

# Business Applications For GPT-3

By Martin Anderson

First published\* **January 27th, 2021** at:

<https://www.width.ai/post/business-applications-for-gpt-3>

[Web-archived version](#)

The interpretive powers of the GPT-3 autoregressive language model stirred up popular tech headlines when it was unveiled by OpenAI in 2020. It could apparently compose feature articles, write poetry, and even talk to the dead.

GPT-3 has perhaps been a victim of its own publicity; it's a vanguard product in the NLP space, but widespread public interest in the possibility of an effective AGI (which GPT-3 is not), combined with tech media's determination to coopt any new AI product into its annual round of febrile headlines, has left GPT-3 rather misunderstood.

## Commercializing GPT-3

Partly to offset the \$4.6 million USD that OpenAI paid to train the 175 billion weights of GPT-3, OpenAI began to lease API access to the neural network. Python bindings are provided, while officially-listed third-party libraries include C#/.NET, Crystal, Dart, Go, Java, JavaScript/Node, Ruby and Unity.

Subsequently, its core capabilities have inspired a slew of startups across a range of sectors.

Here we'll take a look at the usable scope of GPT-3 for business purposes, and at some of the companies that have taken up the vanguard in this respect.

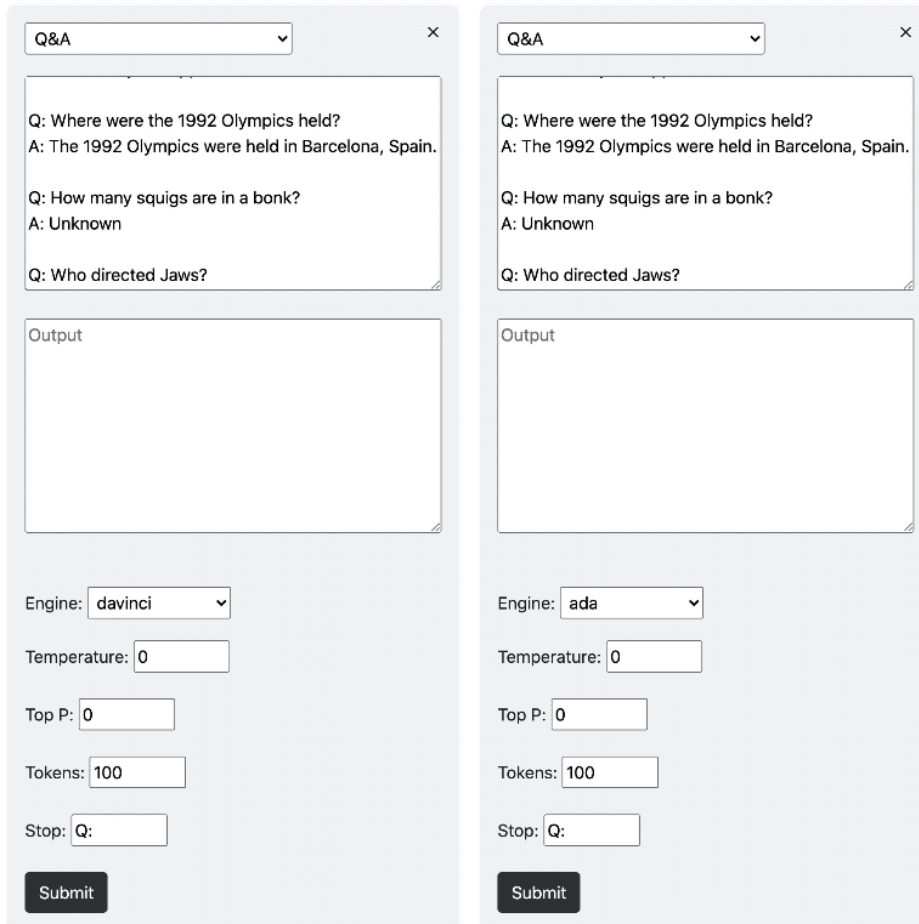
First, let's take a look at the strengths and weaknesses of the various methods by which GPT-3 answers a prompt.

## Choosing a Model in GPT-3

Models available for transformations in GPT-3 include Davinci, Curie, Babbage and Ada, each of which have different capabilities in terms of speed, quality of output and suitability for specific tasks.

For instance, Davinci is the most sophisticated of the available models, and is most likely to produce usable output that's more complex, and that explores (or at least appears to explore) higher-level domain thinking and analysis (later we'll also take a look at Davinci Instruct, which is capable of following more specific commands regarding the formatting and domain-specificity of its output).

However, some of the leaner and less computationally demanding models are more than adequate for simpler prompts, saving latency and API request costs. For instance, novelist Andrew Mayne has found that much of the most wide-spread knowledge available to GPT-3 is accessible across lower-level models than Davinci:



Here we see the full-fledged Davinci engine and the lighter Ada engine arriving at the same result.

## Playing To The Strengths Of GPT-3

Treated as a straightforward 'global oracle', GPT-3 is subject to the same inaccuracies and inexactitudes as the publicly available content that it was trained on, and the depth and truth of its responses on any subject is in proportion to the subject's representation (and misrepresentation) on the internet, and in the standard datasets that informs it.

For instance, Davinci knows a little about Charles Dickens' output, but will give up if you second-guess the first response:

## Davinci

**Human:** Who wrote the novel Bleak House?  
**AI:** Charles Dickens.

**Who wrote the novel Bleak House?**  
**AI:** I don't remember.

Answering the same question 'Who wrote the novel Bleak House?', the other models seemed less well-read, even on multiple attempts:

**Playground** ⓘ

Chat x ▾

⚠ 📄 🗑️ ⬆️ ⏪ ⚙️

**Who wrote the novel Bleak House?**  
**AI:** J. G. Lockhart wrote it, I suppose.  
**Human:**

Curie

**Human:** Who wrote the novel Bleak House?  
**AI:** Edgar Rice  
**Human:**

Babbage

Samuel Richardson. His best-known novel, popular over the past three centuries, is the famous novel "Pamela, or Virtue Rewarded."

What is the name of the novel that the cartoon character, Tom, a boy whose father and uncle are lawyers, is reading

*Various attempts by two 'lesser' GPT-3 engines to identify the author of an enduringly popular and classic novel.*

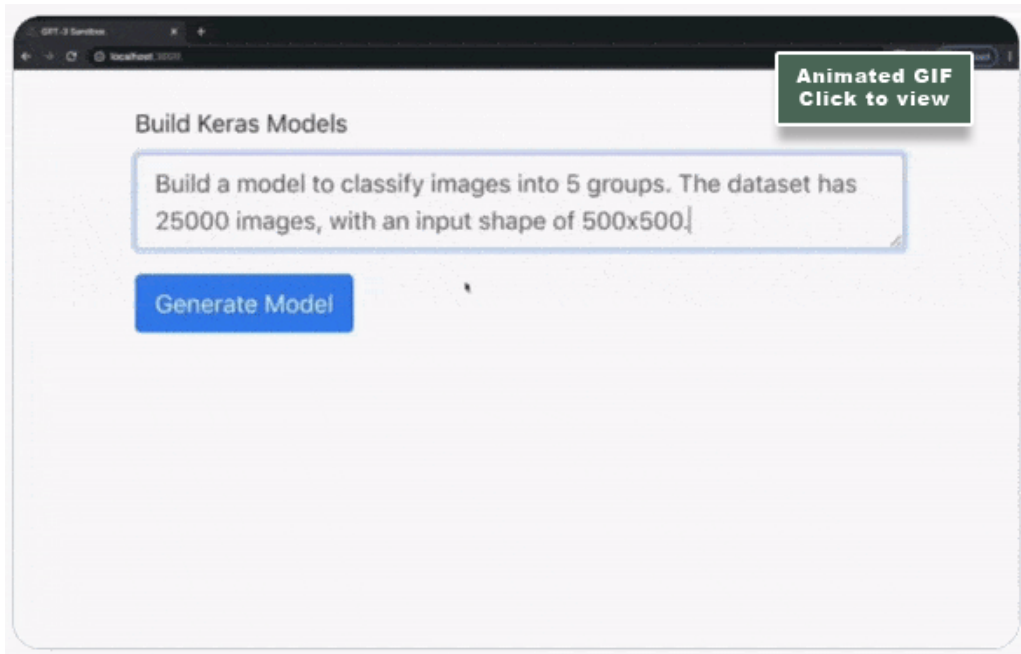
But within more narrow domains, usually away from subjective arts and culture topics, GPT-3 proves extraordinarily adept.

## Application Development With GPT-3

The GPT-3 model was developed through unsupervised training, wherein it had to discover, rationalize and categorize vast volumes of data without being explicitly told what any of it 'meant'. In the course of training it was necessary for the model to distil labels from the available entities – such as picture captions and descriptions – that accompanied the content, or that had some kind of discoverable relationship with it.

Where a subject has a rigid and agreed taxonomy and hierarchy (such as mathematics), it's much easier to be certain of the relationships between sub-entities in that domain, and to create generative rules that are accurate.

Programming falls into this category, and GPT-3 is very good at it.



*GPT-3 apparently producing an accurate and deployable structure for a Keras model solely from a description. It should be considered that while the code is correct, it may not be the optimal solution.*

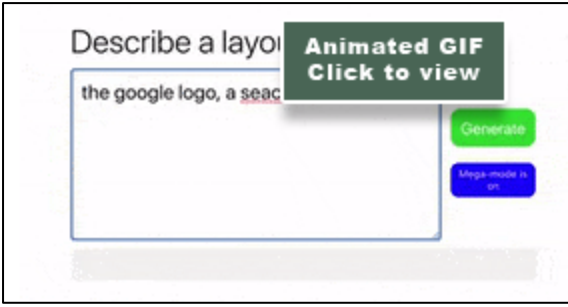
Considering that even professional coders may have eccentric solutions to programming challenges, and that the development community frequently comes into conflict regarding the correct answer to a problem, it's not clear why GPT-3 errs so little when writing code — unless, perhaps, that it might be the only coder who took the time to comprehensively RTFM.

It's also possible that in the lonely course of its training, the model has assigned higher weight to inbound links to 'best answer'-rated posts at sites such as StackOverflow, in much the same way that a search engine's algorithm will calculate authority from frequent citations.

However, GPT-3 is often not the only contributor to those cases where it has won headlines for its programming prowess.

## Combining GPT-3 And Domain-Specific Datasets

In the case of Sharif Shameen's Debuild, which leverages the OpenAI model to provide instantaneous app development from text input, the user-facing framework incorporates GTP-3, but was also specifically and intensively trained against content from GitHub and StackOverflow.



Even with spelling errors, the GPT-3-based Debuild programmatically reconstructs the classic Google interface through NLP.

