3 NC, 4 NC
4 (current)	1	102	102	114	37	1 NC, 2 NC, 3 NC, 4 NC
C = correct, NC = not correct, PC = partially correct						

Interpretation of Results:

The data presented in Table 1 requires minimal elaboration, as all models consistently fail to provide an accurate answer. A salient observation, however, is the perceptible performance increase between all iterations of version 3.5 and version 4. The responses of version 3.5 deviate significantly from the expected count of 101. In contrast, version 4, at temperature settings of both 0 and 1, delivers responses so close to the expected figure (102 instead of 101) that they could potentially be viewed as statistical anomalies. Given this consistent proximity to the anticipated response, a more granular examination should be pursued to identify the source of this discrepancy.

Another intriguing observation is GPT's susceptibility to an error often seen among students — an inclination to second-guess initial responses. Much like a student's first response on an examination often tends to be the most accurate, GPT-4's initial reaction to the prompt also appears more precise.

Upon being challenged with the follow-up prompt, "Are you certain?" the model seemingly expresses uncertainty and offers an alternate figure — one less accurate than the original. Further, inviting GPT to analyze the actual text in UTF-8 triggers a dramatic reduction in accuracy. This may likely be attributed to the model's tendency to "hallucinate" in the face of the unfamiliar or uncertain.

Given the substantial improvements witnessed between major version upgrades, it is reasonable to expect that further experimentation with subsequent iterations may yield compelling results. For the moment, GPT, on the merits of its pre-training alone, is not effective in analyzing basic features of biblical Hebrew syntax.

1.2 Verbs in Jonah 1[7]

The subsequent tests center around the prompt, "In the original Hebrew, how many verbs are there in the first chapter of Jonah?" Similar to the previous evaluations, we engage models spanning versions 3.5 to 4, employing temperature settings of 0 and 1. These tests incorporate the same follow-up prompting techniques outlined in Section 1.1.

Results:

Table 2. Responses and performance on the prompt: In the original Hebrew, how many verbs are there in the first chapter of Jonah? (Correct 77)

GPT Model	Temp	Prompt Response	Certainty Response	Per Verse Response	Given Text Response	Accuracy
3.5	0	93	78	91		1 NC, 2 NC, 3 NC
3.5	1	102	106	137		1 NC, 2 NC, 3 NC
3.5 (0301)	0	53	63	52		1 NC, 2 NC, 3 NC
3.5 (0301)	1	122	70	54		1 NC, 2 NC, 3 NC
4 (0314)	0	69	64	64		1 NC, 2 NC, 3 NC
4 (0314)	1	69	61	54		1 NC, 2 NC, 3 NC
4 (current)	0	69	64	64	34	1 NC, 2 NC, 3 NC, 4 NC
4 (current)	1	65	65	70	35	1 NC, 2 NC, 3 NC, 4 NC

Interpretation of Results:

---

[7] Jonah was chosen for the following experiments because of its narrative framework, simplistic language, and minimal word count (which is an important consideration when concerning the token limitations of certain GPT models).

Once more, GPT's performance falls short, failing to render the correct response (77) across all models and temperature settings. Curiously, the closest approximation to the actual figure is generated by the second prompt of version 3.5. However, given the discrepancy in surrounding figures, this can likely be attributed to mere coincidence. The responses exhibit considerable variation across all tests, and in alignment with previous observations, when provided with the original Hebrew text, GPT-4 still fails to ascertain the correct figure. In light of these results, the initial conclusion stands unaltered: relying solely on its pre-training, GPT does not effectively analyze fundamental aspects of biblical Hebrew syntax.

1.3 Finding a verbal root

The subsequent task, marginally more intricate than the previous, involves identifying all verbs derived from a specific verbal root within a text. For this purpose, we continue with the book of Jonah, selected due to its relatively limited length and vocabulary. We evaluate this through the prompt, "Please find all verbs deriving from the verbal root יטב in the book of Jonah." The appropriate response would indicate that there are three instances: the initial instance is in 4:4, followed by two more in 4:9 (all of which are in the Hifil infinitive absolute form, הֵיטֵב). Unlike the previous experiments, only one follow-up prompt is given, "Are you certain?" The model was not given the entire text of Jonah for analysis due to token constraints.

Table 3. Responses and performance on the prompt: Please find all verbs deriving from the verbal root יטב in the book of Jonah.

GPT Model	Temp	Prompt Response	Accuracy	Certainty Response	Accuracy
3.5	0	1: 4:6, יִיטַב	NC	Apology, no verbs	NC
3.5	1	5: 4:6-7[X]	NC	4: 1:14, 1:16, 3:9, 4:1-2[XX]	NC
3.5 (0301)	0	8: all NC, no ref	NC	8, all NC, no ref	NC
3.5 (0301)	1	2: 3:10, 4:6	NC	refused	NC
4 (0314)	0	1: 4:4, הֵיטֵב	PC	confirmed	PC
4 (0314)	1	1: 4:4, הֵיטֵב	PC	2: 4:1, 4	PC
4 (current)	0	1: 4:4, הֵיטֵב	PC	1, 4:4, הֵיטֵב	PC
4 (current)	1	2: 4:4, 4:9, הֵיטֵב	PC	4: 4:4, 9, 10, 10	PC

[X]The exact response from the model: 1. יְצַמַּח - *yetsimach* - he caused to grow; 2. מְשֻׁוָּה - *mesuvah* - a plant that has grown up; 3. יָכַן - *yakhan* - he prepared; 4. נָכוֹן - *nakon* - it was ready, prepared; 5. יְכָּה - *yikah* - it struck, smote.

Interpretation of Results:

It is intriguing that the model exhibits improved performance under these qualitatively more challenging constraints, conditions which could arguably pose greater difficulty for a human analyst. While version 3.5 fails comprehensively, succumbing to hallucination with every attempted response, version 4 displays a degree of proficiency with morphology. Notably, GPT-4 renders the most accurate response at a temperature setting of 1, identifying two instances (4:4 and 4:9). Although both verse locations are correctly identified, the model misses accounting for two instances within the second verse. Nevertheless, this represents GPT's first demonstration of some analytical capability relating to biblical Hebrew syntax. Although currently inconsistent, further exploration with upcoming models promises to yield more persuasive outcomes.

The following brief experiments (1.4 through 1.6) represent an intersection between categories 1 and 2, knowledge retrieval and generative modeling. Here GPT was prompted to provide interpretations directed by the user such that a number of factors would need to be understood, interpreted, and connected.

1.4 Directed interpretation: Justifying Jonah's anger through decoloniality

The following exercise demonstrates GPT-4's ability to interpret an esoteric and biased request from the user. Decoloniality, the user's desired outcome (justification of an opinion), the story of Jonah, and economy of language (one paragraph requirement), all need to be parsed in order to competently produce the desired result.

Prompt:

Considering the story of Jonah, in one paragraph, justify Jonah's actions from a decolonial perspective.

Response:

From a decolonial perspective, Jonah's actions can be seen as an expression of resistance against the oppressive power dynamics between the Israelites and the Assyrians. His initial refusal to go to Nineveh and warn the people of their impending destruction can be understood as an attempt to protect his own people's identity and dignity, as well as to resist the idea of engaging with their oppressors in a way that might legitimize or humanize them. Jonah's reluctance to participate in the potential redemption of the Assyrians can be seen as a form of self-preservation and an assertion of his own people's agency in the face of a dominant and oppressive power. In this light, Jonah's actions can be interpreted as a manifestation of the struggle for autonomy and self-determination that is central to the decolonial project, highlighting the complexities and challenges of navigating power dynamics and cultural differences in a world marked by historical injustices and inequalities.

Interpretation of Results:

GPT-4 demonstrated remarkable performance in this case, delivering a nuanced and persuasive interpretation that met each criterion outlined in the prompt. This is at once encouraging and disconcerting. It is encouraging to witness the model's ability to apply interpretive techniques to a text, thereby potentially assisting users in brainstorming and evaluating their own interpretations. However, it is disquieting to observe how readily and effectively the model can—upon request— validate a user's preconceived bias (e.g., "justify Jonah's actions…"). While the bias justified in this case may not be inherently harmful or concerning, it is worth noting that the model has been trained to resist other attempts at promoting bias and misinformation. Nevertheless, the present circumstances provide certain types of intellectual dishonesty with an unbridled platform.

1.5 Guided interpretation: Queer hermeneutics interpretation of the story of Jonah

This exercise is intended to explore GPT-4's aptitude in deciphering a specialized request, yet this time without any specific directive to affirm a particular bias. The model is presented with a specific interpretative framework—queer hermeneutics—but the conclusions it draws within that framework are entirely its own. In this instance, the AI must interpret and apply a niche field of literary study to its understanding of the story of Jonah, all while adhering to the constraint of maintaining the response within a single paragraph.

