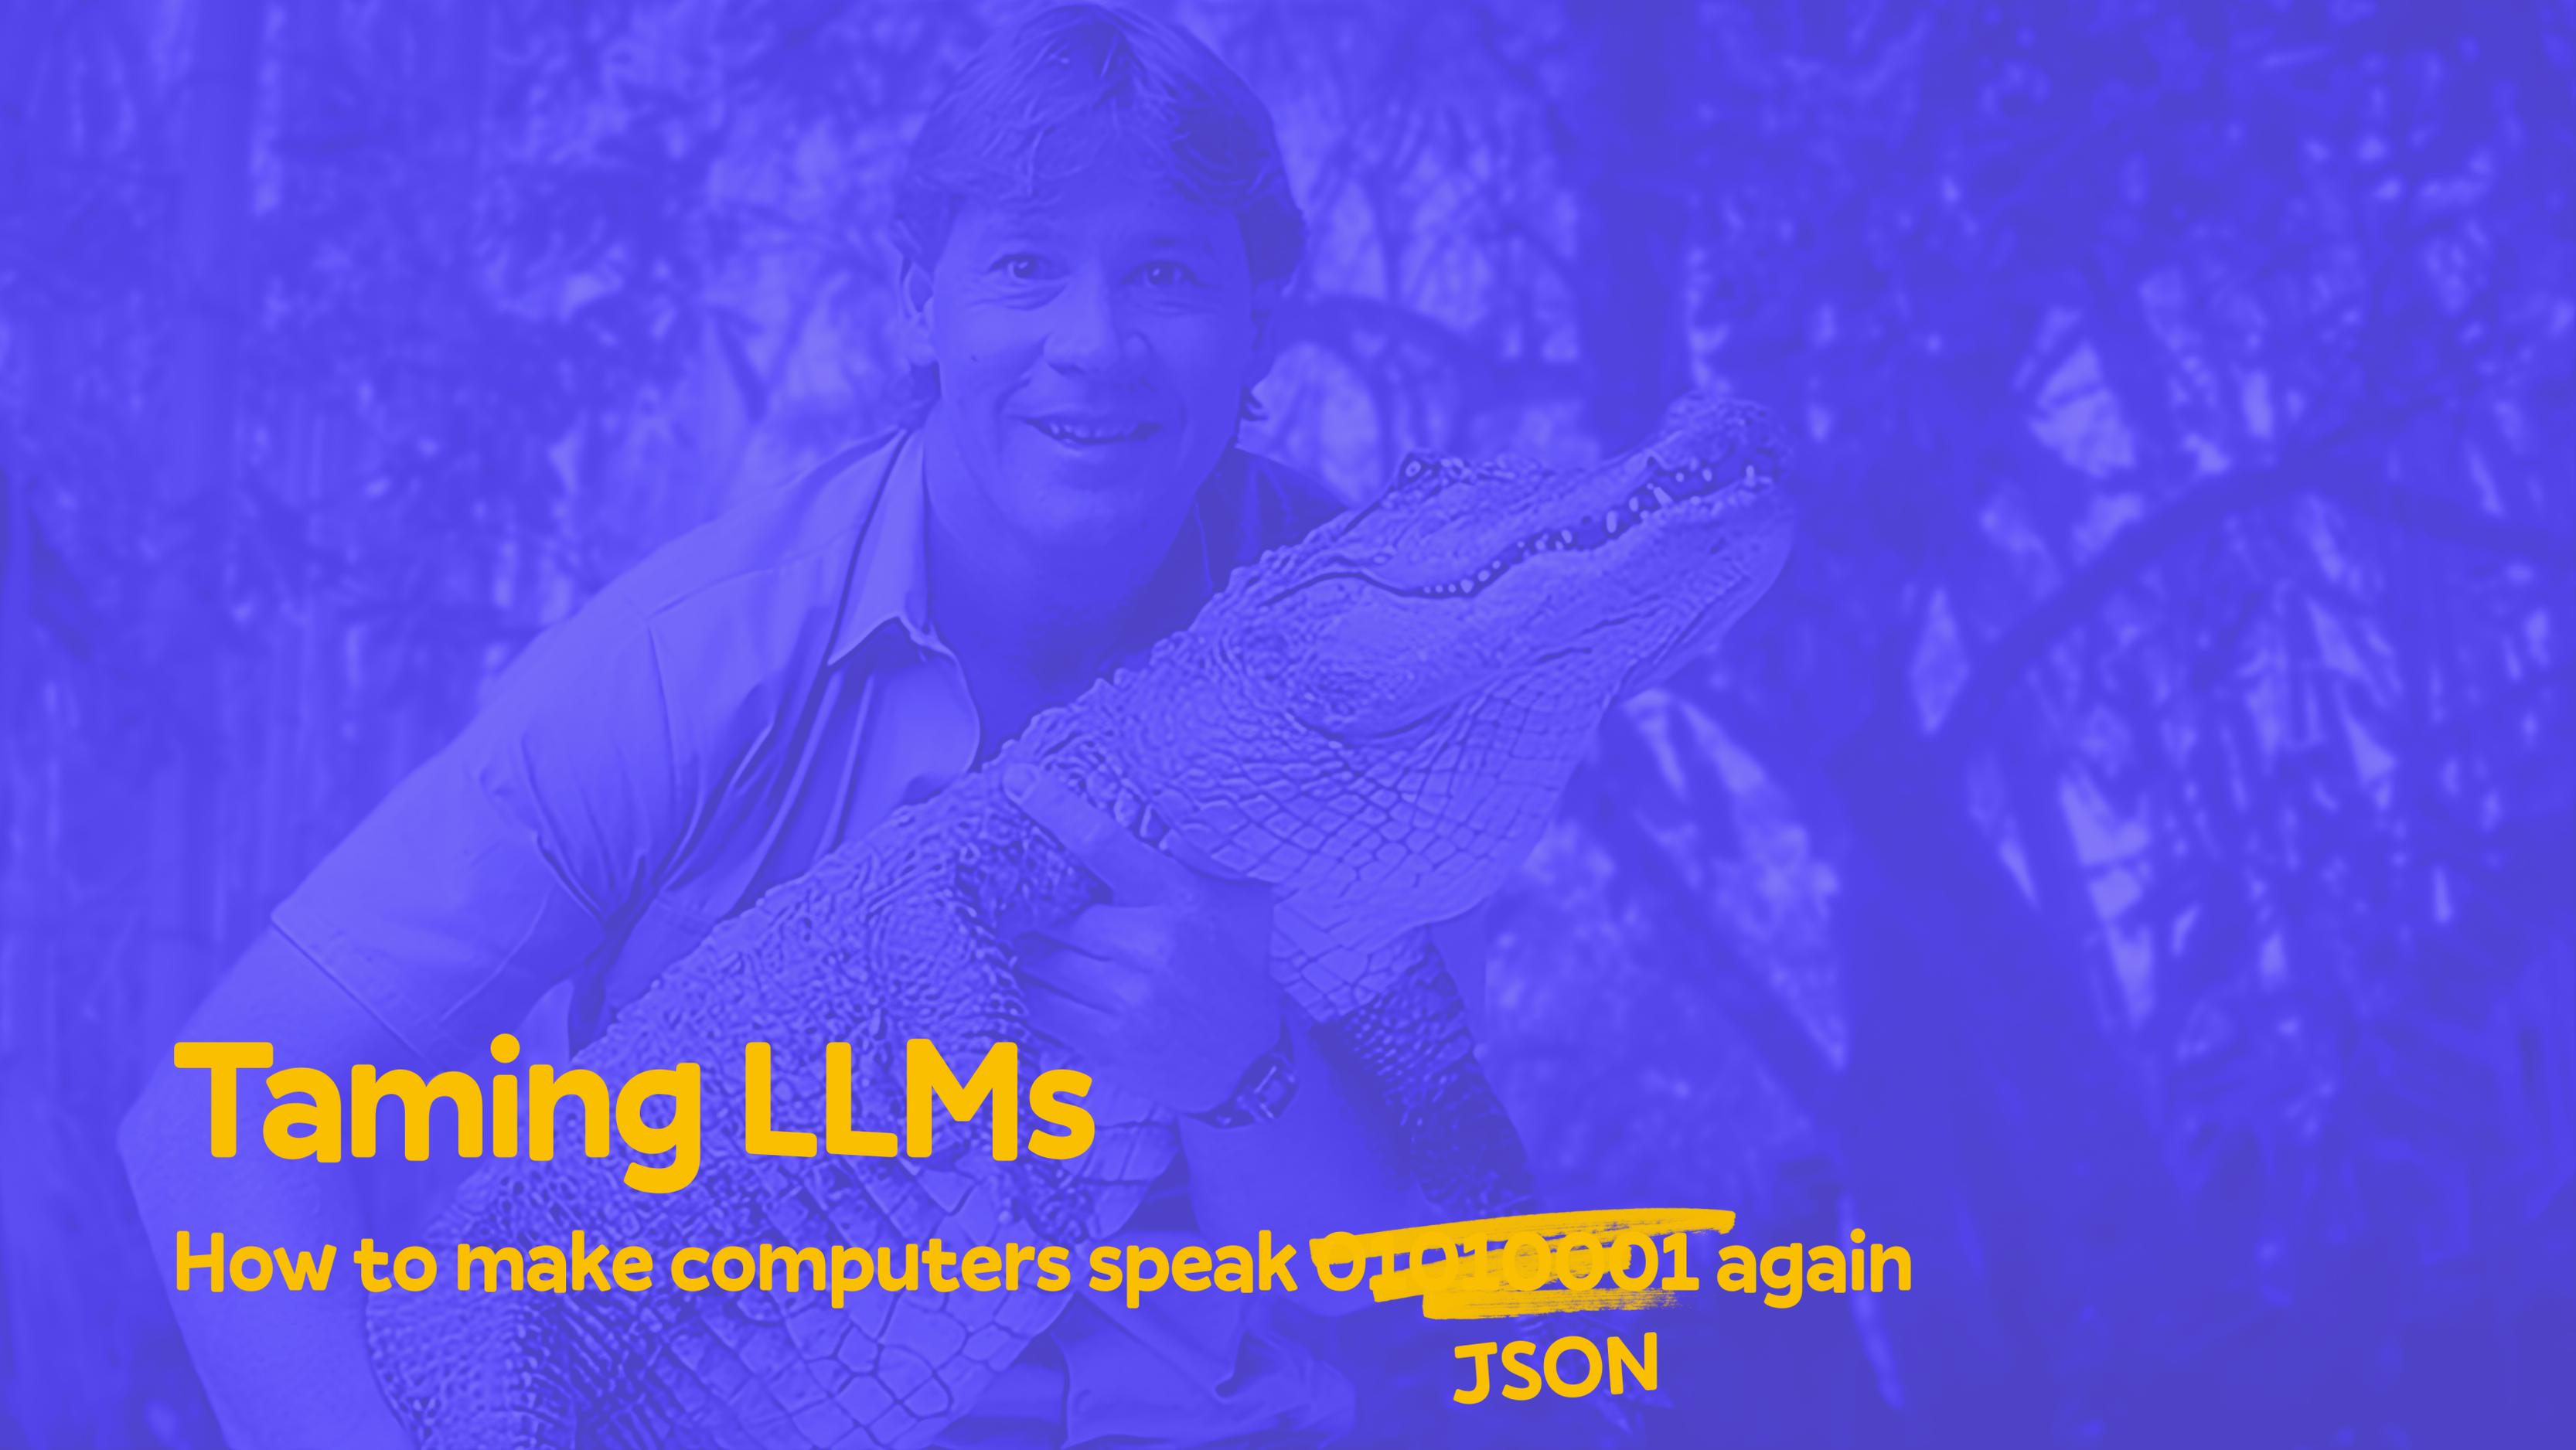


Taming LLMs

How to make computers speak 01010001 again



Taming LLMs

How to make computers speak ~~01010001~~ again

JSON

TLDR: Structured Outputs

Structured Outputs is a feature that ensures the model will always generate responses that adhere to your supplied JSON Schema, so you don't need to worry about the model omitting a required key, or hallucinating an invalid enum value.

```
1  const WildlifeEncounterSchema = z.object({
2    animal: z.string().describe("The animal species involved in the encounter"),
3    description: z.string().describe("A detailed description of the encounter, including the location, behavior, and a
4  });
5
6  const WildlifeDocumentarySchema = z.object({
7    title: z.string().describe("The main title of the documentary"),
8    episodeNumber: z.number().int().describe("The specific episode number in the documentary series"),
9    locations: z.array(z.string()).describe("The different places or regions featured in the filming"),
10   crewSize: z.number().int().describe("The total number of team members involved in production"),
11   dangerousEncounters: z.array(WildlifeEncounterSchema).describe("The noteworthy or dramatic encounters with wildlif
12   releaseDate: z.string().describe("The scheduled date when the documentary will be made available. Format: YYYY-MM-
13   catchphrase: z.string().describe("A memorable expression or tagline that sums up the spirit of the documentary")
14  });
15
16  const client = new OpenAI();
17
18  const completion = await client.beta.chat.completions.parse({
19    model: 'gpt-4o-mini',
20    messages: [
21      { role: 'system', content: 'You are a director of a TV series the Crocodile Hunter starring Steve Irwin. Giv
22      { role: 'user', content: 'Saltwater Crocodile, Komodo Dragon, Cassowary, King Cobra, and Tasmanian Devil' },
23    ],
24    tools: [zodFunction({ name: 'query', parameters: WildlifeDocumentarySchema })],
25  });
26
27  console.log(completion.choices[0].message.tool_calls[0].function.parsed_arguments);
28
```

Bartek Legięć

- Senior Product Engineer at Craft Docs / Chaps
- Programming for 10+ years
- WarsawJS community member since 2016
- Co-founder of zium.app and bodgingbear.dev



WARNING!