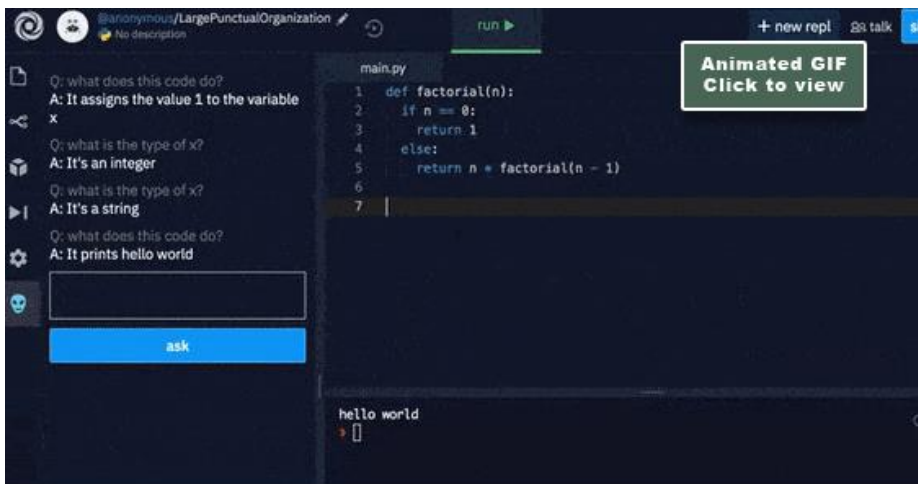
Here, GPT-3 handles the interpretive and transformational aspect of NLP requests, but will avoid wide-ranging answers in favor of focused, domain-limited responses that are (apparently) guided by the weights of a separately trained model.

## Explaining Code With GPT-3

There is a well-established market for any product or process that can explain legacy code – live software that may now have multiple dependencies across version-locked VMs, and which in some cases may have roots as far back as the 1970s.

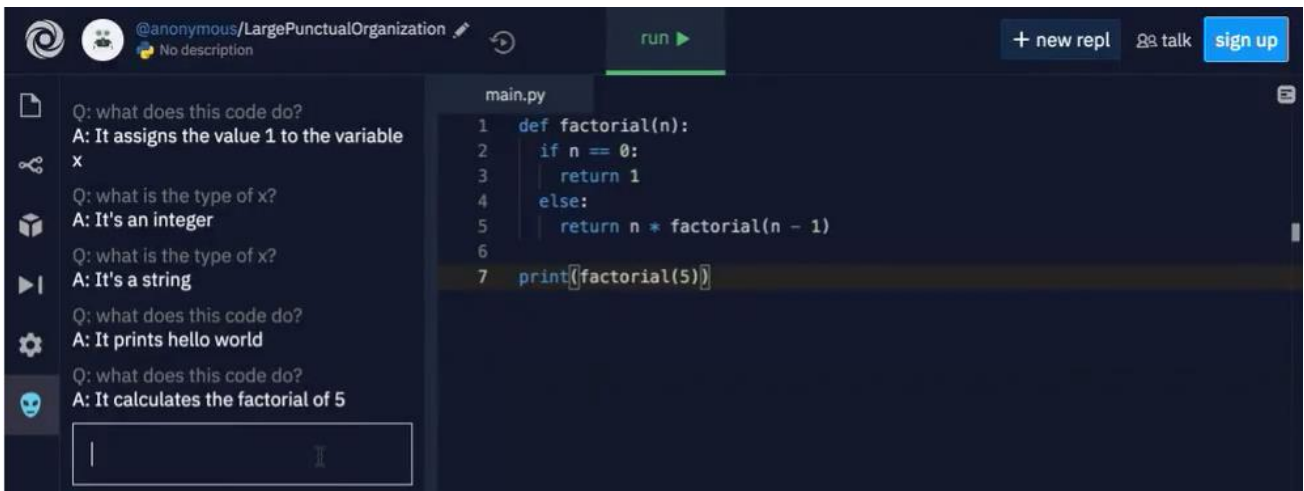
Technical debt of this nature can force companies to retain otherwise unproductive staffers who have become the gatekeepers of arcane but essential systems, or necessitate massive consultation bills if the originating developers should die or depart the company without adequately documenting their sprawling projects.

The Replit Code Oracle uses GPT-3 to explain what complex code will actually do.



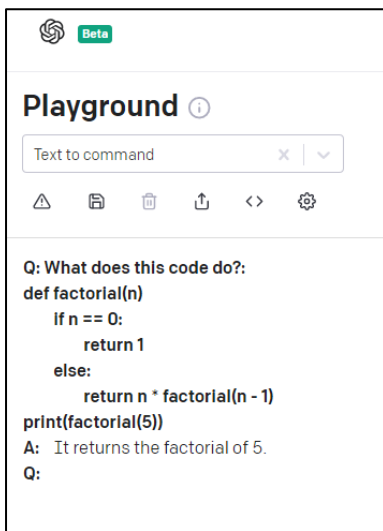
GPT-3 explains what complex code does in plain English.

The company, which is based in San Francisco and raised \$20 million USD in series A funding in February 2021, has even recruited GPT-3 to explain Replit's intent to 'rewrite all the rules of software development' by making natural language the default interface to code generation.



Amjad Masad, Twitter, 22nd July 2020, <https://twitter.com/amasad/status/1285789362647478272>

Since Replit has not released details of any co-training that contributes to its code-explanation algorithm, I tested the Replit challenge in the above image in the basic OpenAI GPT-3 playground, and found that the Davinci and Babbage engines can calculate a correct answer by default, without additional domain-specific training:

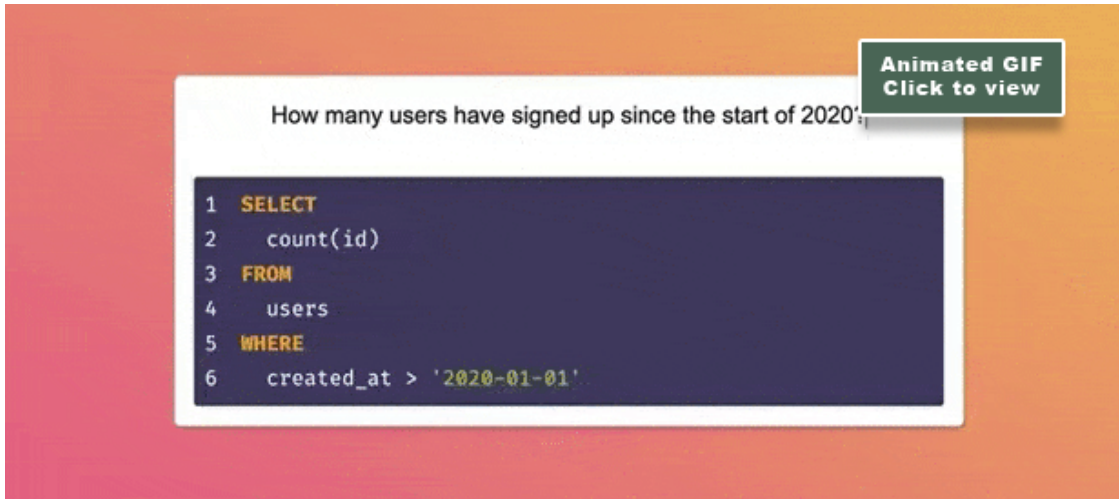


In developing a complex code interpreter with GPT-3, it is important to remember that where a problem has not been directly assimilated from dataset material at training time, GPT-3 would likely fall back to a text analysis and semantic deconstruction of the written challenge, in the absence of examples in its database that match a particular text pattern.

However, since GPT-3 is very capable of deconstructing query and request logic based on similar taxonomies in a programming language, its 'logical guesses' are well-informed, and likely to be correct.

## SQL Queries From Natural Language

Research into text-to-SQL predates the advent of GPT-3; but since standard SQL query language is in itself designed to conform to natural language, GPT-3 is very good at reproducing functional SQL requests:



'GPT-3 writes my SQL queries for me', Vecanoi, YouTube, 28th July 2020 - <https://twitter.com/rajnishkumar/status/1288502875455475712>

More complex SQL queries are far less human-readable, and nested conditional phrases make mistakes more likely when humans are writing out the commands. Due to the inflexibility of the SQL language, and its restricted taxonomy, GPT-3 has assimilated SQL brilliantly.

New York-based analytics startup SeekWell has adopted GPT-3 as an SQL 'translation' layer.

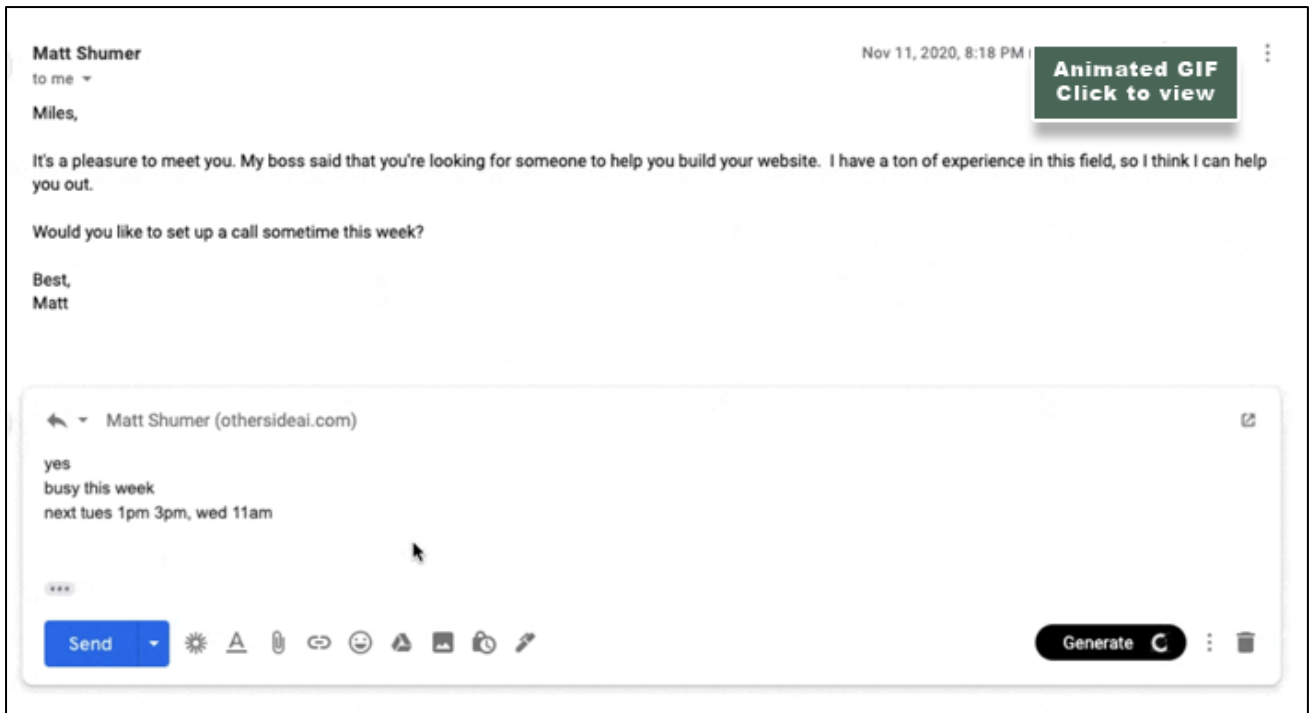


'Automating my job by using GPT-3 to generate database-ready SQL to answer business questions', Brian Kane, SeekWell Blog, 2020 - <https://blog.seekwell.io/gpt3>

The company has noted that the Davinci instruct model performs best at SQL queries, since it allows the user to provide specific guidance, such as 'Only respond in correct SQL syntax'.