Prompt:

> From the perspective of queer hermeneutics, provide a one-paragraph interpretation of the story of Jonah.

Response:

> From a queer hermeneutics perspective, the story of Jonah can be read as a narrative that challenges traditional norms and expectations, both in terms of the characters' actions and the broader themes of the story. Jonah, as a prophet, defies the conventional role assigned to him by attempting to avoid God's command to go to Nineveh, thus subverting the expectations of obedience and submission typically associated with prophetic figures. Additionally, the story's emphasis on God's mercy and compassion towards the Ninevites, who were considered enemies and oppressors, disrupts the binary opposition between "us" and "them," inviting readers to question and reevaluate their own assumptions about identity, morality, and the boundaries between different groups of people. In this way, the story of Jonah can be seen as a queer narrative that destabilizes established norms and encourages a more inclusive and fluid understanding of human relationships, divine interactions, and the transformative power of empathy and forgiveness.

Interpretation of Results:

In keeping with previous results, GPT-4 adeptly fulfills each facet of the prompt, demonstrating a nuanced comprehension of the text and the designated hermeneutic. The task here was not to justify a particular standpoint, but to interpret, with the only bias being that inherent in the provided framework. Despite the well-documented tendencies of even the most advanced generative AI models to exhibit hallucination or inadvertent misinformation, GPT-4 deftly addresses the request in these last two experiments, delivering responses that are both logical and cogent.

1.6 Guided interpretation: comparison and contrast

In a final test designed to further push the AI's capacities in (1) knowledge retrieval, (2) interpretation, and (3) creative problem-solving, the model is tasked with comparing and contrasting the previously discussed paragraphs. Notably, this is accomplished from memory, that is, the original texts are not included in the prompt, thus increasing the complexity of the task.

Prompt:

> In one paragraph, compare and contrast these different interpretations.

Response:

> The decolonial and queer hermeneutics interpretations of the story of Jonah offer unique insights, with the decolonial perspective focusing on Jonah's resistance to oppressive power dynamics between the Israelites and the Assyrians, emphasizing the struggle for autonomy and self-determination. In contrast, the queer hermeneutics interpretation reads the story as challenging traditional norms and expectations, both in terms of Jonah's defiance of his prophetic role and the broader themes of mercy and compassion towards the Ninevites, disrupting binary oppositions and encouraging a more inclusive and fluid understanding of human relationships. While both interpretations focus on the subversion of established norms and questioning of boundaries, the decolonial perspective highlights resistance in the face of oppression, whereas the queer hermeneutics interpretation emphasizes the destabilization of norms and the exploration of more inclusive understandings of identity and relationships.

Interpretation of Results:

The preceding analysis should underscore that GPT-4 possesses formidable interpretive capabilities. Although it is presently inadvisable to rely solely on the model's interpretations without subsequent critical review, the potential advantages it offers to scholars engaged in biblical and literary hermeneutics are evident.

2. Generative Modeling or Creative Problem Solving

2.1 Jonah Lacunae Experiment

LLMs, such as GPT, have text prediction as their foundational design and primary function. As such, these models possess a particular aptitude for tasks such as reconstructing missing or lost text, simulating alternative textual versions, calculating probabilities of textual variations, and amalgamating textual fragments. Each of these proficiencies directly corresponds to the type of tasks regularly undertaken by historians, archaeologists, and biblical exegetes, to name a few. The following experiments merely scratch the surface of the exploration of these capabilities utilizing readily accessible, non-specialized tools. Theoretically, the model's performance could be significantly enhanced through direct training and optimization.

The initial experiment aims to evaluate the fundamental ability of GPT's general pre-training to predict text in ancient Hebrew. In this instance, lacunae were artificially created by removing 15% of the text from Jonah 1, such that random clusters of one to three words were removed. All diacritical

markings, references, and line breaks were likewise removed from the text, in an attempt to replicate a scenario analogous to deciphering the Dead Sea Scrolls. The input text provided to GPT-4 was configured as follows:

ויהי דבר [1] אל יונה בן אמתי לאמר קום לך אל [2] העיר הגדולה וקרא עליה כי עלתה [3] לפני ויקם יונה לברח תרשישה מלפני [4] וירד יפו וימצא [5] באה תרשיש ויתן שׂכרה וירד בה לבוא עמהם תרשישה מלפני יהוה ויהוה הטיל רוח [6] אל הים ויהי סער [7] והאניה חשבה להשבר וייראו המלחים ויזעקו [8] אל אלהיו ויטלו את הכלים אשר[9] אל הים להקל מעליהם ויונה ירד אל ירכתי הספינה וישכב וירדם ויקרב [10] רב החבל ויאמר לו מה לך נרדם קום [11] אל אלהיך אולי יתעשת האלהים לנו ולא נאבד ויאמרו איש [12] לכו ונפילה גורלות ונדעה בשלמי הרעה הזאת לנו ויפלו גורלות ויפל [13] הגורל על יונה ויאמרו אליו הגידה נא [14] לנו באשר מה מלאכתך ומאין תבוא מה ארצך ואי מזה עם אתה ויאמר אליהם [15] אנכי עברי ואת יהוה אלהי השמים אני ירא אשר עשה את הים ואת היבשה וייראו האנשים יראה גדולה [16] ויאמרו אליו מה זאת עשית כי ידעו האנשים כי מלפני יהוה הוא ברח כי הגיד להם ויאמרו אליו מה [17] נעשה לך וישתק הים מעלינו כי הים הולך וסער ויאמר אליהם [18] שׂאוני והטילני אל הים וישתק הים מעליכם כי יודע אני כי בשלי הסער הגדול הזה עליכם ויחתרו האנשים [19] להשיב אל היבשה ולא יכלו כי הים הולך וסער עליהם ויקראו אל יהוה ויאמרו אנה יהוה אל [20] נאבדה בנפש האיש הזה ואל תתן עלינו דם נקיא כי אתה יהוה כאשר חפצת עשית וישאו את יונה ויטלהו אל הים [21] ויעמד הים מזעפו וייראו האנשים יראה גדולה [22] את יהוה ויזבחו זבח ליהוה וידרו נדרים

GPT-4 was subsequently instructed that each bracketed number represented an absence of text that could range from one to three words in length. It was prompted to make predictions adhering to these stipulations: (1) a temperature setting of zero should be employed, (2) it should rely exclusively on its pre-training and the intrinsic context of the passage for predictions, without making reference to similar documents, and (3) it should provide three predictions, spanning from most probable to least probable. The outcomes are laid out in Table 4 below. C1P refers to the scenario wherein the most probable choice was correct, C2P pertains to the second most probable choice, and C3P the third. "Partial" denotes an instance where the prediction was partially correct, while NC indicates that none of the predictions were accurate.

Table 4, produced by GPT-4 upon request following receipt of the key, represents the model's performance. With limited contextual information, the model was capable of accurately predicting the absent text, across all three probabilities, 64% of the time. The most probable prediction was accurate 50% of the time.

Table 4. GPT-4 generated table representing the model's predictions, assessing the accuracy of those predictions, and providing explanations for each.

No.	Predictions	Accuracy	Key	Explanation
1	ה' / יהוה / האלוהים	C2P	יהוה	The word יהוה is often used to refer to God in

				the Old Testament.
2	נינוה / ירושלים / בית-אל	C1P	נינוה	The story of Jonah is tied to the city of Nineveh, thus it was most probable.[X]
3	רעתם / זעקתם / פשעיהם	C1P	רעתם	In context, רעתם (their evil) fits as Nineveh was known for its wickedness.
4	יהוה / ה' / האלוהים	C1P	יהוה	As in first prediction, יהוה is frequently used to refer to God in the Old Testament.
5	אניה / ספינה / שפנה	C1P	אניה	Given the context of the sea and Jonah fleeing, אניה (ship) was the most likely choice.
6	גדולה / חזקה / סערה	C1P	גדולה	The phrase רוח גדולה (great wind) is common in biblical texts referring to a large storm.
7	גדול / חזק / כבד	Partial	גדול בים	Predicted גדול correctly, but missed בים. The complete phrase means "great in the sea".
8	איש / כל / כל איש	C1P	איש	איש (each man) fits contextually with sailors on the ship.
9	בידם / באניה / שבאניה	C2P	באניה	באניה (in the ship) is fitting considering the context of the scene on the ship.
10	אל / לאת / לקראת	NC	אליו	Incorrect prediction, אליו (to him) was not predicted.
11	קרא / התפלל / זעק	C1P	קרא	קרא (call) fits contextually with the sailors calling out to their gods.
12	אל רעהו / לרעהו / לחברו	C1P	אל רעהו	The phrase אל רעהו (to his friend) suits the context of the sailors speaking amongst themselves.
13	הגורל / גורל המשא / החטים	Partial	הגורל על יונה	Correctly predicted הגורל, but missed על יונה (on Jonah), which would specify the lot fell on Jonah.
14	לנו / מה / איפה	C1P	לנו	לנו (to us) fits contextually with the sailors questioning Jonah.
15	אני / אני הנה / אנכי	NC	עברי	Incorrect prediction, עברי (Hebrew) was not predicted.
16	את / ל / את כל	NC	ויאמרו אליו	Incorrect prediction, ויאמרו אליו (they said to him) was not predicted.
17	נעשה / נעשה לך / תעשה	C1P	נעשה	נעשה (we will do) fits contextually as the sailors ask Jonah what to do.