The talk contains a lot of **JSONs** and/or technologies that may cause discomfort and/or convulsions in people with love for **XML**.
Listener discretion is advised.

The App

We want to build an app that can build an outline for an episode of Steve Irwin's "The Crocodile Hunter" TV Series

System Prompt

Extract the animal names from the following messages. Return it as a comma separated list.

Eg. Animal1,Animal2,Animal3

User

The first one will be Saltwater Crocodile. Then Komodo Dragon, Cassowary, King Cobra, and at the very end the Tasmanian Devil.

System Response

Saltwater Crocodile,Komodo Dragon,Cassowary,King Cobra,Tasmanian Devil

Parse the output

```
completion.choices[0].message.content.split(",")
```

HTTP Calls to Wikipedia

```
{"query": "Tasmanian Devil"}
```

System Prompt

Respond in Steve Irwin's (Crocodile Hunter) enthusiastic style!

Wikipedia Article about "Saltwater crocodile":

The saltwater crocodile (*Crocodylus porosus*) is a crocodilian native to saltwater habitats, brackish wetlands and freshwater rivers from India's east coast across Southeast Asia and the Sundaland to northern Australia and Micronesia. It has been listed as Least Concern on the IUCN Red List since 1996. // Rest of the article...

Wikipedia Article about "Komodo dragon":

The Komodo dragon (*Varanus komodoensis*), also known as the Komodo monitor, is a large reptile of the monitor lizard family Varanidae that is endemic to the Indonesian islands of Komodo, Rinca, Flores, Gili Dasami, and Gili Motang. It is the largest extant species of lizard, with the males growing to a maximum length of 3 m (10 ft) and weighing up to 150 kg (330 lb). // Rest of the article...

Function calling

Function calling provides a powerful and flexible way for models to interface with your code or external services.

AI Makes HTTP Requests?

Define JS function on the input?

Function calling

Webhooks?

Magic? 🎱

```
1 {
2   role: "system",
3   content: "You are a helpful assistant that helps create user accounts.",
4 },
5 {
6   role: "user",
7   content: "Create a new user account for John Smith, a 25-year-old admin who prefers dark theme",
8 },
```

```
1  const completion = await openai.chat.completions.create({
2    model: "gpt-4o-mini",
3    messages: [/* ... */],
4    tools: [
5      {
6        type: "function",
7        function: {
8          name: "create_user",
9          description: "Create a new user account with the specified details",
10         parameters: {
11           type: "object",
12           properties: {
13             id: { type: "string", description: "The unique identifier for the user", format: "uuid" },
14             username: { type: "string", description: "The username for the new user (3-20 characters)", minLength: 3
15             /* ... */
16           },
17           required: ["id", "username", "email", "firstName", "lastName", "role"],
18         },
19       },
20     ],
21   },
22 });
```

```
1 {
2   role: "assistant",
3   content: null,
4   tool_calls: [
5     {
6       id: "call_4906dbKffJhKolUFVwHtQSI3",
7       type: "function",
8       function: {
9         name: "create_user",
10        arguments:
11          '{"username":"johnsmith","email":"john.smith@example.com","firstName":"John","lastName":"Smith","role":"ad
12        },
13      },
14    ],
15    refusal: null,
16  }
```

```
1 {
2   role: "tool",
3   content: JSON.stringify({ id: "3edf5176-bedf-4830-a44d-7860cdcfb3ee" }),
4   tool_call_id: "call_4906dbKffJhKolUFVwHtQSI3",
5 },
6 {
7   role: "assistant",
8   content:
9     "The user account for John Smith has been created successfully. His unique identifier is **3edf5176-bedf-4830-a4
10   refusal: null,
11 },
```

```
1 {
2   type: "function",
3   function: {
4     name: "create_user",
5     description: "Create a new user account with the specified details",
6     parameters: {
7       type: "object",
8       properties: {
9         id: { type: "string", description: "The unique identifier for the user", format: "uuid" },
10        username: { type: "string", description: "The username for the new user (3-20 characters)", minLength: 3, ma
11        email: { type: "string", description: "The email address of the user", format: "email" },
12        firstName: { type: "string", description: "The user's first name" },
13        lastName: { type: "string", description: "The user's last name" },
14        role: { type: "string", description: "The user's role in the system", enum: ["user", "admin", "moderator"] },
15        age: { type: "integer", description: "The user's age (optional, must be between 13 and 120)", minimum: 13, m
16      },
17      required: [
18        "id",
19        "username",
20        "email",
21        "firstName",
22        "lastName",
23        "role",
24      ],
25    },
26  },
27 }
```

```
1 {
2   role: "assistant",
3   content: null,
4   tool_calls: [
5     {
6       id: "call_4906dbKffJhKolUFVwHtQSI3",
7       type: "function",
8       function: {
9         name: "create_user",
10        arguments:
11          '{"username":"johnsmith","email":"john.smith@example.com","firstName":"John","lastName":"Smith","role":"ad
12        },
13      },
14    ],
15    refusal: null,
16  }
```