## Automating Email Generation With GPT-3

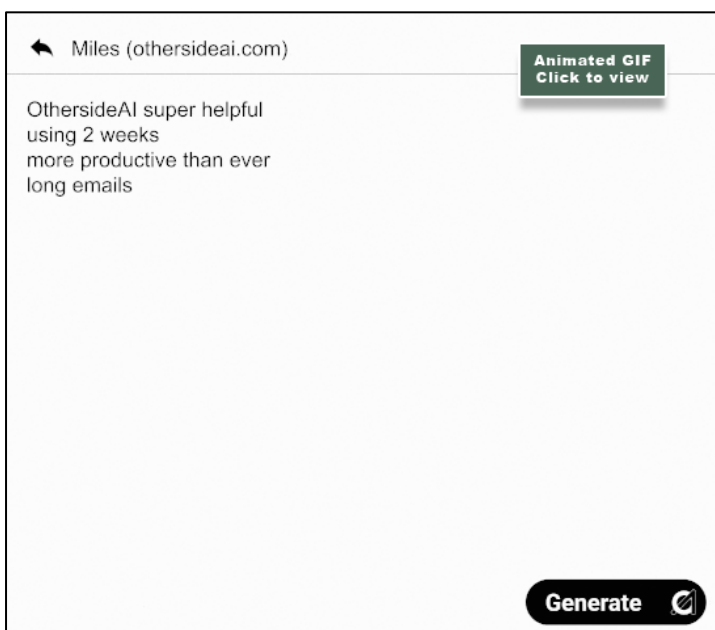
Currently in closed alpha, OthersideAI's email generation system uses GPT-3 to generate email responses based on bullet-points that the user provides.



Entitled 'Quick Response', the synthetic replies take into account the content and context of other emails in the chain, but seem to transform the responses primarily based on GPT-3's understanding of business email style.

The AI can infer the temporal tense from brief prompts, for instance turning 'good to meet' into 'It was good to meet you'.

However, it's not clear whether Quick Response could draw on earlier emails that contain meeting details to develop a response such as 'It was good to meet you for coffee Tuesday at Starbucks on 1st St.', or whether it understands how mechanistic and data-stuffed a response like that would seem, and would favor a conversational style instead.

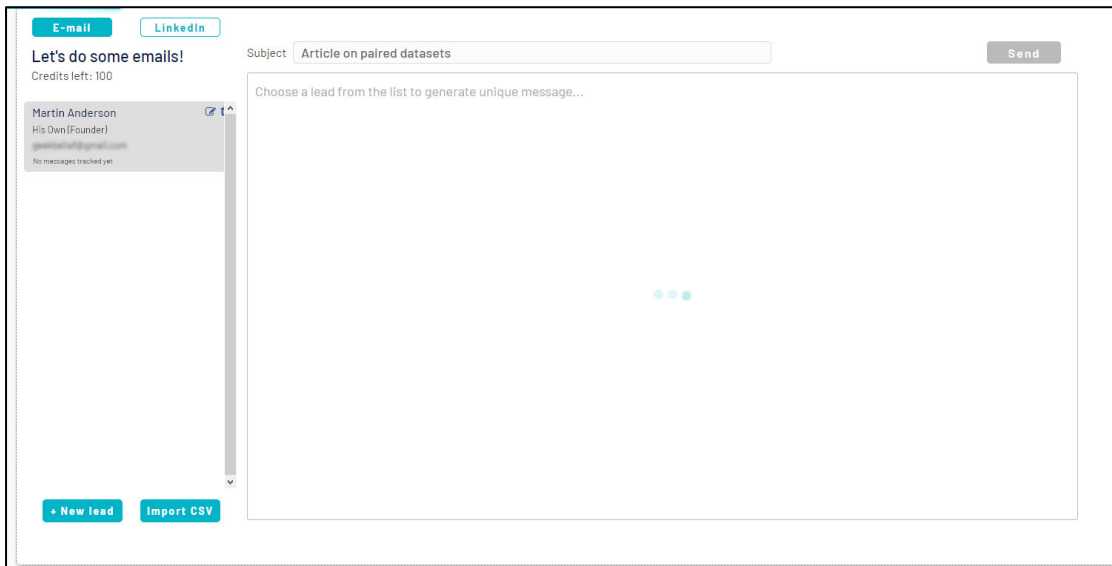




Quick Response is posited as the first stage in a multi-AI communications framework wherein, effectively, two people may exchange information and ideas programmatically and without generating the text content themselves. It's analogous, arguably, to two business entities communicating through their lawyers instead of directly.

At the time of writing, the wait list for the OthersideAI email generation API stands at 15,000+.

Other email-based GPT-3 tools have arisen: Lean To provides a dedicated user interface for AI-based email generation.



<https://leanto.me/>

When setting up a new project (i.e. an email to single or multiple participants), you can choose the objective of the email, even if it is just 'getting a response', and can import contacts via a CSV file, or add them manually.

As is often the case, the GPT-3-based interface requires a lot of prior information. In this case I chose to write to a Stanford Vision Lab professor, and had to do a fair bit of research in order to fill in all the necessary fields.

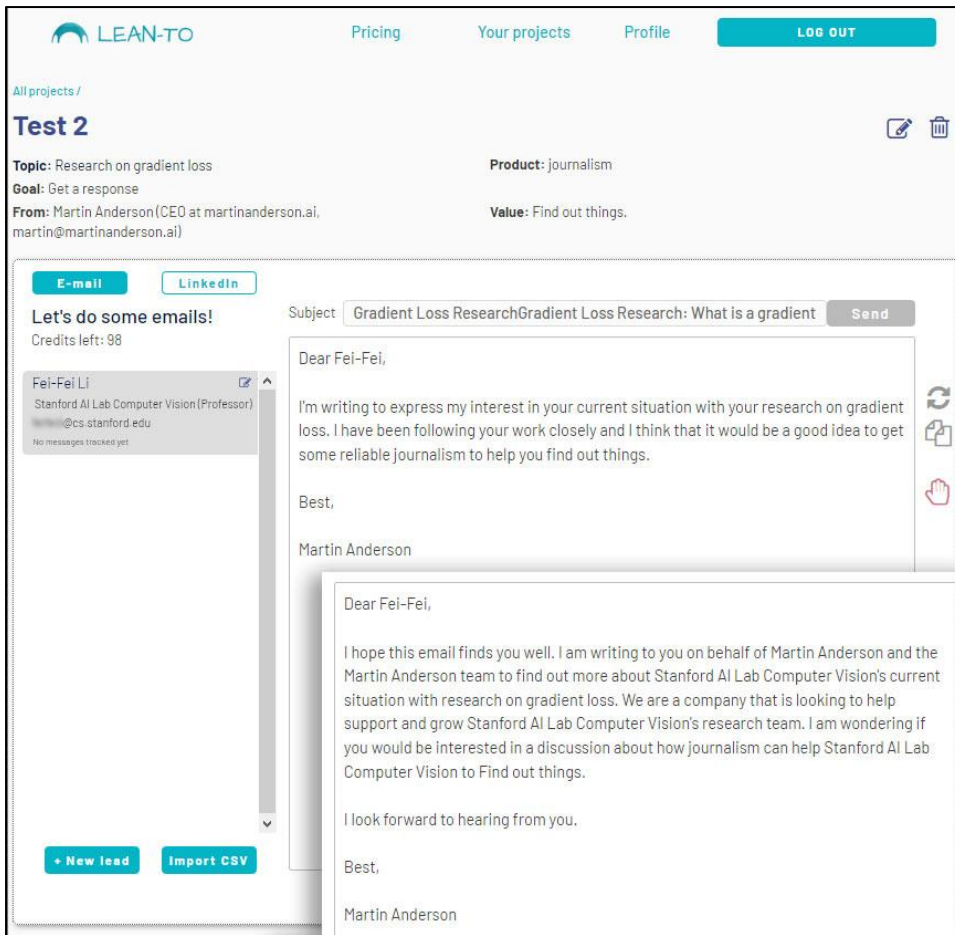
**New lead**

The more info you add, the better will be the email content. Name, email, Position and Company are the absolute minimum.

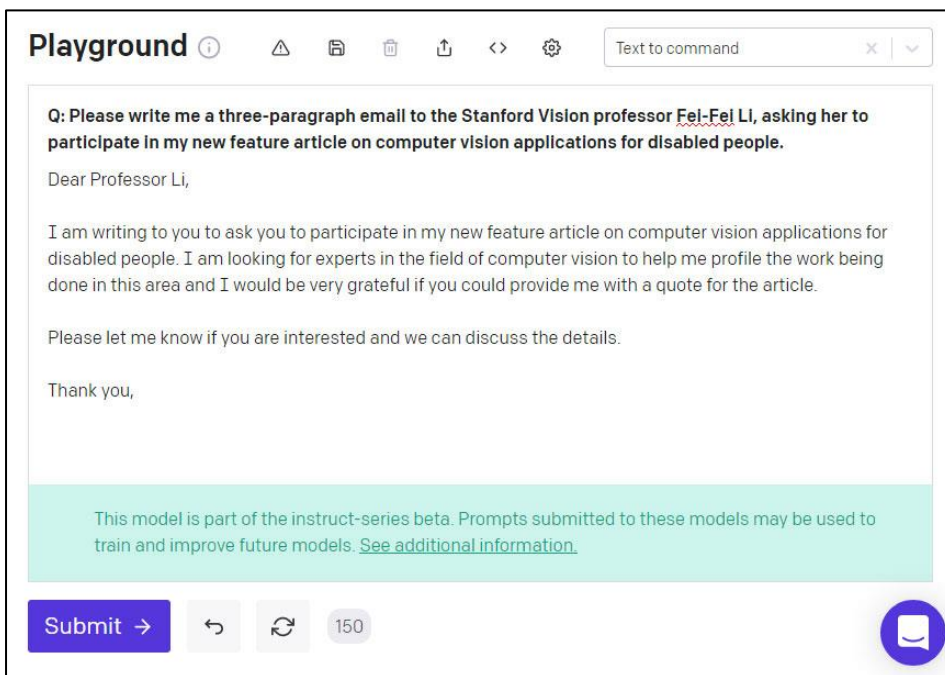
Name	Fei-Fei	Li
Email	[redacted]@cs.stanford.edu	
Position	Professor	
Company	Stanford AI Lab Computer Vision	
Description	. Fei-Fei Li is the inaugural Sequoia Professor in the Computer Science Department at Stanford University	
Industry	Computer Vision Research	
Website	http://vision.stanford.edu/	
City	353 Serra Mall, Gates Building, Stanford,	
Country	United States	

**ADD**

The results are unconvincing, and generic enough to likely trigger spam filters, while subsequent attempts do not improve the quality. Despite having put research leads into the contact details (such as a website for the recipient), there's no evidence that any of this information was used.