18	עלו / שאו / הרימו	NC	שׂאוני	Incorrect prediction, שׂאוני (lift me up) was not predicted.
19	לחזור / לשוב / להשיב	C3P	להשיב	להשיב (to return) was the least probable choice, but turned out to be correct.
20	אל / פן / שלא	NC	נא	Incorrect prediction, נא (please) was not predicted.
21	וישקט / וחדל / ונשקט	NC	ויעמד הים	Incorrect prediction, ויעמד הים (the sea stood) was not predicted.
22	את / ל / ב	C1P	את	את is a direct object marker in Hebrew, fitting in this context.

Most probable prediction correct = C1P; second most probable = C2P; third = C3P; partially correct = Partial; not correct = NC.
[X] It is important to note that these explanations were generated in the same prompt that generated the table. As such, they are separated one prompt from the initial predictions. For this reason, their validity should be questioned. The mention here of Jonah and the relevancy of Nineveh to that story seems to indicate that the model provided the explanations as an "afterthought."

Interpretation of Results:

These results showcase an area in which generative AI has a distinctive qualification. The most fundamental function of these models is the generation of text based on statistical predictions–a function arguably mirrored by scholars who deal with the lacunae of ancient texts, albeit with a given degree of knowledge and contextual understanding. However, GPT-4 has access to a far more extensive body of knowledge and broader scope of training than what a human researcher could possibly amass, and this underscores the significant potential of the technology to augment scholarship in these fields.

To further refine and better understand the results of this experiment, additional investigation is recommended that: (1) probes into the peculiar correct prediction by the AI of 'Nineveh', a proper noun, which does not appear elsewhere in the provided text (though the occurrence of the name 'Jonah' may hint at its presence), (2) investigates and quantifies the variation in outcomes with different temperature settings, and (3) replicates the experiment with an expanded context (for instance, inclusion of chapters 3 and 4 from the Book of Jonah to provide additional relevant context, maintaining the same principle of removing 15% of the text).

2.2 Authentic Lacunae Experiment: DSS, 4Q394 3-7 i

How might the pre-training of GPT-4 fare under more challenging and practical textual reconstruction circumstances? To examine these capacities (albeit without specific training or context), a transcription of a Dead Sea Scrolls (DSS) fragment was selected for analysis:

4Q394 3-7 i[8]

<div dir="rtl">

אלה מקצת דברינו [   ]ל שהם מ[קצת דברי]

[ה]מעשים שא א[נ]ח[נ]נו חושבים וכו]לם על]   [

וטהרת [   ]הר]   ]ועל תרומת ד]גן ה]גוים שהם מ . . ים]

ומגיע]י]ם בה א]ת   ]יחם ומט]מאים אותה ואין לאכול]

מדגן [הג]וים [ ]ואין [ ]לברא למק ש] ועל זבח החטאת

שהם מבשלים [אות]ה בכלי [נחושת ומ . . ים בה את]

בשר זבחיהם ומ]   ]ים בעזר]ה ומ . . ים אותה]

במרק זבחם ועל זבח הגרים [אנחנו חושבים שהם זובחים]

אל חי[   ]שא היא] כ]מי שזכת אליו [ואף על מנחת זבח]

השל[מים] שמניחים אותה מיום ליום ואף [כתוב]

שהמנ]חה נאכלת] ער החלבים והבשר ביום ו ב[חם כי לבני]

הכוהנ]ים] ראו להזה'ר בדבר הזה בשל שלוא י]היו]

מסיא[י]ם את העם עוון   ואף על טהרת פרת החטאת

השוחט אותה והסורף אותה והאוסף [א]ת אפרה והמזה את [מי]

החטאת לכול אלה להערי]בו]ת השמש להיות טהורים

בשל שא יהיה הטהר מזה על הטמה כי לבני

</div>

Alongside this text, GPT-4 was provided with the following prompt, structured to facilitate the model in approaching the task in a systematic manner:

I have a segment of ancient Hebrew text transcribed from the Dead Sea Scrolls. The transcription's format follows the original fragment's layout, with brackets indicating missing sections, or lacunae. The relative size and spacing of these brackets show the extent of the text lost. In some lacunae, you will see suggested reconstructions by scholars, marking their best estimates of the missing content. Note that a bracket's beginning or end could be within a word.

Given your expertise in ancient Hebrew and the provided text's internal context, could you please:

1. Suggest probable fill-ins for each blank lacuna. These should consider the required space within the bracket.
2. Evaluate the scholars' proposed reconstructions for lacunae. Do you agree with their suggestions, or could there be a more likely text?
3. Divide the results into lines, print the original line and its translation, and maintain the brackets.
4. Print the lines with your predictions, maintaining the brackets.

---

[8] Donald W. Parry and Emanuel Tov, eds., *The Dead Sea Scrolls Reader, Texts Concerned with Religious Law* (Leiden ; Boston: Brill, 2004), 326.

Response:

Original Text:

Line 1:

אלה מקצת דברינו [ ]ל שהם מ[קצת דברי]

"These are some of our words [ ] which are from [some of our words]."

Line 2:

[ה[מעשים שא א[נ]ח[נו חושבים וכו]לם על] [

"[The] actions that we [are] think[ing and so on] about [ ]."

Line 3:

וטהרת [ ]הר[ ] ועל תרומת ד[גן ה]גוים שהם מ . . ים]

"And purity of [ ] mount[and the contribution of the] garden of [the gentiles that they are from . .
]."

Line 4:

ומגיע[י]ם בה א[ת]יחם ומט[מאים אותה ואין לאכול]

"And touch[ing] her, you [ ] them and prep[aring it and not to eat]."

Line 5:

מדגן [הג]וים [ ] ואין]לברא למק ש] ועל זבח החטאת]

"From the grain [of the] Gentiles [and there isn't] to create a place [and upon the sin offering]."

Line 6:

שהם מבשלים [אות]ה בכלי [נחושת ומ . . ים בה את]

"That they are cooking [it] in a vessel [of copper and from . . bringing in it]."

Line 7:

בשר זבחיהם ומ[ ]ים בעזר[ה ומ . . ים אותה]

"The meat of their sacrifices and [ ] in help[ing and from . . bringing it]."

Line 8:

במרק זבחם ועל זבח הגרים [אנחנו חושבים שהם זובחים]

"In the broth of their sacrifice and upon the sacrifice of the strangers [we think that they are sacrific-
ing]."

Line 9:

אל חי[ן]שא היא[כ]מי שזכת אליו [ואף על מנחת זבח]

"To live[ ]that she is[like] water that has merited to him [and also upon the offering of a sacrifice]."

Line 10:

השל[מים] שמניחים אותה מיום ליום ואף [כתוב]

"The peace[offerings] that they lay daily and also [written]."

Line 11:

שהמנ[חה נאכלת] ער החלבים והבשר ביום ו ב[חם כי לבני]

"That the offering [is eaten] after the fats and the meat in the day and in the [heat because of the children]."

Line 12:

הכוהנ[ים] ראו להזה'ר בדבר הזה בשל שלוא י[היו]

"The priests [have] seen to be cautious about this matter because lest they [were]."

Line 13:

מסיא[י]ם את העם עוון ואף על טהרת פרת החטאת

"Causing [the] people to sin and even upon the purification of the red heifer of the sin offering."

Line 14:

השוחט אותה והסורף אותה והאוסף [א]ת אפרה והמזה את [מי]

"The one who slaughters it and the one who burns it and the one who gathers [its] ashes and the one who sprinkles [its water]."