OpenAI Structured Output

```
1 // POST /v1/chat/completions
2 {
3   "model": "gpt-4o-mini",
4   "messages": [
5     { "role": "system", "content": "You are a director of a TV series the Crocodile Hunter starring Steve Irwin. Given"
6     { "role": "user", "content": "Saltwater Crocodile, Komodo Dragon, Cassowary, King Cobra, and Tasmanian Devil" },
7   ],
8   "response_format": {
9     "type": "json_schema",
10    "json_schema": {
11      "name": "query",
12      "strict": true,
13      "schema": {
14        "type": "object",
15        "properties": {
16          "title": {
17            "type": "string",
18            "description": "The main title of the documentary"
19          },
20          "episodeNumber": {
21            "type": "integer",
22            "description": "The specific episode number in the documentary series"
23          },
24          "locations": {
25            "type": "array",
26            "items": {
27              "type": "string"
28            }
29          }
30        }
31      }
32    }
33  }
```

JSON Schema

JSON Schema is a declarative language for annotating and validating JSON documents' structure, constraints, and data types.

It helps you standardize and define expectations for JSON data.

Internet Engineering Task Force
Internet-Draft
Updates: 4120 (if approved)
Intended status: Standards Track
Expires: June 8, 2010

K. Zyp, Ed.
SitePen (USA)
December 5, 2009

A JSON Media Type for Describing the Structure and Meaning of JSON
Documents
draft-zyp-json-schema-00

Abstract

JSON (JavaScript Object Notation) Schema defines the media type application/schema+json, a JSON based format for defining the structure of JSON data. JSON Schema provides a contract for what JSON data is required for a given application and how to interact with it. JSON Schema is intended to define validation, documentation, hyperlink navigation, and interaction control of JSON data.

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that

```
1 const completion = await openai.chat.completions.create({
2   model: "gpt-4o-mini",
3   messages: [
4     {
5       role: "system",
6       content: "You are a director of a TV series the Crocodile Hunter starring Steve Irwin. Given a list of animals
7     },
8     {
9       role: "user",
10      content: "Saltwater Crocodile, Komodo Dragon, Cassowary, King Cobra, and Tasmanian Devil",
11    },
12  ],
13  tool_choice: { type: "function", function: { name: "query" } },
14  tools: [
15    {
16      "type": "function",
17      "function": {
18        "name": "query",
19        "parameters": {
20          "type": "object",
21          "properties": {
22            "title": {
23              "type": "string",
24              "description": "The main title of the documentary"
25            },
26            "episodeNumber": {
27              "type": "integer",
28              "description": "The specific episode number in the documentary series"
```