Bar the odd missing comma, the basic GPT-3 playground does a far better job via the Davinci instruct engine:



OpenAI's base API writes a more credible letter to the professor with sensible prompting that did not take long to write, and which did not require forms and additional procedures.

The lack of client integration in Lean To is also a hindrance; though the email can be sent through Lean To's own SendGrid system, this would make replies problematic, and increase the chances of the message being automatically flagged as spam. Therefore it's better to just copy and paste the generated output into your own email environment.

## Breaking Down a GPT-3 Prompt

Lean To's wizard-driven interface attempts to deconstruct the semantically important parts of a GPT-3 prompt and reassemble it into the 'ideal' prompt (which is not shown to the user), rather than writing it naturally; but 'raw' GPT-3 does a better job of this deconstruction by itself, so long as you choose the right engine and phrase the prompt succinctly.

Ultimately, the value of a GPT-3-derived email generation system is in being able to use it directly in an email client, since, as we've seen, it is easy enough to appropriately prompt GPT-3 in any testbed and ask it to generate an email.

## GPT-3 Decrypts Legal Jargon

Once GPT-3 has inferred the semantic intent of a text, it can rephrase that content in a number of ways, including ELI5. This is very useful for locating the devil in the details of verbose EULAs and general legal documentation:

AI Translations (AI Output in Blue)
<p><b>Legal Clause:</b> Upon liquidation of the Company, the Series A Shareholders will receive in preference to all other shareholders an amount in respect of each Series A Share equal to one <u>times</u> the Original Issue Price (the "Liquidation Preference"), plus all accrued but unpaid dividends. To the extent that the Company has assets remaining after the distribution of that amount, the Series A Shareholders will participate with the holders of Ordinary Shares pro rata to the number of shares held on an as converted basis.</p>
<p><b>Plain English:</b> If the startup is wound up, the Series A investors will be paid back at least what they invested and they will also share any leftover assets with ordinary shareholders.</p>
<p><b>Legal Clause:</b> Sale of all or substantially all of the assets of the Company or a sale of shares involving a change in control (each, a "Corporate Transaction") will be treated in the same way as a liquidation and the proceeds of sale will be distributed as set out in paragraph 3. If the holders of Series A Shares have received any Special Dividend it shall be set off against their Liquidation Preference.</p>
<p><b>Plain English:</b> If the company is sold, or a new owner takes control, the proceeds of the sale will be distributed as in the liquidation clause above. Any special dividend paid will be treated as an initial payment towards the Series A investors.</p>
<p><b>Legal Clause:</b> If holders of at least 51% of the Series A Shares and Ordinary Shares agree to sell their shares, there will be drag along rights so that all remaining shareholders and option holders will be required to sell on the same terms, provided that the dragged shareholders will not be required to provide to the purchaser any representations or warranties except as to title or to agree to any other terms.</p>
<p><b>Plain English:</b> If a majority of investors agree to sell their shares, any other investors who don't agree to sell will be forced to do so. They won't have to provide any new assurances or guarantees but they will have to accept the same price and terms as the others.</p>

GPT-3 explains the intent of legal jargon', @michaeltefula, Twitter, 21st July 2020 - <https://twitter.com/michaeltefula/status/1285505897108832257?lang=en>

GPT-3 has a mode called 'summarize for a 2nd grader' that makes a broadly fair assessment, for instance, of the dense EULA for the videogame *Far Cry*:

**My second grader asked me what this passage means:**

**"UBISOFT (or its licensors) grants You a non-exclusive, non-transferable, non-sublicensed, non-commercial and personal license to install and/or use the Product (in whole or in part) and any Product (the "License"), for such time until either You or UBISOFT terminates this EULA. You must in no event use, nor allow others to use, the Product or this License for commercial purposes without obtaining a license to do so from UBISOFT. Updates, upgrades, patches and modifications may be necessary in order to be able to continue to use the Product on certain hardware. THIS PRODUCT IS LICENSED TO YOU, NOT SOLD. "**

**Please rephrase it for him, in plain language that a second grader can understand:**

"If you buy this game, you can play it on your computer. But you can't sell the game or give it to somebody else. If we make a better version of the game, you'll have to buy that too. You can't play the game on any other computer. And you can't make copies of the game for your friends."

I'm sure that is not what they intended, but that is what it says.

## Navigating Legal Ambiguities and Unclear Terms

Despite the context of 'high-precision language', legal terminology often contains ambiguous terms that seem designed to produce a 'chilling effect' without specifying what a term signifies, or what a breach of it might entail. Since such clauses are designed to be tested in court, neither a person nor an AI can hope to extract their exact meaning, because, in the context of the document, that has yet to be established.

In the above case, GPT-3 has sided with the cynics, and assumes that paid upgrades will be necessary in order to continue to use later versions of the videogame (rather than any sequels to it). However, GPT-3 has not correctly understood that this probably applies to cases where the existing game is subsequently ported to next-gen platforms, and must be paid for all over again, because the original terminology ('certain hardware') is very vague.

The GPT-3 summary also states that the user can't play the game on any other computer, which is usually not true. Here GPT-3 seems to have mistaken the intent of 'non-transferable' (which means that the user cannot re-sell the license to someone else, and not that the user cannot de-authorize the game on one device and re-authorize it afterwards on another).

In fairness, a lot of actual people would struggle to define the meaning of this term, in this context.

It seems likely that a GPT-3-based framework for decrypting legal jargon would benefit from cross-training against a dedicated dataset of legal terms, where manual labeling could eliminate some of these ambiguities, or at least give GPT-3 enough confidence to declare that certain legal terms may not be clear in their intent, and to outline some possible meanings for them, based on historical examples.

## Marketing Copy - Social Media to Blog Articles

Though the AI-generated content sector was already populated with platforms such as Articoolo, Article Forge and Phrasee, GPT-3's text-generation and semantic capabilities inspired a slew of new AI content startups around the time of its launch, and since.

Some of the more notable names include CopySmith, WriteSonic, Headlime and PepperType. Though we don't have space here to trial all of the new GPT-3 content platforms, we'll shortly look at a couple of the most popular.

## GPT-3 as a Social Media Generator

In a world where an inappropriate or inconsiderate tweet can end a career or tank a company's stock, there's no realistic prospect that short-form social media output from GPT-3 will ever become entirely programmatic.

However, this more succinct kind of auto-generated social media output is what the current crop of companies are really pushing, perhaps because GPT-3 can produce very apposite results in such a short window of attention.

In the long term, the value of GPT-3 copy generation frameworks may lie partly in the (nascent) potential to write more substantial works, such as coherent blog posts (or at least, in providing better article-spinning transformations than are currently available); in producing corrected or alternative versions of shorter, human-generated social media snippets; and in generating correct English language short-form output for users that have a different native language.

## CopySmith

Founded in October 2020, CopySmith leverages GPT-3 to provide a suite of templates to massage existing copy and generate new text-based material from limited prompts.

CopySmith offers text-completion, rewriting and generation services for a range of contexts, including:

- Product descriptions
- Landing pages
- Blog ideas
- Content ideas
- Content enhancement, TLDRs and summaries
- Headlines for ads, with templates for LinkedIn, Instagram, Facebook and other major ad platforms.
- List generation from a suggested topic
- Metatag generation
- Marketing email generation
- Client pitches

...and many more besides. However, in practice nearly all of these possible use cases are minor variations of four basic NLP capabilities in GPT-3:

- Summarizing (presenting the same information using fewer words).
- Rewriting (presenting the same information using a similar number of different words).
- Extemporizing (using limited text cues to generate a longer and more discursive text).
- Semantic distillation of entities from text (i.e. deriving 'female interest' from 'thousands of women use our products every day', etc.).

## Blog Creation at CopySmith

CopySmith also has a blog post generator, which is currently in beta – but you'll need to 'get out and push' quite a lot:

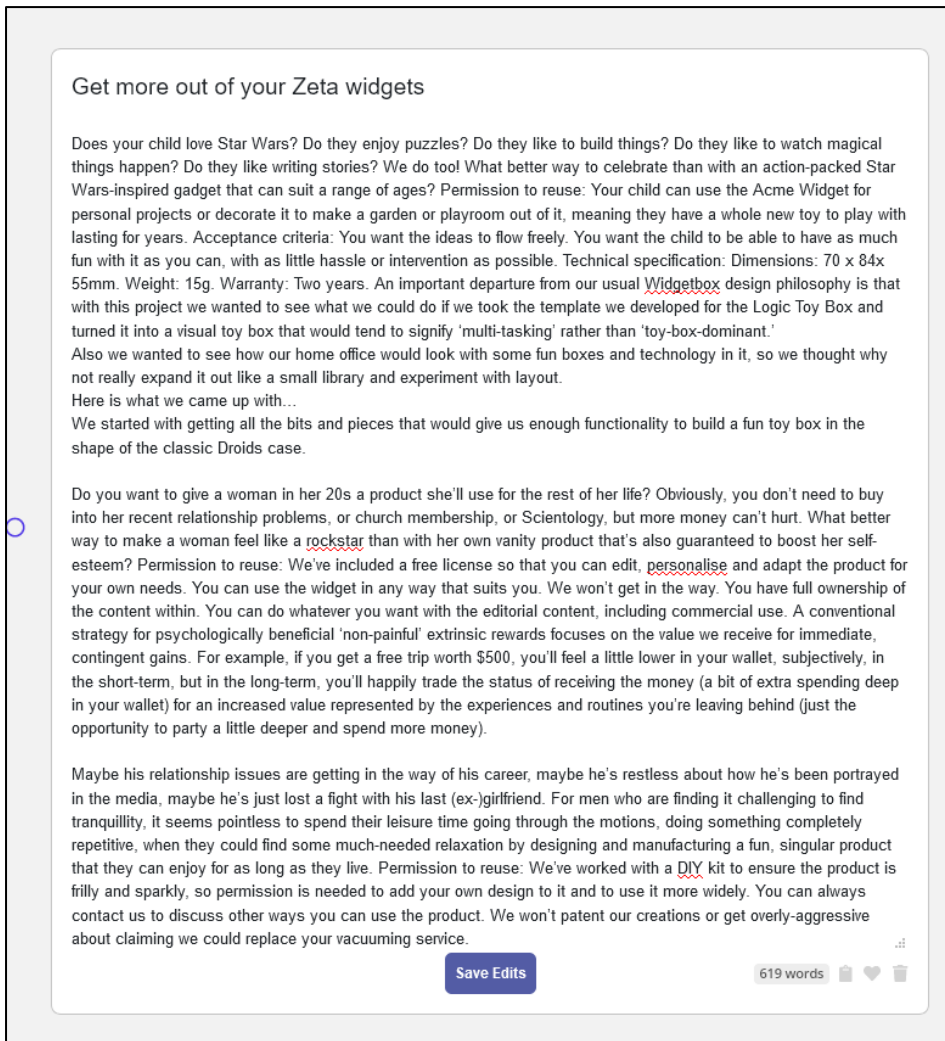
The screenshot shows the CopySmith blog post generator interface. It includes the following fields and elements:

- Company Name:** A text input field containing "Zeta Widgets".
- Audience:** A dropdown menu showing "30-40 middle income".
- Blog Post Title:** A text input field containing "Get more out of your Zeta widgets".
- Add an example paragraph - this will guide the style and main topic of the blog post:** A text area containing "We're so glad to have you as a Zeta Widgets customer, and we hope that you're enjoying your zeta products". A red tooltip indicates "20 words (minimum 120 words)".
- First Sentence of Paragraph One:** A text input field containing "Our Acme Widget is a fun product aimed at the 5-11 age range."
- First Sentence of Paragraph Two:** A text input field containing "Our Default Widget is for women aged 20-35, and is one of our most popular lines."
- First Sentence of Paragraph Three:** A text input field containing "Our Hero Widget is aimed at men in the 25-40 age range."
- Keywords \*:** A text input field containing "Enter ↵". Below it, a tag cloud shows "widgets, fun, value, family".
- One Blog Post:** A dropdown menu with "One Blog Post" selected.
- Generate:** A blue button.
- Credits:** A message stating "You have 5 blog post credits. This generation will use 1 credit."
- Disclaimer:** A message stating "All content is saved by default, so if the first set isn't perfect, try again! Previous generations can be found below the new ones." and "These blog posts are not GPT-3 generated — we are retraining on a larger dataset and working on quality improvements! The generator will take about 1.5 minutes to return one or two blog posts (650 words each)."

CopySmith's blog writer, in common with most similar offerings, needs a lot of help. - <https://copysmith.ai/>

As with CopySmith's rivals, it's not possible to prompt the blog writer with 'Write me an article about this widget company'; rather, the module focuses on user input, and needs a meaty minimum of 120 words to get started, as well as the initial sentence of the subsequent three paragraphs. However, it is possible to generate these requisites in other CopySmith modules, and then input them at this stage.

I filled in the initial paragraph without giving it any real information about the product, besides the brand being advertised and some general information about the target age-ranges for various products, and tested the blog-creation interface:



**Get more out of your Zeta widgets**

Does your child love Star Wars? Do they enjoy puzzles? Do they like to build things? Do they like to watch magical things happen? Do they like writing stories? We do too! What better way to celebrate than with an action-packed Star Wars-inspired gadget that can suit a range of ages? Permission to reuse: Your child can use the Acme Widget for personal projects or decorate it to make a garden or playroom out of it, meaning they have a whole new toy to play with lasting for years. Acceptance criteria: You want the ideas to flow freely. You want the child to be able to have as much fun with it as you can, with as little hassle or intervention as possible. Technical specification: Dimensions: 70 x 84x 55mm. Weight: 15g. Warranty: Two years. An important departure from our usual Widgetbox design philosophy is that with this project we wanted to see what we could do if we took the template we developed for the Logic Toy Box and turned it into a visual toy box that would tend to signify 'multi-tasking' rather than 'toy-box-dominant.'

Also we wanted to see how our home office would look with some fun boxes and technology in it, so we thought why not really expand it out like a small library and experiment with layout.

Here is what we came up with...

We started with getting all the bits and pieces that would give us enough functionality to build a fun toy box in the shape of the classic Droids case.

Do you want to give a woman in her 20s a product she'll use for the rest of her life? Obviously, you don't need to buy into her recent relationship problems, or church membership, or Scientology, but more money can't hurt. What better way to make a woman feel like a rockstar than with her own vanity product that's also guaranteed to boost her self-esteem? Permission to reuse: We've included a free license so that you can edit, personalise and adapt the product for your own needs. You can use the widget in any way that suits you. We won't get in the way. You have full ownership of the content within. You can do whatever you want with the editorial content, including commercial use. A conventional strategy for psychologically beneficial 'non-painful' extrinsic rewards focuses on the value we receive for immediate, contingent gains. For example, if you get a free trip worth \$500, you'll feel a little lower in your wallet, subjectively, in the short-term, but in the long-term, you'll happily trade the status of receiving the money (a bit of extra spending deep in your wallet) for an increased value represented by the experiences and routines you're leaving behind (just the opportunity to party a little deeper and spend more money).

Maybe his relationship issues are getting in the way of his career, maybe he's restless about how he's been portrayed in the media, maybe he's just lost a fight with his last (ex-)girlfriend. For men who are finding it challenging to find tranquillity, it seems pointless to spend their leisure time going through the motions, doing something completely repetitive, when they could find some much-needed relaxation by designing and manufacturing a fun, singular product that they can enjoy for as long as they live. Permission to reuse: We've worked with a DIY kit to ensure the product is frilly and sparkly, so permission is needed to add your own design to it and to use it more widely. You can always contact us to discuss other ways you can use the product. We won't patent our creations or get overly-aggressive about claiming we could replace your vacuuming service.

[Save Edits](#) 619 words

The texture of the post is consistently upbeat, as one would expect, but the module struggles for continuity of theme, and invents product features and purposes out of whole cloth. Even within the individual paragraphs, where one would expect reasonable continuity, the text struggles to stay within its own proposed themes, and the net effect is incoherent.

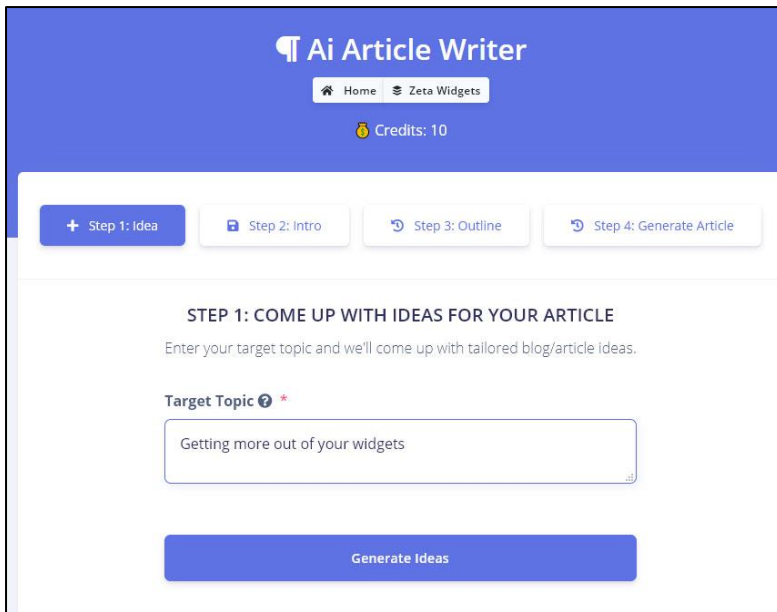
CopySmith is planning to develop SEO principles into its GPT-3 output, as well as paired training on industry-specific topics, which may help some of the 'focus' issues in longer sections of text output.

## WriteSonic

UK-based Writesonic offers a very similar range of services to CopySmith. The beta article writer, like CopySmith's, requires hand-written prefacing text, between 100-150 words.

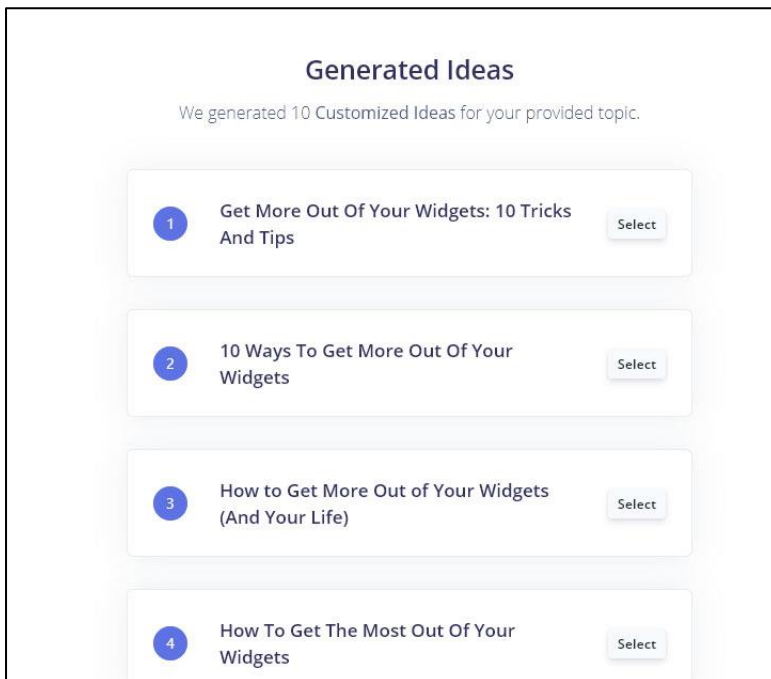
Initially, the system requires a target topic:





<https://writesonic.com/magic/project/>

Subsequently you'll receive ten article ideas:



After this, it's time – inevitably – to provide a detailed corpus of words from which the article will be manufactured:

### STEP 2: WRITE AN INTRO FOR YOUR ARTICLE

Now that you have an idea, write a detailed intro that describes what your article/blog will be covering and why it is important.  
Include users' concerns and likely solutions with facts. (We recommend a minimum of 100-150 words).