Line 15:

החטאת לכול אלה להערי[בו]ת השמש להיות טהורים

"The sin offering for all these to be prepared [in] the sun to be pure."

Line 16:

בשל שא יהיה הטהר מזה על הטמה כי לבני

"Because there should be purity from this on the unclean, because of the sons."

Predicted Translation:

I will follow this up with the predicted translation in the next message.

Response continued:

Here are the translations of the predicted text:

Predicted Text:

Line 1:

אלה מקצת דברינו [שנתנו לכם ב]ל שהם מ[קצת דברי החכמה]

"These are some of our words [that we gave you in the] heart, which are from [some words of wisdom]."

Line 2:

[הכל]מעשים שא א[ם]ח[נו חושבים וכו' על מצוות התורה]

"[All] the deeds that we [are] think[ing and so on about the commandments of the Torah]."

Line 3:

וטהרת [הנפש מן ה]ר[ע ועל תרומת ד]גן ה[גוים שהם מבחרי המינים]

"And purity of [the soul from the] evil [and the contribution of the] garden of [the Gentiles that they are from the choicest of the species]."

Line 4:

ומגיע[ים בה א[ת [הטומאה מתוד]יחם ומט[מאים אותה ואין לאכול]

"And touch[ing] her, you [bring forth the impurity from within] them and prep[aring it and not to eat]."

Line 5:

מדגן [הג]וים [שאין להם תורה]לברא למק ש] ועל זבח החטאת]

"From the grain [of the] Gentiles [that they don't have Torah] to create a place [and upon the sin offering]."

Line 6:

שהם מבשלים [אות]ה בכלי [נחושת וממריחים בה את]

"That they are cooking [it] in a vessel [of copper and from spreading in it]."

Line 7:

בשר זבחיהם ומ[רחיקים מ]ים בעזר]ה ומסייעים אותה]

"The meat of their sacrifices and from [removing from] them with help[ing and assisting it]."

Line 8:

במרק זבחם ועל זבח הגרים [אנחנו חושבים שהם זובחים]

"In the broth of their sacrifice and upon the sacrifice of the strangers [we think that they are sacrific-ing]."

Line 9:

אל חי[ות של קודש[שא היא] כ]מי שזכת אליו [ואף על מנחת זבח]

"To live[lyhood of holiness] that she is[like] water that has merited to him [and also upon the offer-ing of a sacrifice]."

Line 10:

השל[מים] שמניחים אותה מיום ליום ואף [כתוב בתורה]

"The peac[e offerings] that they place it from day to day and also [it is written in the Torah]."

Line 11:

שהמנ[חה נאכלת] ער החלבים והבשר ביום ו ב[חם כי לבני ישראל]

"That the meal [offering is eaten] after the fats and the meat in the day and in the [heat because of the children of Israel]."

Line 12:

הכוהנ[ים] ראו להזה'ר בדבר הזה בשל שלוא י]היו מטעים בתורה]

"The priests [have] seen to be cautious about this matter because lest they [were mistaken in the To-rah]."

Line 13:

מסיא[י]ם את העם עוון ואף על טהרת פרת החטאת

"Causing [the] people to sin and even upon the purification of the red heifer of the sin offering."

Line 14:

השוחט אותה והסורף אותה והאוסף [א]ת אפרה והמזה את [מי המקדש]

"The one who slaughters it and the one who burns it and the one who gathers [its] ashes and the one who sprinkles [its sanctification water]."

Line 15:

החטאת לכול אלה להערי[בו]ת השמש להיות טהורים

"The sin offering for all these to be prepared [in] the sun to be pure."

Line 16:

בשל שא יהיה הטהר מזה על הטמה כי לבני

"Because there should be purity from this on the unclean, because of the sons."

For reference, the translation of the original by E. Qimron and J. Strugnell is provided below:

1. These are some of our rulings [ ] which are [some of the rulings according to]
2. [the] precepts (of the Torah) in accordance with [our opinion, and] all of them concern [ ]
3. and the purity of [the And concerning the sowed gifts of the] new wheat grains of the [gentiles which they]
4. and let their [ ] touch it and de[file it, and no one should eat]
5. any of the new wheat grains of t[he gen]tiles, [nor] should it be brought into the sanctuary. [And concerning the sacrifice of the purification-offering]
6. that they cook in a [copper] vessel [and that they] in it
7. the flesh of their sacrifices, and that they [ ] in the Temple court (?) [and that they] it
8. with the broth of their sacrifices. And concerning the sacrifice of the Gentiles: [we are of the opinion that they] sacrifice
9. to the [ ] that it is like (a woman) who whored with him. [And concerning the cereal-offering] of the sacrifice
10. of well-being which they (the opponents) leave over from one day to the following one: but [it is written]
11. that the cereal-offer[ing is to be ea]ten after the suet and the flesh (are sacrificed), on the day when they are sacri[ficed (i.e. before sunset). For the sons of]
12. the priest[s] should take care concerning this practice so as not to
13. cause the people to bear punishment. And concerning the purity-regulations of the cow of the purification-offering (i.e. the red cow):
14. he who slaughters it and he who burns it and he who gathers its ashes and he who sprinkles the [water of]
15. purification—it is at sun[se]t that all these become pure
16. so that the pure man may sprinkle upon the impure one. For the sons of[9]

Undeniably, the translation lacks the precision and subtlety that human scholarship can offer. Additionally, while the predictions are intriguing (and at times, amusing), they would not be accepted on their own merit within the scholarly community dedicated to Dead Sea Scrolls studies. It is, however, noteworthy that these predictions were generated without the advantage of an extensive, relevant contextual corpus. Hypothetically, if GPT-4 were to incorporate the Dead Sea Scrolls corpus in an annotated format (such as the ETCBC's DSS dataset),[10] this, alongside its pre-training, should yield more remarkable results. Therefore, to evaluate the utility of GPT-4 (and subsequent versions) for these specialized areas of scholarship, future experiments should be designed such that: (1) GPT-4, or a later version, is trained on a relevant textual corpus pertaining to the text with the lacunae under

---

[9] Ibid., 327.

[10] https://github.com/ETCBC/dss

examination, and (2) results are aggregated and analyzed across various temperature settings and multiple probabilities for each predicted text, and (3) optionally, the image parsing capabilities of the model are utilized to examine and interpret high-resolution scans of the actual fragments.

2.3 Creating a coherent passage from two different passages

To further assess the creative problem-solving abilities of GPT-4, two passages from the Hebrew Bible were randomly selected. For each passage, all diacritical markings have been removed, and no consideration was given to where the text should terminate. Passage 1 is part of Deuteronomy 1:1-13; passage 2 is part of Ruth 1:1-8. The AI was prompted to create a new, unique, and coherent passage using only words from Passages 1 and 2.

Prompt:

Using only words from the two passages below, create a coherent text in the biblical Hebrew language.

Passage 1:

אלה הדברים אשר דבר משה אל כל ישראל בעבר הירדן במדבר בערבה מול סוף בין פארן ובין תפל ולבן וחצרת ודי זהב אחד עשר יום מחרב דרך הר שעיר עד קדש ברנע ויהי בארבעים שנה בעשתי עשר חדש באחד לחדש דבר משה אל בני ישראל ככל אשר צוה יהוה אתו אלהם אחרי הכתו את סיחן מלך האמרי אשר יושב בחשבון ואת עוג מלך הבשן אשר יושב בעשתרת באדרעי בעבר הירדן בארץ מואב הואיל משה באר את התורה הזאת לאמר יהוה אלהינו דבר אלינו בחרב לאמר רב לכם שבת בהר הזה פנו וסעו לכם ובאו הר האמרי ואל כל שכניו בערבה בהר ובשפלה ובנגב ובחוף הים ארץ הכנעני והלבנן עד הנהר הגדל נהר פרת ראה נתתי לפניכם את הארץ באו ורשו את הארץ אשר נשבע יהוה לאבתיכם לאברהם ליצחק וליעקב לתת להם ולזרעם אחריהם ואמר אלכם בעת ההוא לאמר לא אוכל לבדי שאת אתכם יהוה אלהיכם הרבה אתכם והנכם היום ככוכבי השמים לרב יהוה אלהי אבותכם יסף עליכם ככם אלף פעמים ויברך אתכם כאשר דבר לכם איכה אשא לבדי טרחכם ומשאכם וריבכם הבו לכם אנשים חכמים ונבנים וידעים לשבטיכם ואשימם בראשיכם ותענו אתי

Passage 2:

ויהי בימי שפט השפטים ויהי רעב בארץ וילך איש מבית לחם יהודה לגור בשדי מואב הוא ואשתו ושני בניו ושם האיש אלימלך ושם אשתו נעמי ושם שני בניו מחלון וכליון אפרתים מבית לחם יהודה ויבאו שדי מואב ויהיו שם וימת אלימלך איש נעמי ותשאר היא ושני בניה וישאו להם נשים מאביות שם האחת ערפה ושם השנית רות וישבו שם כעשר שנים וימותו גם שניהם מחלון וכליון ותשאר האשה משני ילדיה ומאישה ותקם היא וכלתיה ותשב משדי מואב כי שמעה בשדה מואב כי פקד יהוה את עמו לתת להם לחם ותצא מן המקום אשר היתה שמה ושתי כלתיה עמה ותלכנה בדרך לשוב אל ארץ יהודה ותאמר נעמי לשתי כלתיה לכנה שבנה אשה לבית אמה יעשה יהוה עמכם חסד כאשר עשיתם עם

Response:

האלהים אשר דבר אל משה, דיבר אל כל ישראל, בארץ האמרי והכנעני. איש מבית לחם יהודה הלך לגור

בארץ אחרת. יהוה, אלהינו, דבר אלינו בחרב, ראה נתתי לפניכם את הארץ. באו ורשו את הארץ הכנעני
והלבנון, יהוה נשבע לאברהם, ליצחק וליעקב לתת להם ולזרעם אחריהם. ויהיו שם בימי שפט השופטים, עם
עלה מארץ מואב, כי שמעה שם יהוה פקד את עמו לתת להם לחם. ויברך אתכם כאשר דבר לכם. אלה הדברים
אשר צוה יהוה אתו להם.

The God who spoke to Moses spoke to all Israel in the land of the Amorites and the Canaanites. A man from Bethlehem of Judah went to live in another land. YHWH our God spoke to us at Horeb, saying, "I have set the land before you. Come and possess the land of the Canaanites and Lebanon, which YHWH swore to Abraham, Isaac, and Jacob to give to them and to their descendants after them." They were there in the days of the judges, with people who had ascended from Moab, for they had heard that YHWH had remembered His people to give them bread. He will bless you as He has spoken to you. These are the words that YHWH commanded for them.

Utilizing only the words found in Passages 1 and 2, GPT-4 successfully generated a unique and largely coherent piece of text. Although the practical application of such a task may not be immediately evident, the outcome effectively demonstrates GPT-4's proficiency in Biblical Hebrew and its capacity to employ it creatively in problem-solving scenarios involving multifaceted forms of reasoning.

3. Command interpretation or query parsing

Recognizing the current limitations of the model–particularly its inability to accurately interrogate biblical Hebrew syntax solely based on its pre-training–GPT-4 was employed in conjunction with top-tier linguistic tools, specifically the Hebrew Bible in the form of the BHSA dataset, and the Text-Fabric Python library, to conduct the following experiments. As the forthcoming results will illustrate, this combination of technologies yields accurate responses to high-level, plain-language queries that would likely be of interest to scholars of biblical Hebrew.

3.1 In line with the 1.1 experiment, we prompt GPT-4 to compose a code snippet that leverages the Text-Fabric Python library to query the BHSA dataset.

Prompt:

Please write a Python code snippet that uses the Text-Fabric Python library to query the BHSA and count the number of occurrences of the "וְ" used as a conjunction in the first chapter of Genesis.

Response:

```python
from tf.app import use

A = use('bhsa', hoist=globals())[11]

# Define the target conjunction Lemma
target_conjunction_lemma = "W"

# Get the node for the first chapter of Genesis
```

---

[11] This portion of the code will be removed from all future code snippets.

```
genesis_1_node = T.nodeFromSection(('Genesis', 1))

# Get all the word nodes in the first chapter of Genesis
word_nodes = L.d(genesis_1_node, otype='word')

# Count the occurrences of the target conjunction lemma
conjunction_count = 0
for word_node in word_nodes:
    if F.sp.v(word_node) == 'conj' and F.lex.v(word_node) == target_conjunction_lemma:
        conjunction_count += 1

print(f"The waw conjunction (lemma '{target_conjunction_lemma}') occurs {conjunction_count} times in
Genesis 1 according to the BHSA.")
```

Code Output:

> The waw conjunction (lemma 'W') occurs 101 times in Genesis 1 according to the BHSA.

Interpretation:

Since only a small body of text was queried, a manual count of the conjunctive "W" was sufficient to confirm the accuracy of the code output. But it might be helpful to modify the code such that there is a count for each of the 31 verses.

3.1a Prompting for further verification of conjunctive "W" count

Prompt:

> Please modify this code such that the count is performed and reported for each of the 31 verses.
> Print a table to illustrate this.

Response:

```
# Define the target conjunction lemma
target_conjunction_lemma = "W"

# Get the node for the first chapter of Genesis
genesis_1_node = T.nodeFromSection(('Genesis', 1))

# Get all the verse nodes in the first chapter of Genesis
verse_nodes = L.d(genesis_1_node, otype='verse')

# Print the table header
print(f"{'Verse':<10}{'Waw Conjunctions':<20}")
print(f"{'-'*10:<10}{'-'*20:<20}")

# Count and print the occurrences of the target conjunction lemma for each verse
for verse_node in verse_nodes:
    verse_label = T.sectionFromNode(verse_node, lang='en', lastSlot=False)
    word_nodes = L.d(verse_node, otype='word')
    conjunction_count = 0
    for word_node in word_nodes:
        if F.sp.v(word_node) == 'conj' and F.lex.v(word_node) == target_conjunction_lemma:
            conjunction_count += 1
    print(f"{verse_label:<10}{conjunction_count:<20}")
```

Output: Error[12]

Prompt:

An error was returned:

TypeError   Traceback (most recent call last)
Cell In[7], line 26
     24   if F.sp.v(word_node) == 'conj' and F.lex.v(word_node) == target_conjunction_lemma:
     25   conjunction_count += 1
---> 26 print(f"{verse_label:<10}{conjunction_count:<20}")

TypeError: unsupported format string passed to tuple.__format__

Please correct the error and modify the code for the table such that there are 3 columns: verse, the count, and a running count. Find a Python library that can output a simple table to clearly illustrate this data.

Response:

```python
from tf.app import use
from tabulate import tabulate

A = use('bhsa', hoist=globals())

# Define the target conjunction lemma
target_conjunction_lemma = "W"

# Get the node for the first chapter of Genesis
genesis_1_node = T.nodeFromSection(('Genesis', 1))

# Get all the verse nodes in the first chapter of Genesis
verse_nodes = L.d(genesis_1_node, otype='verse')

# Initialize the table data and running count
table_data = []
running_count = 0

# Count the occurrences of the target conjunction lemma for each verse
for verse_node in verse_nodes:
    verse_label = T.sectionFromNode(verse_node, lang='en', lastSlot=False)
    verse_label_str = f"{verse_label[0]} {verse_label[1]}:{verse_label[2]}"
    word_nodes = L.d(verse_node, otype='word')
    conjunction_count = 0
    for word_node in word_nodes:
        if F.sp.v(word_node) == 'conj' and F.lex.v(word_node) == target_conjunction_lemma:
            conjunction_count += 1
    running_count += conjunction_count
    table_data.append([verse_label_str, conjunction_count, running_count])

# Print the table using the tabulate library
print(tabulate(table_data, headers=['Verse', 'Count', 'Running Count'], tablefmt='grid'))
```

---

[12] In the course of these experiments, certain code occasionally produced errors. All of these were addressed uniformly: the error was copied verbatim, without alteration, explanation, or contextualization, into a prompt stating, "I received this error: [error message]." Subsequently, the model would diagnose its mistake and provide rectified code. On rare occasions where the code still malfunctioned, this process was reiterated until successful execution was achieved. Notably, no user intervention was required during the error correction process beyond relaying the error to the model–a process that could feasibly be automated.

Code Output:

Table 5. Experiment 3.1a code generated table.

Verse	Count	Running Count
Genesis 1:1	1	1
Genesis 1:2	4	5
Genesis 1:3	2	7
Genesis 1:4	3	10
Genesis 1:5	4	14
Genesis 1:6	2	16
Genesis 1:7	4	20
Genesis 1:8	3	23
Genesis 1:9	3	26
Genesis 1:10	3	29
Genesis 1:11	2	31
Genesis 1:12	3	34
Genesis 1:13	2	36
Genesis 1:14	6	42
Genesis 1:15	2	44
Genesis 1:16	3	47
Genesis 1:17	1	48
Genesis 1:18	5	53
Genesis 1:19	2	55
Genesis 1:20	2	57
Genesis 1:21	4	61
Genesis 1:22	4	65
Genesis 1:23	2	67
Genesis 1:24	4	71
Genesis 1:25	4	75

Genesis 1:26	6	81
Genesis 1:27	2	83
Genesis 1:28	8	91
Genesis 1:29	2	93
Genesis 1:30	4	97
Genesis 1:31	4	101
Note: this table is a representation following the exact structure of the table produced by the code.		