Zod

```
1 import { z } from 'zod';
2
3 const WildlifeEncounterSchema = z.object({
4   animal: z.string()
5     .describe('The animal species involved in the encounter'),
6   description: z.string()
7     .describe('A detailed description of the encounter, including the location, behavior, and any other relevant
8 details'),
9 });
10
11 const WildlifeDocumentarySchema = z.object({
12   title: z.string().describe('The main title of the documentary'),
13   episodeNumber: z.number().int()
14     .describe('The specific episode number in the documentary series'),
15   locations: z.array(z.string())
16     .describe('The different places or regions featured in the filming'),
17   crewSize: z.number().int()
18     .describe('The total number of team members involved in production'),
19   dangerousEncounters: z.array(WildlifeEncounterSchema)
20     .describe('The noteworthy or dramatic encounters with wildlife that add risk and excitement'),
21   releaseDate: z.string()
22     .describe('The scheduled date when the documentary will be made available. Format: YYYY-MM-DD'),
23   catchphrase: z.string()
24     .describe('A memorable expression or tagline that sums up the spirit of the documentary'),
25 });
```

```
1 import { zodFunction } from 'openai/helpers/zod';
2
3 const WildlifeEncounterSchema = z.object({
4   animal: z.string().describe("The animal species involved in the encounter"),
5   description: z.string().describe("A detailed description of the encounter, including the location, behavior, and a
6 });
7
8 const WildlifeDocumentarySchema = z.object({
9   title: z.string().describe("The main title of the documentary"),
10  episodeNumber: z.number().int().describe("The specific episode number in the documentary series"),
11  locations: z.array(z.string()).describe("The different places or regions featured in the filming"),
12  dangerousEncounters: z.array(WildlifeEncounterSchema).describe("The noteworthy or dramatic encounters with wildlif
13  releaseDate: z.string().describe("The scheduled date when the documentary will be made available. Format: YYYY-MM-
14  catchphrase: z.string().describe("A memorable expression or tagline that sums up the spirit of the documentary")
15 });
16
17 const client = new OpenAI();
18
19 const completion = await client.beta.chat.completions.parse({
20   model: 'gpt-4o-mini',
21   messages: [
22     { role: 'system', content: 'You are a director of a TV series the Crocodile Hunter starring Steve Irwin. Giv
23     { role: 'user', content: 'Saltwater Crocodile, Komodo Dragon, Cassowary, King Cobra, and Tasmanian Devil' },
24   ],
25   tools: [zodFunction({ name: 'query', parameters: WildlifeDocumentarySchema })],
26 });
27
28 console.log(completion.choices[0].message.tool_calls[0].function.parsed_arguments);
```