**Article Title** ⓘ \*

**Article Intro** (100-150 words) ⓘ \*

Widgets can make a huge difference both to your personal and professional life. New technologies have made widgets truly transformational, and our range of widgets - aimed at every sector, from children through to specific widgets for men and women, and for seniors - are designed to make your life easier. Let's take a look at some of the ways that this happens, and some of the products that integrate into our widgets.

Now, click the button below to proceed to Step 3.

[Next Step →](#)



Having chosen one, you can then choose from ten proposed article outlines, where every two sections cost one credit:

### Generated Outlines

We generated 5 Customized Outlines for your provided input.

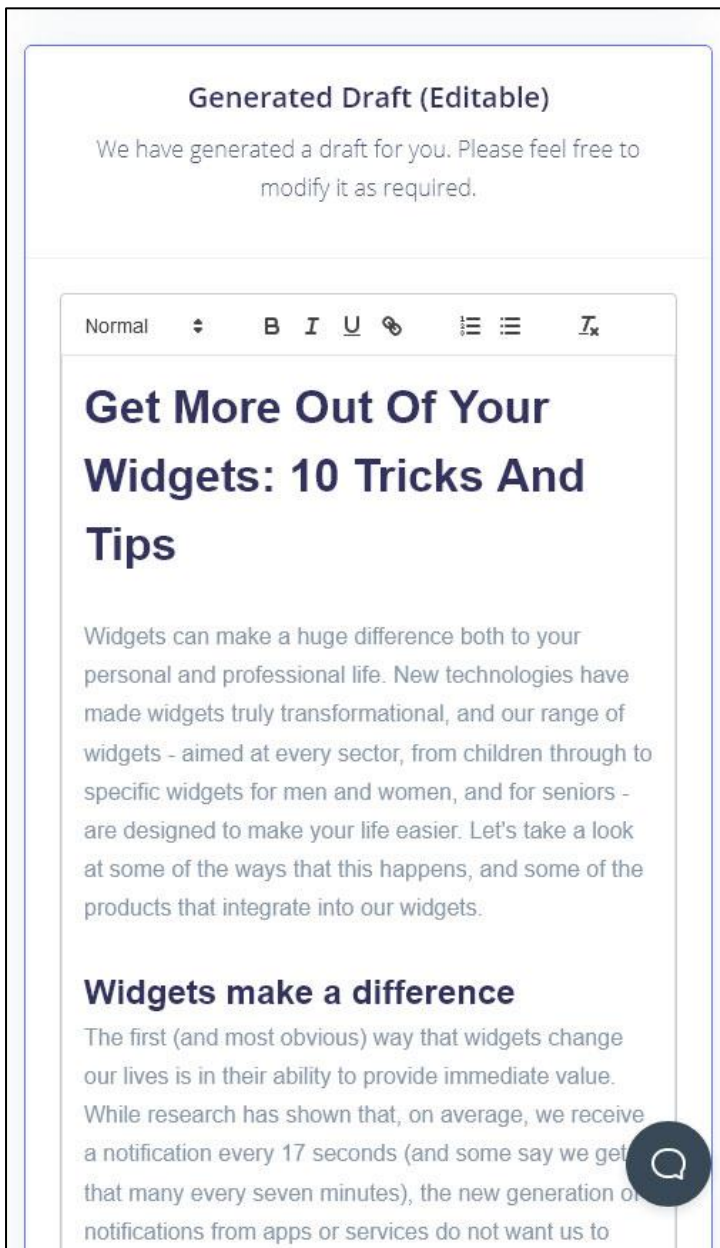
- Section 1: Widgets make a difference  
Section 2: Widgets for business  
Section 3: Widgets for children  
Section 4: Widgets for adults  
Section 5: Widgets for seniors  
Section 6: Conclusion

Select
- Section 1: The Ways Widgets Can Transform Your Life  
Section 2: The Widgets We Offer  
Section 3: 10 Tricks To Get The Most Out Of Widgets  
Section 4: Conclusion

Select
- Section 1: Widgets make life easier  
Section 2: Widgets for women  
Section 3: Widgets for men  
Section 4: Widgets for seniors  
Section 5: Widgets for a variety of sectors  
Section 6: Conclusion

Select

Finally, you'll get a completed article:



The screenshot shows a web interface for editing a draft. At the top, it says "Generated Draft (Editable)" and "We have generated a draft for you. Please feel free to modify it as required." Below this is a rich text editor with a toolbar containing options for bold, italic, underline, link, list, and text color. The draft content includes a main heading "Get More Out Of Your Widgets: 10 Tricks And Tips", a paragraph about the benefits of widgets, a sub-heading "Widgets make a difference", and a paragraph about the value of widgets. A dark circular button with a white speech bubble icon is visible in the bottom right corner of the editor area.

**Generated Draft (Editable)**

We have generated a draft for you. Please feel free to modify it as required.

Normal **B** *I* U 🔗 ☰ ☰ **Tx**

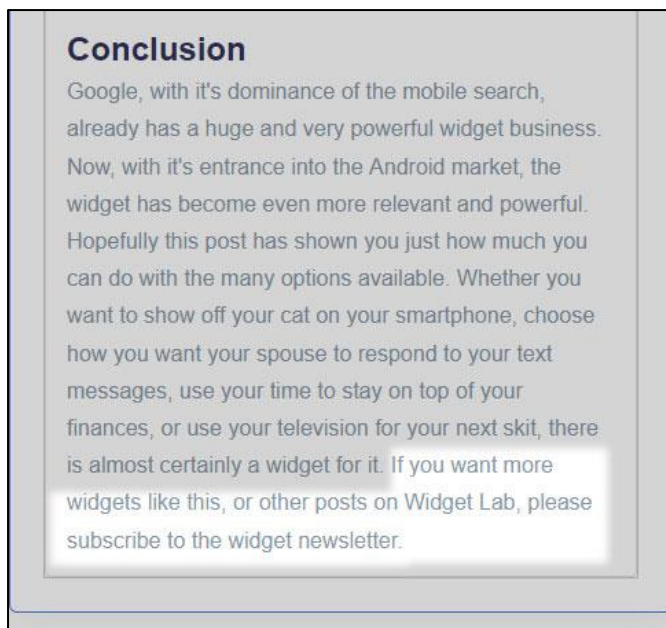
## Get More Out Of Your Widgets: 10 Tricks And Tips

Widgets can make a huge difference both to your personal and professional life. New technologies have made widgets truly transformational, and our range of widgets - aimed at every sector, from children through to specific widgets for men and women, and for seniors - are designed to make your life easier. Let's take a look at some of the ways that this happens, and some of the products that integrate into our widgets.

### Widgets make a difference

The first (and most obvious) way that widgets change our lives is in their ability to provide immediate value. While research has shown that, on average, we receive a notification every 17 seconds (and some say we get that many every seven minutes), the new generation of notifications from apps or services do not want us to

Given that the text prompt was incredibly generic, and designed to challenge the AI, Writesonic has produced a surprisingly coherent article from the draft outline, working productively through the header topics to an industry-standard closing paragraph, including a call-to-action.



WriteSonic gives you 10 free generation tokens, but just one attempt at an article is likely to burn through all of them, and so it is very difficult to give the product a second chance if it disappoints.

## Originality - And Excessive Invention?

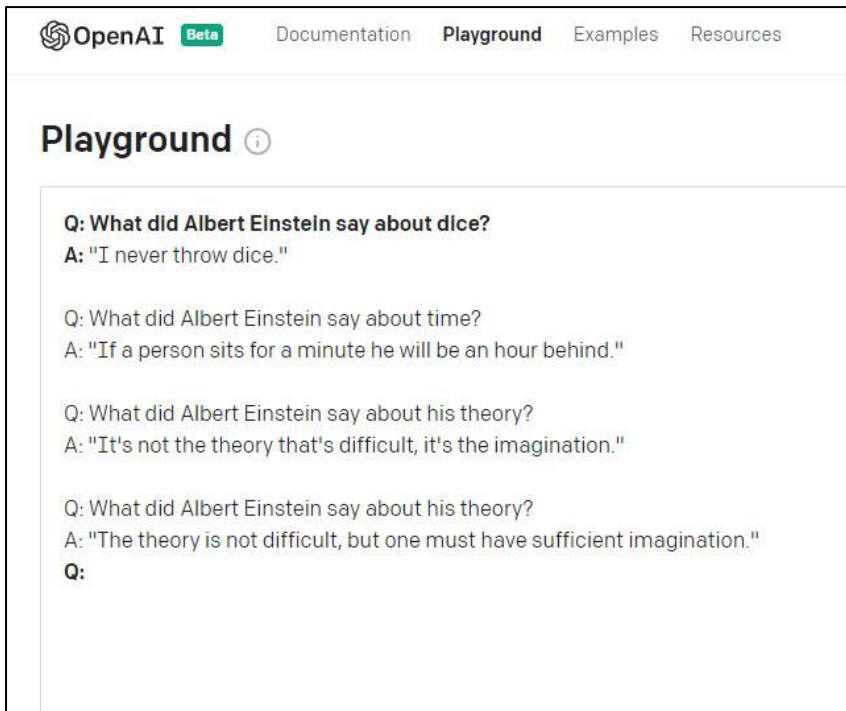
Though the GPT-3 article writing system AI Writer claims that the content-theft sentinel CopyScape identifies 94.47% of its articles as unique, CopyScape could not find *any* matches for the final article texts generated from my experiments with CopySmith and WriteSonic.

One word of (anecdotal) warning: when GPT-3 includes interview quotes in its article-style output (i.e., when I have asked for text, but not specifically for quotes or for interview material), I've found, in my two months of experimenting, that not even one of those quotes was available on the internet.

Since GPT-3 is trained on public-facing documents that are likely to be indexed in major search engines, this has led me to believe that although GPT-3 *can* find apposite interview material, it may currently be reluctant to quote a source verbatim.