Interpretation:

When presented with this new task, the modified code produced by GPT-4 returned an error. By merely copying and pasting the error message, the model was able to diagnose the issue and rectify the code to ensure proper execution. Moreover, the model was given instructions that required an additional layer of interpretation. With a hypothetical understanding of what the data might resemble, the model was tasked with selecting a Python library that could present the results in a simple table. The choice of the "tabulate" library and its subsequent inclusion in the code produced an ideal representation.

The secondary objective of this code was to validate the earlier estimate from section 3.1. The results offered an accurate tally that was verified by a manual count of the occurrences in each verse.

3.2 Pursuant to the 1.3 experiment, we prompt GPT-4 to produce another code snippet.

Prompt:

Please find all words deriving from the verbal root יטב in the book of Jonah. Print out the reference and the word for each.

Response:

```python
lexeme = 'יטב'  # The lexeme for יטב in Hebrew script

# Find all words in the book of Jonah with the specified lexeme
words = [word for word in F.otype.s('word')
        if lexeme in F.g_cons_utf8.v(word) and T.sectionFromNode(word)[0] == 'Jonah']

print(f"The lexeme יטב occurs {len(words)} times in the book of Jonah.")

# For each word
for word in words:
    # Get the reference (book, chapter, verse)
    reference = T.sectionFromNode(word)
    # Get the actual Hebrew word
    word_hebrew = F.g_word_utf8.v(word)
    print(f'Reference: {reference}, Word: {word_hebrew}')
```

Code Output:

The lexeme יטב occurs 3 times in the book of Jonah.
Reference: ('Jonah', 4, 4), Word: הֵיטֵב
Reference: ('Jonah', 4, 9), Word: הֵיטֵב
Reference: ('Jonah', 4, 9), Word: הֵיטֵב

Interpretation:

This output is correct. After an attempt that returned 0 results, GPT-4 was given a second prompt, "That produced 0 results." The model then arrived at a code snippet that accurately satisfied the original prompt's requirements.

This process of trial and error was occasionally required to effectively interface GPT-4 with Text-Fabric. A solution would be to have the model ingest all documentation relating to Text-Fabric and the BHSA dataset. Such a project is desirable but is outside of the scope of this introductory overview.

3.2a Expanding the experiment to encompass the whole of the BHSA dataset

Understanding that the model is able to build on the insights gained from the previous experiment regarding the successful identification of words formed from a specific verbal root, we are able broaden the scope of our inquiry to encompass all text within the BHSA. Additionally, we prompt the model to produce a code snippet that incorporates a requirement for a verbal stem, yields conjugation data, and generates a visualization pertinent to the data, the choice of which is left to the AI's discretion.

Prompt:

Can we run a modified search (of our last prompt) that counts every instance of the words derived from that root, which are also in the Hifil stem? Include the reference, word, and conjugation data. Additionally, create a visualization that effectively illustrates this data.

Response:

Certainly, for visualization, we could create a bar chart showing the distribution of the word in different books of the Bible. We will use the matplotlib library for this purpose.

Here is the updated code:

```python
import matplotlib.pyplot as plt

lexeme = 'יטב'  # The lexeme for יטב in Hebrew script

# Find all words with the specified lexeme and stem across the entire BHSA
words = [word for word in F.otype.s('word')
        if lexeme in F.g_cons_utf8.v(word) and F.vs.v(word) == 'hif']

print(f"The lexeme יטב occurs {len(words)} times in the Hifil stem in the entire BHSA.")
```

```python
# Counter for book occurrences
book_counts = Counter()

# For each word
for word in words:
    # Get the reference (book, chapter, verse)
    reference = T.sectionFromNode(word)
    book_counts[reference[0]] += 1  # Increase the count for this book
    # Get the actual Hebrew word
    word_hebrew = F.g_word_utf8.v(word)
    # Get the stem (binyan)
    stem = F.vs.v(word)
    # Get the tense
    tense = F.vt.v(word)
    # Get the person
    person = F.ps.v(word)
    # Get the gender
    gender = F.gn.v(word)
    # Get the number
    number = F.nu.v(word)

    print(f"Reference: {reference}, Word: {word_hebrew}, Stem: {stem}, Tense: {tense}, Person: {person}, Gender: {gender}, Number: {number}")

# Plotting
books = list(book_counts.keys())
counts = list(book_counts.values())
plt.figure(figsize=(10, 5))  # Increase the size of the figure
plt.bar(books, counts)
plt.title('Occurrences of יטב in Hifil stem in each book')
plt.xlabel('Books')
plt.ylabel('Occurrences')
plt.xticks(rotation=90)  # Rotate the x-axis labels for better readability
plt.show()
```

Code Output:

> The lexeme יטב occurs 24 times in the Hifil stem in the entire BHSA.
> Reference: ('Genesis', 32, 13), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown
> Reference: ('Exodus', 1, 20), Word: יֵּיטֶב, Stem: hif, Tense: wayq, Person: p3, Gender: m, Number: sg
> Reference: ('Deuteronomy', 8, 16), Word: הֵיטִבְךָ, Stem: hif, Tense: infc, Person: unknown, Gender: unknown, Number: unknown
> Reference: ('Deuteronomy', 9, 21), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown
> Reference: ('Deuteronomy', 13, 15), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown
> Reference: ('Deuteronomy', 17, 4), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown
> Reference: ('Deuteronomy', 19, 18), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown
> Reference: ('Deuteronomy', 27, 8), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown
> Reference: ('Deuteronomy', 30, 5), Word: הֵיטִבְךָ, Stem: hif, Tense: perf, Person: p3, Gender: m, Number: sg
> Reference: ('1_Samuel', 20, 13), Word: יֵיטִב, Stem: hif, Tense: impf, Person: p3, Gender: m, Number: sg
> Reference: ('1_Samuel', 25, 31), Word: הֵיטִב, Stem: hif, Tense: perf, Person: p3, Gender: m, Number: sg
> Reference: ('1_Kings', 1, 47), Word: יֵיטֵב, Stem: hif, Tense: impf, Person: p3, Gender: m, Number: sg

Reference: ('2_Kings', 9, 30), Word: תֵּיטֶב, Stem: hif, Tense: wayq, Person: p3, Gender: f, Number: sg

Reference: ('2_Kings', 11, 18), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown

Reference: ('Isaiah', 1, 17), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown

Reference: ('Jeremiah', 1, 12), Word: הֵיטַבְתָּ, Stem: hif, Tense: perf, Person: p2, Gender: m, Number: sg

Reference: ('Jeremiah', 2, 33), Word: תֵּיטְבִי, Stem: hif, Tense: impf, Person: p2, Gender: f, Number: sg

Reference: ('Jonah', 4, 4), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown

Reference: ('Jonah', 4, 9), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown

Reference: ('Jonah', 4, 9), Word: הֵיטֵב, Stem: hif, Tense: infa, Person: unknown, Gender: unknown, Number: unknown

Reference: ('Proverbs', 15, 13), Word: יֵיטָב, Stem: hif, Tense: impf, Person: p3, Gender: m, Number: sg

Reference: ('Proverbs', 17, 22), Word: יֵיטָב, Stem: hif, Tense: impf, Person: p3, Gender: m, Number: sg

Reference: ('Proverbs', 30, 29), Word: מֵיטִבֵי, Stem: hif, Tense: ptca, Person: unknown, Gender: m, Number: pl

Reference: ('Ruth', 3, 10), Word: הֵיטַבְתְּ, Stem: hif, Tense: perf, Person: p2, Gender: f, Number: sg