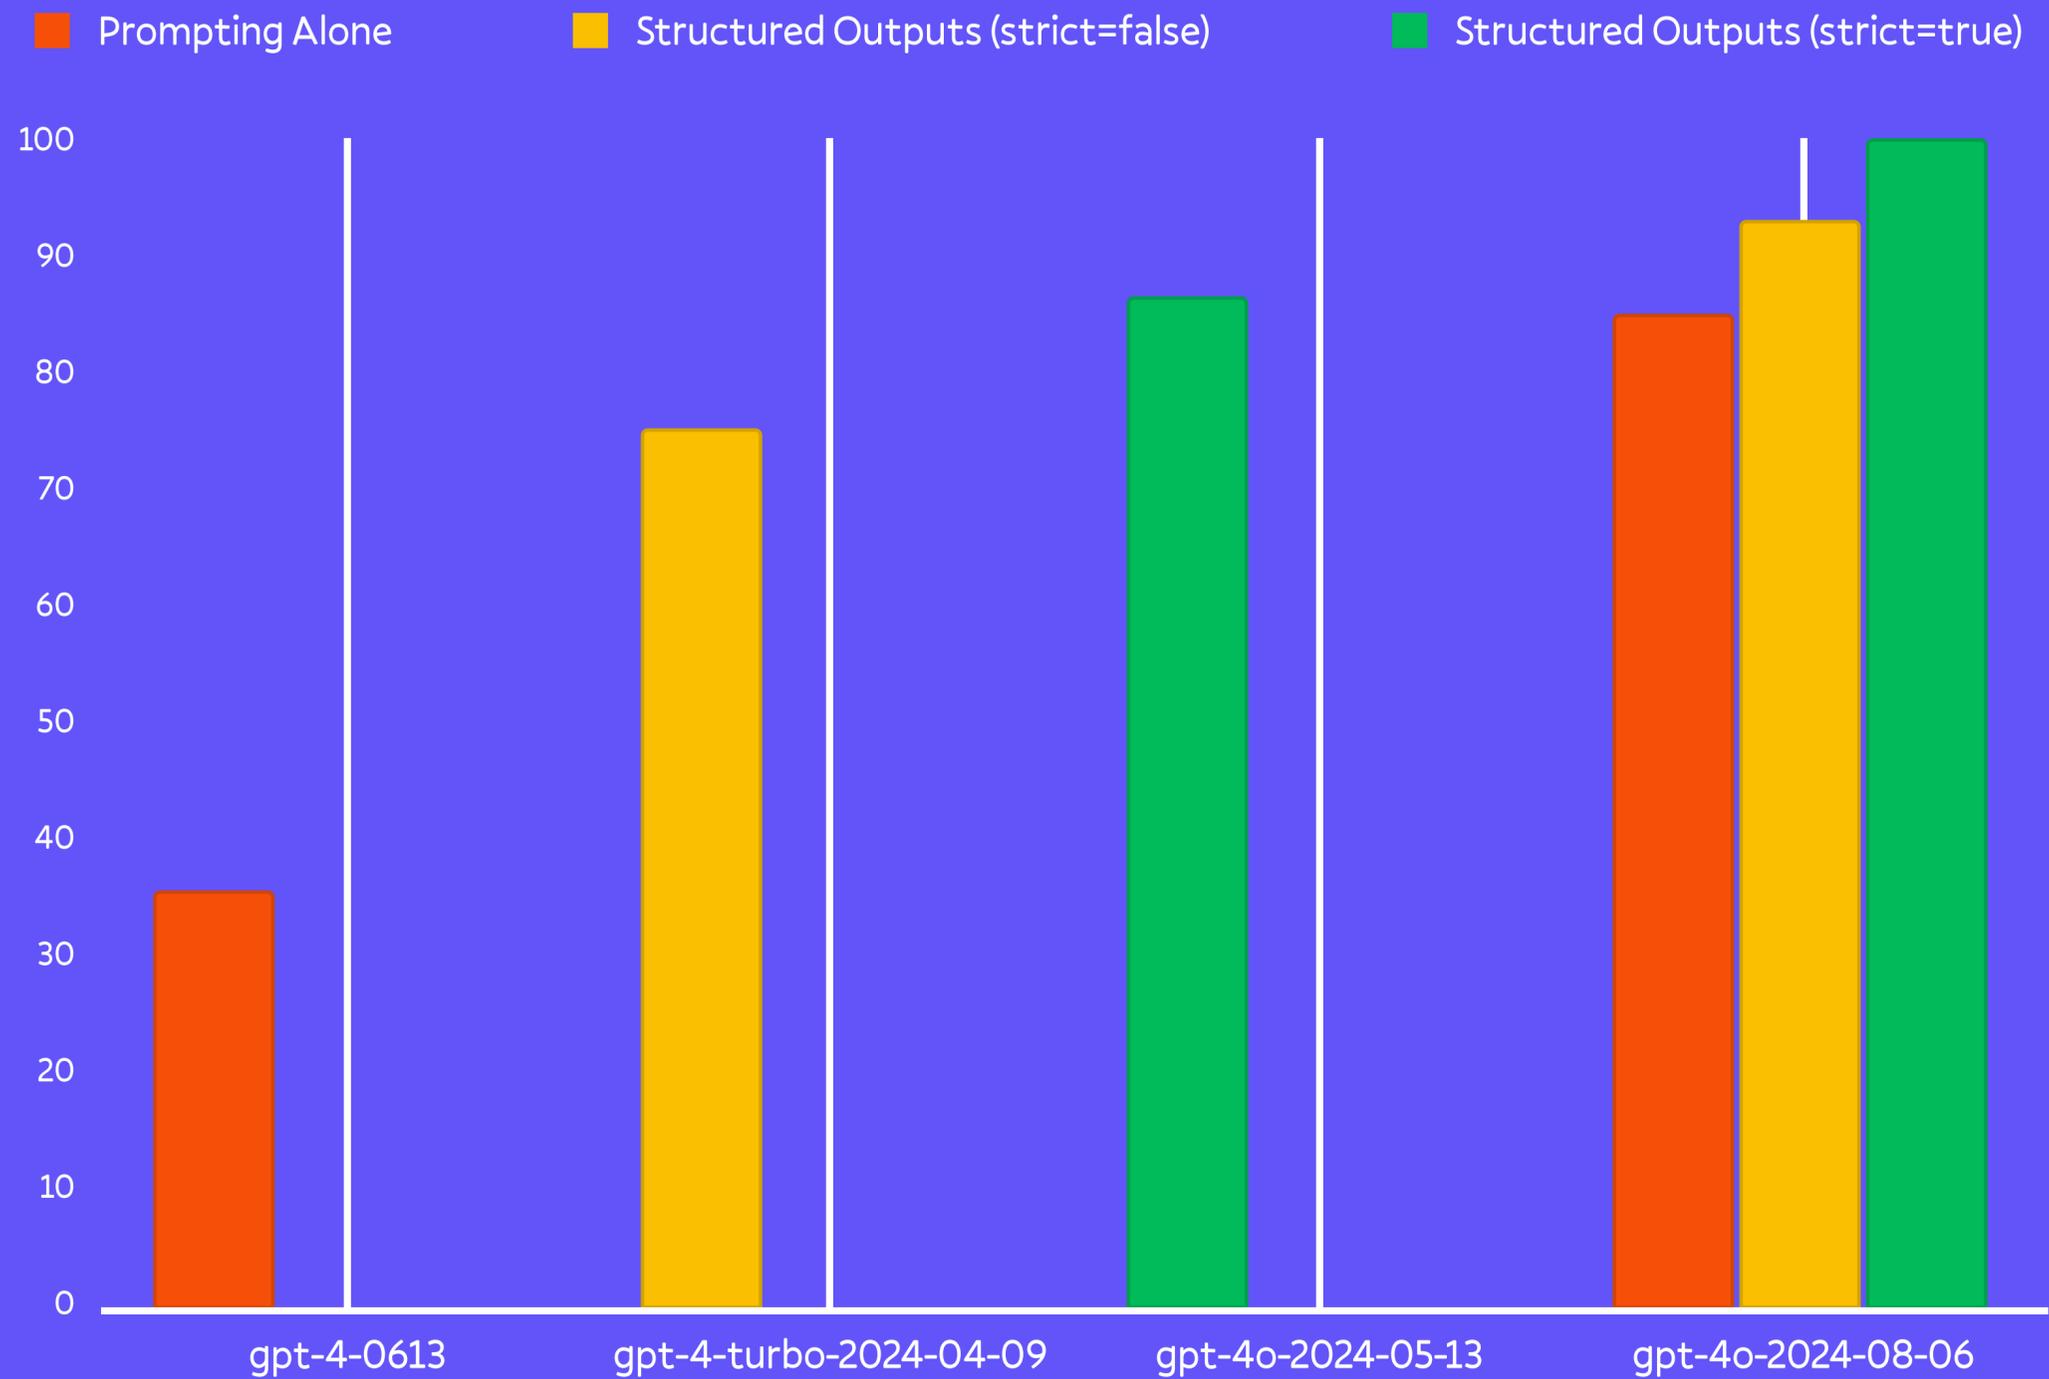
zod-to-json-schema

```
import { z } from "zod";
import { zodToJsonSchema } from "zod-to-json-schema";

const mySchema = z
  .object({
    myString: z.string().min(5),
    myUnion: z.union([z.number(), z.boolean()]),
  })
  .describe("My neat object schema");

const jsonSchema = zodToJsonSchema(mySchema, "mySchema");
```

```
{
  "$schema": "http://json-schema.org/draft-07/schema#",
  "$ref": "#/definitions/mySchema",
  "definitions": {
    "mySchema": {
      "description": "My neat object schema",
      "type": "object",
      "properties": {
        "myString": {
          "type": "string",
          "minLength": 5
        },
        "myUnion": {
          "type": ["number", "boolean"]
        }
      },
      "additionalProperties": false,
      "required": ["myString", "myUnion"]
    }
  }
}
```