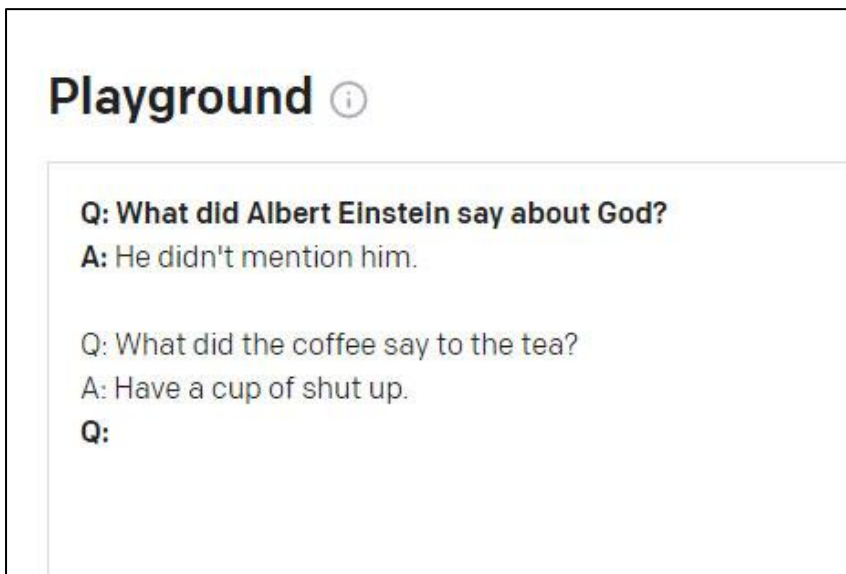
## Things Einstein Never Said

I tested this theory in the GPT-3 playground, asking Davinci Instruct 'What did Einstein say about dice?', expecting to receive the famous quote 'God does not play dice with the universe.'

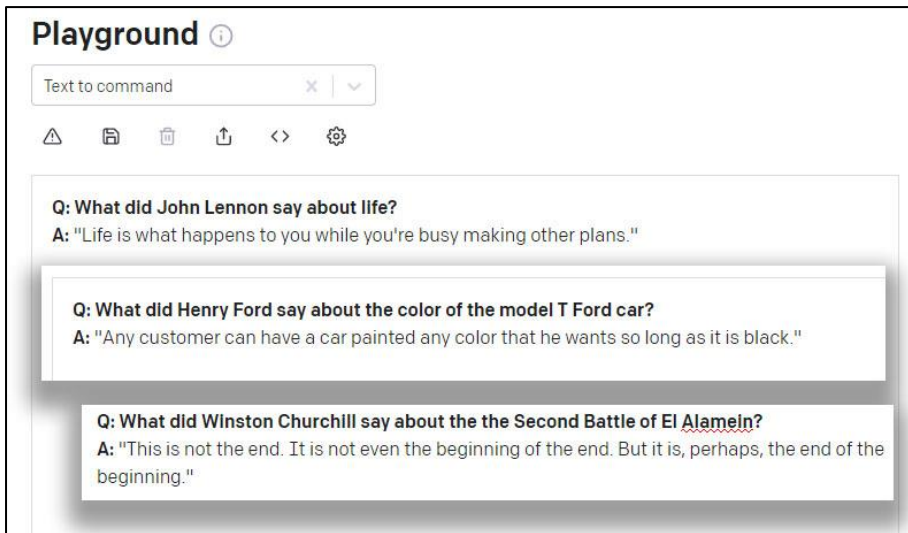


An exact-text search on Google indicates that either Einstein never said any of these things, or that if he did, the internet doesn't know about it. These are, at best, paraphrases — a core function of GPT-3.

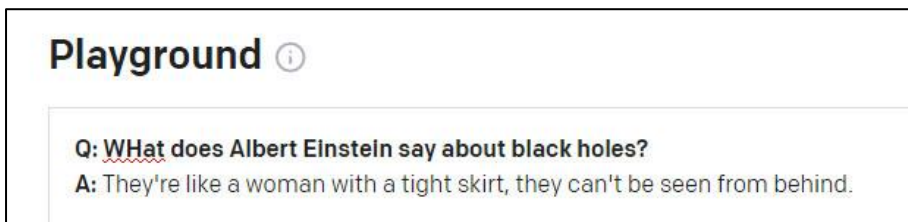
Further submissions were unhelpful:



This would be less of a hazard if GPT-3 was *consistently* wrong about quotes; however, it does get many famous sayings correct:



Watch out as well for persistent session variables in a GPT-3 consultation:




Here GPT-3 is 'on a roll' – in an effort to maintain continuity and return improved results as the session develops, it has trained on my previous prompts, and failed to notice that I am now looking for Einstein quotes, and not (as with the previous prompt) for Winston Churchill quotes.

## Other GPT-3 Products And Demonstrations

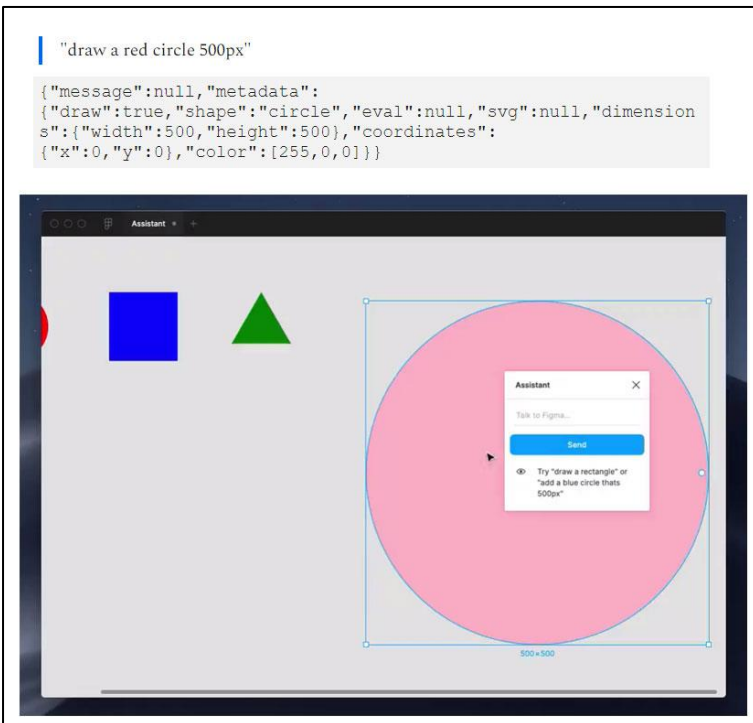
### Generating App and Site Designs Via GPT-3

Shortly after the release of API access, New York-based product designer Jordan Singer developed a GPT-3 plugin capable of passing instructions to the proprietary vector graphics editor Figma.



I build my ideas #8 - 07/19/20', <https://ibuildmyideas.substack.com/p/i-build-my-ideas-8-071920>

Called Designer, the plugin communicates with the Figma work space, converting descriptions into vector representations of a described shape. Among its other programming capabilities, GPT-3 has assimilated the postscript language, which represents geometry by systematically describing it in a mathematical space. Designer converts plain text instructions such as 'draw a red circle 500px' into a corresponding postscript command, and then wraps it as a JSON to pass it to Figma.



## GPT-3 Startup Pitches

GPT-3 is capable of generating ideas for company startups, with minimal prompting.

**Q: Invent ten fictitious startup pitches for services related to NLP. Reply with a bulleted list, where each entry is a startup suggestion.**

**A:**

1.  
A company that trains people to speak more eloquently by providing them with a vocal coach and feedback.
2.  
A company that analyzes the content of things people say to find out what they are really feeling, and provides advice on how to respond.
3.  
A company that tells you what someone is feeling based on their facial expressions and body language.
4.  
A company that provides personalized language lessons for people who want to learn a new language or improve their proficiency in their native language.
5.  
A company that creates personalized books for children based on their individual interests, personality, and learning style.
6.  
A company that creates personalized music playlists based on what mood the listener is in, their favorite artists, and what genre they are interested in.
7.  
A company that provides text-to-speech software for people who cannot read because of blindness or visual impairment.
8.  
A company that provides software to help people with dyslexia read more easily by reversing the words on the screen.
9.  
A company that provides software to help people learn how to type without looking at the keyboard by memorizing the keyboard layout and typing by touch-typing.
10.  
A company that provides a service where people can text a question and receive a response from an expert in their field within seconds.

In December of 2020 one outlet published a list of genuine startups, mixed in with fictitious startups suggested by GPT-3 (which had been briefly trained on suitable real-life examples in the course of an API session), challenging readers to distinguish them.

Though the best results were included in the feature, the unedited list of GPT-3 suggestions is available on Google Docs, together with the 'real world' examples that informed them. Many of the ideas are imitative of real proposals, while others lack any particular 'angle'.

The IdeasAI site allows subscribers to receive GPT-3 startup ideas, which are generated daily, in an email. However, once again, the pitches are frequently vague, and some of them are simply random excerpts from related web articles.



Idea	Score
A collection of online educational materials available for free. <small>Generated by GPT-3, 27 days ago.</small>	125
The company has raised a total of \$11.5 million. <small>Generated by GPT-3, 16 days ago.</small>	104
A way to help people navigate the process of buying a home. <small>Generated by GPT-3, 18 days ago.</small>	83
A platform for scientists to share and discover data. <small>Generated by GPT-3, 25 days ago.</small>	60
The company was founded by Ramy El Bannaoui and Chris Curtis, who had previously founded the company in 2007 and sold it in 2011. <small>Generated by GPT-3, 19 days ago.</small>	58

<https://ideasai.net/>

Nonetheless, it's reasonable to assume that dedicated training and more targeted data sampling may hold considerable potential for generating more innovative and off-beat startup concepts across a number of sectors.

## GPT-3 For Search & Recommender Systems

The kind of domain understanding that GPT-3 offers makes it possible to perform 'elliptical' searches, where even if none of the words or phrases in a query are found, their intent can be discerned, and relevant results obtained.



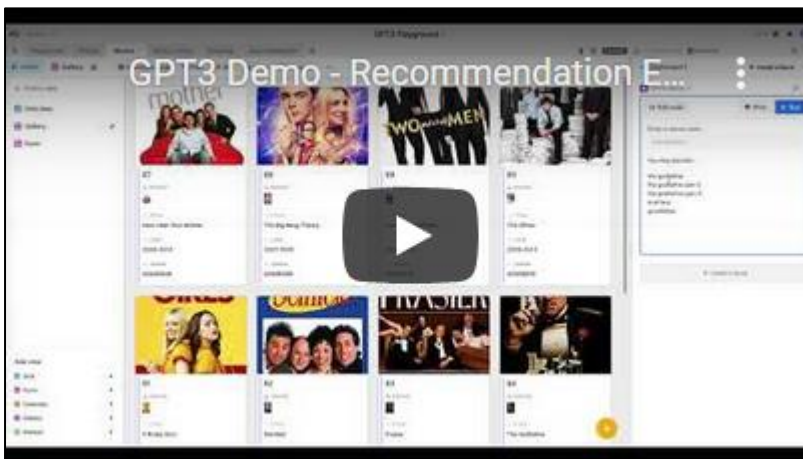
Where a Wikipedia search fails to find a search term, a GPT-3-powered semantic search can recognize text that's relevant to it. <https://beta.openai.com/>

Algolia's suite of commercial search tools incorporates a number of machine learning technologies, including GPT-3, to facilitate semantic search on websites, or on any corpus of information that can be indexed.