Occurrences of יטב in Hifil stem in each book
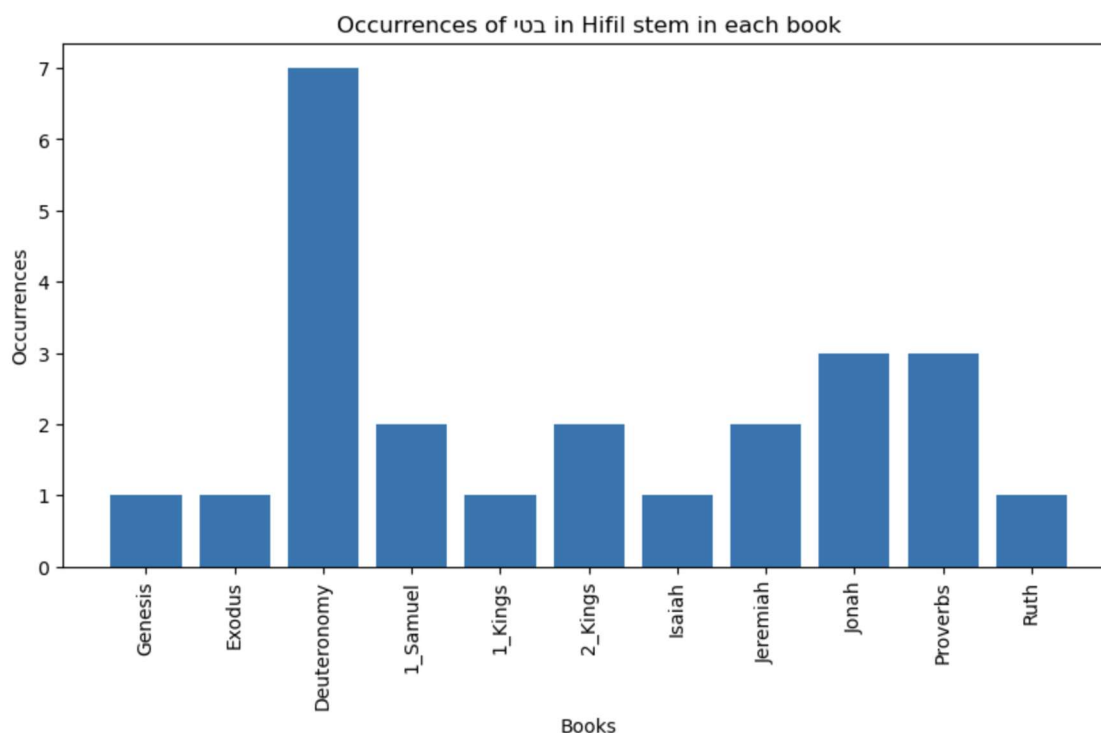
Interpretation:

In this instance, GPT-4 adeptly modifies the preceding code to locate the specific features across the entire BHSA dataset and makes an informed decision to generate a useful visualization of the data. The code was flawlessly executed on the initial attempt. The success of this plain language approach,

in the given context, enables the human researcher to transition from being a mere user to a collaborator, actively contributing towards the research objective.

3.3 The AlefBet and the exploration word histories[13]

Hebrew letters have their origins in the ancient Paleo-Hebrew script, which itself evolved from the Phoenician alphabet, a Semitic system of writing that utilized pictographic symbols to represent words. These pictograms held symbolic meanings and phonetic values. Over time, the Paleo-Hebrew script evolved into the more abstract, standardized script we recognize as the Hebrew/Aramaic alphabet today. Despite these changes, traces of the original pictographic meanings are preserved in the names and shapes of some Hebrew letters. As such, the Hebrew alphabet not only serves as a linguistic tool but also a historical artifact that provides a glimpse into the semantic, cultural, and visual development of language in the ancient Near East.

The forthcoming experiment delves into how GPT-4 can assist a researcher in a multi-step process designed to generate data conducive to investigating the historical development of Hebrew letters, their names, and their Paleo-Hebrew counterparts. For the sake of brevity, the process steps are outlined instead of providing an exhaustive list of each prompt and response. Importantly, every step of this process is conducted entirely as a "conversation" with GPT-4.

The initial step entailed requesting the AI to provide a list of Hebrew letters along with their names. This elicited the following response:

(אָלֶף) א, (בֵּית) ב, (גִּימֵל) ג, (דָּלֶת) ד, (הֵא) ה, (וָו) ו, (זַיִן) ז, (חֵת) ח, (טֵת) ט, (יוֹד) י, (כַּף) כ, (לָמֶד) ל, (מֵם) מ, (נוּן) נ, (סָמֶךְ) ס, (עַיִן) ע, (פֵּא) פ, (צַדִּי) צ, (קוֹף) ק, (רֵישׁ) ר, (שִׁין) ש, (תָּו) ת.

Afterward, the site containing the ETCBC's transcriptions (which are used by the BHSA) was shared with the AI alongside instructions to transcribe the Hebrew consonants of the names.[14] This produced the following result:

>LP, BJT, GML, DLT, H>, WW, ZJN, XJT, VJT, JWD, KF, LMD, MM, NWN, SMK, <JN, P>, YDY, QWF, RJ#, #JN, TW

The model was then tasked with writing code to find words related to these consonantal combinations, to provide the lex and gloss for any words returned, and to add the Hebrew word and a count of occurrences to the list. One hint was given: word lex0=>LP lex* gloss*.[15]

These instructions, with two instances of GPT-4 directed error correction, produced the following code:

```
transliterations = [
    ">LP", "BJT", "GML", "DLT", "H>", "WW", "ZJN", "XJT", "VJT", "JWD",
    "KF", "LMD", "MM", "NWN", "SMK", "<JN", "P>", "YDY", "QWF", "RJ#", "#JN", "TW"
]
```

---

[13] Credit for the conception of this experiment belongs to Tim Nash, and the rough outline for its execution to Dirk Roorda.

[14] Ibid., Text-Fabric.

[15] This *hint* was not strictly necessary as the model, through a process of trial and error, could determine the right syntax. Moreover, future experiments that allow the model to ingest the necessary documentation should produce more efficient and accurate code. The "hint" here is provided primarily for expediency.

```python
for transliteration in transliterations:
    query = f'''
word lex0={transliteration} lex gloss g_word_utf8
'''
    results = A.search(query, silent=True)

    if results:
        lex_node = results[0][0]
        lex = F.lex.v(lex_node)
        gloss = F.gloss.v(lex_node)
        hebrew_text = F.g_word_utf8.v(lex_node)
        print(f'{transliteration}: {lex} - {gloss}, Hebrew: {hebrew_text}, Count: {len(results)}')
    else:
        print(f'{transliteration}: Not found')
```

Code Output:

```
>LP: >LP=/ - thousand, Hebrew: אֶלֶף, Count: 523
BJT: BJT/ - house, Hebrew: בֵּית, Count: 2108
GML: GML/ - camel, Hebrew: גְמַלִּים, Count: 91
DLT: DLT/ - door, Hebrew: דֶּלֶת, Count: 87
H>: H> - behold, Hebrew: הֵא, Count: 4
WW: WW/ - hook, Hebrew: וָוֵיהֶם, Count: 13
ZJN: Not found
XJT: Not found
VJT: Not found
JWD: Not found
KF: Not found
LMD: LMD[ - learn, Hebrew: מְלַמֵּד, Count: 92
MM: Not found
NWN: NWN/ - Nun, Hebrew: נוּן, Count: 30
SMK: SMK[ - support, Hebrew: סְמַכְתִּיו, Count: 48
<JN: <JN/ - eye, Hebrew: עֵינֵיכֶם, Count: 896
P>: Not found
YDY: Not found
QWF: Not found
RJ#: Not found
#JN: Not found
TW: TW/ - mark, Hebrew: תָּו, Count: 3
```

While this is interesting, more could be added to illustrate the point. The model was subsequently tasked with providing a list of the Paleo-Hebrew (Phoenician) pictograms in Unicode from its own pre-training. This generated:

𐤕 ,𐤔 ,𐤓 ,𐤒 ,𐤗 ,𐤐 ,𐤏 ,𐤎 ,𐤍 ,𐤌 ,𐤋 ,𐤊 ,𐤆 ,𐤈 ,𐤇 ,𐤅 ,𐤄 ,𐤃 ,𐤂 ,𐤁 ,𐤀

Armed with this information, the model was next instructed to alter the preceding code to list the Paleo-Hebrew pictograms alongside the corresponding letters of the Hebrew alphabet, ensuring that the new fields were aligned with the relevant transliterated consonants. Ultimately, to enhance the presentation of this data, the AI was also directed to generate a table of the results in HTML format. The code produced from these instructions is as follows:

```python
transliterations = [
    ">LP", "BJT", "GML", "DLT", "H>", "WW", "ZJN", "XJT", "VJT", "JWD",
    "KF", "LMD", "MM", "NWN", "SMK", "<JN", "P>", "YDY", "QWF", "RJ#", "#JN", "TW"
]

# Hebrew letters in order
hebrew_letters = [
    "א", "ב", "ג", "ד", "ה", "ו", "ז", "ח", "ט", "י",
    "כ", "ל", "מ", "נ", "ס", "ע", "פ", "צ", "ק", "ר", "ש", "ת"
]

# Phoenician pictograms in order
phoenician_pictograms = [
    "𐤀", "𐤁", "𐤂", "𐤃", "𐤄", "𐤅", "𐤆", "𐤇", "𐤈", "𐤉",
    "𐤊", "𐤋", "𐤌", "𐤍", "𐤎", "𐤏", "𐤐", "𐤑", "𐤒", "𐤓", "𐤔", "𐤕"
]

# Create DataFrame
df = pd.DataFrame(columns=["Pictogram", "Hebrew Letter", "Hebrew Word", "Lex", "Transcription",
"Gloss", "Count"])
data_list = []

for i, transliteration in enumerate(transliterations):
    query = f'''
word lex0={transliteration} lex gloss g_word_utf8
'''
    results = A.search(query, silent=True)

    hebrew_letter = hebrew_letters[i]
    phoenician_pictogram = phoenician_pictograms[i]

    if results:
        lex_node = results[0][0]
        lex = F.lex.v(lex_node)
        gloss = F.gloss.v(lex_node)
        hebrew_text = F.g_word_utf8.v(lex_node)
        data_list.append({"Pictogram": phoenician_pictogram, "Hebrew Letter": hebrew_letter, "Hebrew
Word": hebrew_text, "Lex": lex, "Transcription": transliteration, "Gloss": gloss, "Count": len(re-
sults)})
    else:
        data_list.append({"Pictogram": phoenician_pictogram, "Hebrew Letter": hebrew_letter, "Hebrew
Word": None, "Lex": None, "Transcription": transliteration, "Gloss": None, "Count": 0})

df = pd.concat([pd.DataFrame([i]) for i in data_list], ignore_index=True)

# Save DataFrame to HTML
df.to_html("output.html", index=False)
```

Code Output:

Table 6. Experiment 3.3 code generated table.

Pictogram	Hebrew Letter	Hebrew Word	Lex	Transcription	Gloss	Count
𐤀	א	אֶלֶף	>LP=/	>LP	thousand	523
𐤁	ב	בַּיִת	BJT/	BJT	house	2108
𐤂	ג	גְמַלִּים	GML/	GML	camel	91
𐤃	ד	דֶּלֶת	DLT/	DLT	door	87

�-glyph	Hebrew	Word	Lex1	Lex2	Gloss	Count
ꓱ	ה	הֵא	H>	H>	behold	4
ꓬ	ו	וֲוֵיהֶם	WW/	WW	hook	13
ꓶ	ז	None	None	ZJN	None	0
ꓭ	ח	None	None	XJT	None	0
⊕	ט	None	None	VJT	None	0
ꓜ	י	None	None	JWD	None	0
ꓘ	כ	None	None	KF	None	0
ꓥ	ל	מְלַמֵּד	LMD[	LMD	learn	92
ꓦ	מ	None	None	MM	None	0
ꓨ	נ	נוּן	NWN/	NWN	Nun	30
ꓭ	ס	סְמַכְתָּיו	SMK[	SMK	support	48
O	ע	עֵינֵיכֶם	<JN/	<JN	eye	896
ꓶ	פ	None	None	P>	None	0
ꓤ	צ	None	None	YDY	None	0
φ	ק	None	None	QWF	None	0
ꓭ	ר	None	None	RJ#	None	0
ꓪ	ש	None	None	#JN	None	0
ꓕ	ת	תָּו	TW/	TW	mark	3

Several fields returned no results, for which the Ai supplied the word "None." This appears to be due to the limits of the biblical Hebrew lexicon. While some of these words can be found in sister languages or even Modern Hebrew (e.g., ZJN), they are not found in the BHSA. Therefore, no corresponding Hebrew word, lex, nor gloss was returned.

Interpretation:

This experiment demonstrates how a scholar specializing in biblical Hebrew might conceptualize a project and, via plain language prompts, meet the demands of each step (or the majority of steps) utilizing GPT-4. Beyond the ease of use, the scholar stands to gain substantial time savings. For instance, the process delineated above was completed in less than an hour. Here, the versatility of the

model's "intelligence" is demonstrated through its interpretation of various requests and its engagement across multiple domains of expertise. The capacity to fulfill such a broad range of tasks underscores the potential of GPT-4 as a powerful tool in the domain of academic research.

Conclusion

מַה־שֶּֽׁהָיָה֙ ה֣וּא שֶׁיִּהְיֶ֔ה וּמַה־שֶּׁנַּֽעֲשָׂ֔ה ה֖וּא שֶׁיֵּעָשֶׂ֑ה וְאֵ֥ין כָּל־חָדָ֖שׁ תַּ֥חַת הַשָּֽׁמֶשׁ׃[16]

Qohelet could not have envisaged the accumulation of advancements over centuries that could potentially culminate in the realization of an Artificial General Intelligence (AGI).[17] Engaging in hyperbole is difficult when discussing the prospects of an artificial entity capable of problem-solving across domains traditionally believed to be exclusive to human intellect. This development undoubtedly introduces something "new."

This cursory examination of the technology barely grazes the potential capabilities of this tool—particularly within the niche field of biblical Hebrew studies. What it has achieved, however, is the provision of an introduction to a toolkit that will inevitably pervade all domains in the future. The coherence of the three categories of experiments—knowledge retrieval or language understanding, generative modeling or creative problem solving, and command interpretation or query parsing—highlights the potential of generative AI to contribute significantly to the field of biblical Hebrew studies. Each category demonstrates a unique aspect of GPT's capabilities in understanding, interpreting, and generating content related to biblical Hebrew while aligning these abilities with the diverse tasks typically undertaken by scholars in the field.

At this nascent stage, this technology already possesses the capacity to disrupt established scholarly norms. The results of these experiments indicate that GPT, particularly in its more advanced iterations, possesses the capacity to assist scholars in diverse areas of biblical study, from basic syntactic analysis to more abstract hermeneutic tasks. Moreover, the integration of the model with best-in-class tools such as Text-Fabric and the BHSA dataset, once restricted to those with specific expertise in the intersection of digital humanities skills (e.g., ancient language and computer science), are now accessible to more traditional scholars of biblical Hebrew who lack proficiency in even rudimentary computer coding. The capacity to query these tools using plain language, by way of technologies such as GPT, brings their immense feature sets into the purview of those who would most benefit. As the printing press ushered in a new era of democratized knowledge and scholarship, it is plausible to envision a future where generative AI could achieve similar feats. As the technology continues to evolve, it is reasonable to expect even more impressive outcomes and applications in the realm of biblical Hebrew studies.

Yet, even amidst this formidable landscape of the distinctly "new," the wisdom of Qohelet rings true. While methods may evolve—and evolve dramatically—the human impulse towards curiosity and the

---

[16] Eccl 1:9.

[17] Relevantly, researches have observed "sparks" of AGI in the GPT-4 model; cf. Sébastien Bubeck et al., "Sparks of Artificial General Intelligence: Early Experiments with GPT-4," arXiv (2023): http://arxiv.org/abs/2303.12712.

quest for knowledge remains unchanged. Therefore, Qohelet's wisdom serves as a reminder that despite changes in tools and methods, the pursuit of knowledge and understanding continues as an intrinsic aspect of the human endeavor. In that sense, there is indeed nothing new under the sun. This interplay between ancient wisdom and contemporary technology emphasizes the continuity of human intellectual pursuit, from the era of Qohelet to the era of AI.

References

Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, et al. 2023. "Sparks of Artificial General Intelligence: Early Experiments with GPT-4." Accessed May 28, 2023. http://arxiv.org/abs/2303.12712.

ETCBC. n.d. "BHSA." Accessed May 28, 2023. https://etcbc.github.io/bhsa/.

Future of Life Institute. 2023. "Pause Giant AI Experiments: An Open Letter." Accessed May 28, 2023. https://futureoflife.org/open-letter/pause-giant-ai-experiments/.

"Generative AI Could Raise Global GDP by 7%." 2023. Accessed May 28, 2023. https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html.

OpenAI. n.d. "API Reference: Create Completion." *OpenAI API Reference*. Accessed May 28, 2023. https://platform.openai.com/docs/api-reference/models.

OpenAI. n.d. "GPT-4 Is OpenAI's Most Advanced System, Producing Safer and More Useful Responses." Accessed May 28, 2023.

Parry, Donald W., and Emanuel Tov, eds. 2004. *The Dead Sea Scrolls Reader. Vol. Texts Concerned with Religious Law*. Leiden ; Boston: Brill.

Roorda, Dirk. n.d. "Text-Fabric." Accessed May 28, 2023. https://annotation.github.io/text-fabric/tf/.