With Structured Outputs, gpt-4o-2024-08-06 achieves 100% reliability in our evals, perfectly matching the output schemas.

zod-stream

Sprawdź z AI

Która z poniższych nie jest technologią działającą po stronie serwera:

NodeJS

BackboneJS

Mongo

ExpressJS

Sprawdź

Elements Console Sources Network Performance Memory Application >>

top Filter Default levels No Issues 1 hidden

>

```
1  const ExplanationSchema = z.object({
2    answerIndex: z.number().describe('0 based index of an answer'),
3    explanation: z
4      .string()
5      .describe(
6        'If you see code or SQL, wrap it in code block. Please make sure to use the correct language for the code block.',
7      ),
8  });
9
10 const AiTeacherResponseSchema = z.object({
11   questionLanguage: QuestionLanguageSchema,
12   correctAnswersIndex: CorrectAnswersIndexSchema,
13   correctExplanations: z.array(ExplanationSchema),
14   incorrectExplanations: z.array(ExplanationSchema),
15 });
```

@instructor-ai/instructor-js

```
1 import Instructor from "@instructor-ai/instructor"
2 import OpenAI from "openai"
3 import { z } from "zod"
4
5 const UserSchema = z.object({
6   age: z.number(),
7   name: z.string().refine(name => name.includes(" "))
8 })
9
10 const client = new OpenAI({
11   baseURL: "https://api.together.xyz/v1",
12   apiKey: process.env.TOGETHER_API_KEY
13 })
14
15 const instructor = Instructor({
16   client: client,
17   mode: "TOOLS"
18 })
19
20 const user = await instructor.chat.completions.create({
21   messages: [{ role: "user", content: "Jason Liu is 30 years old" }],
22   model: "mistralai/Mixtral-8x7B-Instruct-v0.1",
23   response_model: { schema: UserSchema, name: "User" },
24   max_retries: 4
25 })
26
27 console.log(user);
```

chaps

Apps you can chat with.





bibixx



bibix1999



bibixx.com



bartoszlegiec



legiec.io/taming-llms