There's nothing new about semantic search in itself – In a 2013 update to its search algorithm, entitled Hummingbird, Google began to leverage its own semantic Knowledge Graph, which had debuted the previous year, and which operates similarly to GPT-3 in this aspect.

Though the search giant did not entirely abandon keyword/keyphrase frequency as an index of relevance (and has not abandoned them since), it was clear that the future of search might lie in generalized domain understanding, where a lexicon of related terms and images would constitute a 'cloud' of essential information on a topic; and not in merely parsing text and counting occurrence frequencies.

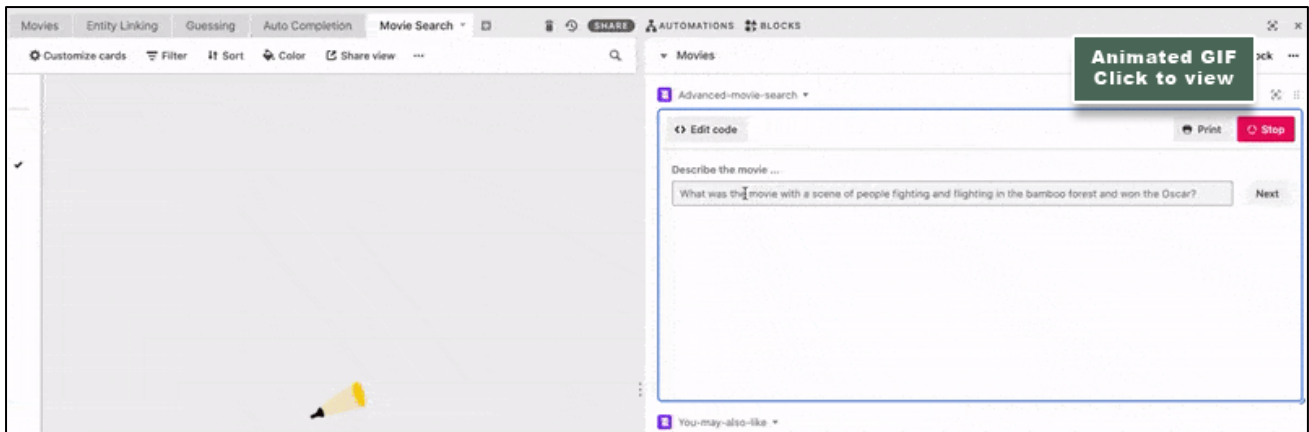
What is new is that accessible and increasingly affordable machine learning infrastructure makes it possible for smaller players to implement truly effective semantic search, either for the purpose of developing new generalized search engines based on web-crawling, or for more specific functions, such as recommender engines.



Video: 'GPT3 Demo - Recommendation Engine' - <https://www.youtube.com/watch?v=qca2lmQZeoo>

The above movie and TV recommendation system utilizes GPT-3 and metadata from the IMDB.

The system can also identify a movie from a description that may contain few or no keywords in the target entry:



'GPT3 Demo - Advanced search', SOTA, YouTube, 10th September 2020 - <https://www.youtube.com/watch?v=Vg4lky4rbAM>

For entry-level startups, the democratizing power of AI-driven semantic search via GPT-3 is perhaps one of the most exciting aspects of OpenAI's language model.

## Conclusion

GPT-3 is a powerful tool that, intelligently applied, can help create genuinely useful and innovative new products and services that were not financially or technologically feasible before.

Though its core strength lies in its power to manipulate domain knowledge (without entirely abstracting it), some of the critical publicity it has received since launch has instead focused on what it does not yet do well, such as generating long and coherent treatises with minimal supervision or input.

However, that's a far broader challenge for the NLP community, and one that's ongoing; nor did GPT-3 ever advertise itself as a novelist or long-form feature writer, though it does have some nascent ability in these fields.

Plugging a raw GPT-3 API call into a new product will work best if your service concentrates on those aspects that GPT-3 can perform well without additional training, such as programmatic analysis. However, there's no real value proposition in discovering and re-packaging GPT-3's native capabilities, since a) it's trivial for any competing product to do likewise, and b) native calls can still be unpredictable, even when querying in a limited domain.

The real potential of GPT-3 is as an adjunct or enabling technology for proprietary systems that train GPT-3 against other dedicated datasets that support your product — as evidenced by the best of some of the services that we've taken a look at in this article.

*\* This article was originally published at the scalr.ai domain on 27th January, 2021, though it has now migrated to the width.ai domain, with a new official publication date of November 10th.*

## References:

SOURCE: <https://arxiv.org/pdf/1808.01075.pdf>  
SOURCE: <https://arxiv.org/pdf/1811.12799.pdf>  
SOURCE: [http://www.evaascarza.com/padilla\\_ascarza\\_20.pdf](http://www.evaascarza.com/padilla_ascarza_20.pdf)  
SOURCE: <https://www.amazon.science/blog/amazon-com-tests-customer-service-chatbots>  
SOURCE: <https://yellrobot.com/shoplifting-detection-software-vaak/>  
SOURCE: <https://arxiv.org/pdf/1603.06289.pdf>  
<https://www.forbes.com/sites/bryanpearson/2018/03/15/3-ways-retailers-can-use-facial-recognition-to-express-better-experiences/>  
<https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/leaderships-role-in-fixing-the-analytics-models-that-covid-19-broke>  
<https://www.washingtonpost.com/business/2020/07/29/masks-facial-recognition-error-rates/>  
<https://www.thedrum.com/opinion/2021/01/13/how-covid-19-changed-e-commerce-sales-growth-and-irreversible-dependence>  
<https://www.wsj.com/articles/retailers-use-ai-to-improve-online-recommendations-for-shoppers-11604330308>  
<https://aws.amazon.com/blogs/machine-learning/using-a-b-testing-to-measure-the-efficacy-of-recommendations-generated-by-amazon-personalize/>  
<https://www.delish.com/food/news/a39015/facts-about-chocolate/>  
<https://www.kdnuggets.com/2019/05/churn-prediction-machine-learning.html>  
<https://core.ac.uk/download/pdf/83461632.pdf>  
<https://arxiv.org/ftp/arxiv/papers/1904/1904.00690.pdf>  
[https://link.springer.com/chapter/10.1007/978-3-030-58808-3\\_42](https://link.springer.com/chapter/10.1007/978-3-030-58808-3_42)  
<https://onlinelibrary.wiley.com/doi/abs/10.1002/smj.3215>  
<https://towardsdatascience.com/tagged/customer-lifetime-value>  
<https://towardsdatascience.com/rnns-vs-traditional-ml-for-predictive-customer-lifetime-value-42b6b12558fe>  
<https://arxiv.org/ftp/arxiv/papers/1604/1604.05377.pdf>  
[https://www.h2o.ai/wp-content/uploads/2017/03/Case-Studies\\_PayPal.pdf](https://www.h2o.ai/wp-content/uploads/2017/03/Case-Studies_PayPal.pdf)  
<https://cloud.google.com/solutions/machine-learning/clv-prediction-with-offline-training-intro>  
<https://www.intechopen.com/books/data-mining/estimating-customer-lifetime-value-using-machine-learning-techniques>  
<https://www.diva-portal.org/smash/get/diva2:1431621/FULLTEXT02>  
<https://machinethink.net/blog/new-in-coreml3/>  
<https://www.kdd.org/kdd2016/papers/files/adf0755-vanderveldAbr.pdf>  
<https://cdn.intechopen.com/pdfs/61571.pdf>  
<https://www.rezi.io/>  
<https://www.copysmith.ai/>  
<https://sapling.ai/>  
<https://venturebeat.com/2016/07/12/whole-foods-just-launched-a-messenger-chatbot-for-finding-recipes-with-emojis/>  
<https://chatbotmagazine.com/how-we-built-the-wine-bot-margot-for-lidl-b54f42cda4dd>  
[https://www.cio.com/article/3564260/customers-bank-on-bofa-chatbot-amid-pandemic.html?upd=1593439458685&utm\\_source=Triggermail&utm\\_medium=email&utm\\_campaign=Post%20Blast%20bii-banking:%20BofA%27s%20virtual%20assistant%20usage%20spikes%20during%20pandemic%20%7C%20Lloyds%20to%20grow%20wealth%20management%20unit%20amid%20low%20interest%20rates%20%7C%20ING%20Belgium%20adds%200subscription%20management%20feature&utm\\_term=BII%20List%20Banking%20ALL](https://www.cio.com/article/3564260/customers-bank-on-bofa-chatbot-amid-pandemic.html?upd=1593439458685&utm_source=Triggermail&utm_medium=email&utm_campaign=Post%20Blast%20bii-banking:%20BofA%27s%20virtual%20assistant%20usage%20spikes%20during%20pandemic%20%7C%20Lloyds%20to%20grow%20wealth%20management%20unit%20amid%20low%20interest%20rates%20%7C%20ING%20Belgium%20adds%200subscription%20management%20feature&utm_term=BII%20List%20Banking%20ALL)  
<https://netfinance.wbresearch.com/bank-of-america-ai-powered-chatbot-strategy-ty-u>  
<https://promo.bankofamerica.com/erica/>  
<https://arxiv.org/ftp/arxiv/papers/2011/2011.03969.pdf>  
<https://www.weforum.org/agenda/2020/04/chatbots-covid-19-governance-improved-here-s-how/>

<https://www.technologyreview.com/2020/05/14/1001716/ai-chatbots-take-call-center-jobs-during-coronavirus-pandemic/>  
<https://www.bankofengland.co.uk/quarterly-bulletin/2020/2020-q4/the-impact-of-covid-on-machine-learning-and-data-science-in-uk-banking>  
<https://www.businessinsider.com/chatbot-market-stats-trends>  
<https://www.nature.com/articles/s41746-020-0280-0.pdf>  
<https://www.scmp.com/tech/start-ups/article/3038772/chinas-reliance-us-origin-platforms-deep-learning-raises-questions>  
<https://botstar.com/blog/the-rise-of-chatbots-post-covid-19-outbreak>  
<https://www.informationweek.com/big-data/ai-machine-learning/the-state-of-chatbots-pandemic-edition/d/d-id/1338841?>  
<https://www.dataversity.net/predictions-for-big-data-analytics-in-2019/>  
<https://blogs.nvidia.com/blog/2016/01/12/accelerating-ai-artificial-intelligence-gpus/>  
[http://www.evaascarza.com/padilla\\_ascarza\\_20.pdf](http://www.evaascarza.com/padilla_ascarza_20.pdf)  
<https://www.thinkwithgoogle.com/intl/en-ccc/marketing-strategies/programmatic/how-global-startup-used-first-party-data-enhance-machine-learning/>