IT Services

India

Sector View: **Neutral** NIFTY-50**: 24,596** **August 08, 2025**

GPT-5 unleashed: Presents risks and opportunities

OpenAI's GPT-5 might not present a step-jump in capabilities such as GPT-4, but it claims to offer a meaningful improvement in reasoning and coding tasks, among others. Hallucination rates have dropped significantly, a positive for gen AI and agentic AI adoption. GPT-5 will increase the adoption of gen AI in software development, exposing IT services firms to revenue deflation risks, but it is also a step forward in opening up new opportunities from (1) cloud and data foundation, (2) legacy modernization and (3) AI for business use cases.

GPT-5 lands finally but not a step jump compared to recent predecessors

OpenAI released GPT-5 series of models after much anticipation. Models have 400k context length, with 128k max output tokens. The models incorporate elements of predecessors such as reasoning but do not seem to possess capabilities that could be categorized as a step-jump. The models carry several improvements that aim to provide a better user experience. GPT-5 claims significant advances in reducing hallucinations, improving instruction following, and minimizing sycophancy, while leveling up performance in three of ChatGPT's most common uses: writing, coding and health.

Presents improvements in coding tasks

GPT-5 is OpenAI's strongest coding model to date. OpenAI claims particular improvements in complex front-end generation and debugging larger repositories and notes that it is state-of-the-art (SOTA) across key coding benchmarks, scoring 74.9% on SWE-bench Verified and 88% on Aider polyglot. According to OpenAI, GPT-5 was trained to be a true coding collaborator. It excels at producing high-quality code and handling tasks such as fixing bugs, editing code and answering questions about complex codebases. The model is steerable and collaborative—it can follow very detailed instructions with high accuracy and can provide upfront explanations of its actions before and between tool calls. GPT-5 also excels at long-running agentic tasks—achieving SOTA results on τ2-bench telecom (96.7%), a tool-calling benchmark released just two months ago. GPT-5's improved tool intelligence lets it reliably chain together dozens of tool calls—both in sequence and in parallel—without losing its way, making it far better at executing complex, real-world tasks end to end. It also follows tool instructions more precisely, is better at handling tool errors, and excels at long-context content retrieval. Prominent code editors such as Cursor and Windsurf have rated GPT-5 highly relative to others (Exhibit 2).

Reduction in hallucination rates and lower pricing are positives for AI adoption

Hallucination rates have reduced significantly (Exhibit 1) compared to predecessors and thus will help quell worries of hallucinations increasing with an increase in model capabilities. Lower hallucinations help increase the reliability of generative AI tools and are a positive for adoption across use cases, especially in agentic AI. Improvements in other categories such as intelligence also help in this regard. Pricing per token is lower compared to frontier models from other key US-based companies such as Anthropic and Google (Exhibit 3).

Kawaljeet Saluja
kawaljeet.saluja@kotak.com
+91-22-4336-0860

Sathishkumar S
sathishkumar@kotak.com
+91-22-4336-0879

Vamshi Krishna
vamshi.krishna@kotak.com
+91-22-6166-1801

Raises risks and opportunities for IT services—expect net headwinds in the near term

We believe that GPT-5 is another step toward increasing generative AI adoption among enterprises. Given (1) the focus of AI labs to increase coding capabilities, (2) the willingness of enterprises to adopt generative AI in software development, (3) the increasing developer usage of generative AI tools and (4) the higher focus of enterprises to get productivity from AI adoption, we believe that the adoption of generative AI in software development will increase and present revenue deflation risks for IT services providers. The same applies for customer service BPO as well.

At the same time, reduction in hallucinations, lower pricing, better reasoning and intelligence, among other improvements, also open the door for more and faster adoption of AI for business use cases. These can present opportunities for Indian IT. They can help enterprises set up cloud and data foundation required to reap the full value of generative AI. Effort, costs and risks around legacy modernization can be reduced, leading to accelerated modernizations. IT services firms can be partners providing end-to-end services for viable AI-enabled use cases for business. Higher velocity and lower costs of software development can catalyse higher custom application builds.

We believe that new opportunities from generative AI adoption will offset revenue deflation in existing volumes over time. However, we expect a lag in the pickup of new opportunities and for savings in software development to be deployed into these new opportunities, leading to a period of net headwinds. We believe gen AI can impact revenue growth for Indian IT by 2-3% for a period of 2-3 years on a net basis. Our current revenue growth assumptions build in some portion of this impact already.

Can boost custom app development, providing an alternative to SaaS; may require push from Indian IT

Increased velocity and lower costs of software development bode well for custom application development. Custom apps help enterprises utilize contextual knowledge and enable differentiation as opposed to SaaS, which promotes standardization. Enterprises will also benefit from lower security/data issues, lower vendor lock-in risks and higher customization possibilities from reduced reliance on SaaS. Boosting custom app development might require evangelism on the part of Indian IT vendors while not offending SaaS partners.

We expect IT services to continue to be relevant even in a future with high gen AI adoption

At a fundamental level, service provides one-on-one customizable solutions compared to software/AI, which provides a one-to-many solution. One can argue that a sufficiently capable AI can itself provide customized solutions, provided it is well fed with enterprise data and context and can take on the role played by services. We do not expect everything in the enterprise to be measurable and hence possible to be given as an input to AI or for what is measured to be truly representative of the underlying due to the potential of it being gamed. There will continue to be scope for services, perhaps it can get redefined over a period of time.

The value that gets captured by services compared to other technology elements (software/cloud infra/AI model providers, etc.) will be a key point of debate, though. We do note that improvements in generative AI capabilities by an AI lab tends to get quickly replicated by others despite the high costs involved and scarcity of quality AI talent. There is intense competition among at least 5-6 players globally (mainly in the US and China), ensuring there is no lasting moat. Improvements in hardware will also likely lead to a reduction in training and inference costs. Such an environment is ideal for services, since it will ensure that AI models do not end up capturing disproportionately high value (assuming perfect competition).

Exhibit 1: Performance comparison of GPT-5 and predecessor models on various benchmarks

Detailed benchmarks

Intelligence

	GPT-5 (high)	GPT-5 mini (high)	GPT-5 nano (high)	OpenAI o3 (high)	OpenAI o4-mini (high)	GPT-4.1	GPT-4.1 mini	GPT-4.1 nano
AIME '25 (no tools)	94.6%	91.1%	85.2%	86.4%	92.7%	46.4%	40.2%	-
FrontierMath (with python tool only)	26.3%	22.1%	9.6%	15.8%	15.4%	-	-	-
GPQA diamond (no tools)	85.7%	82.3%	71.2%	83.3%	81.4%	66.3%	65.0%	50.3%
HLE[1] (no tools)	24.8%	16.7%	8.7%	20.2%	14.7%	5.4%	3.7%	-
HMMT 2025 (no tools)	93.3%	87.8%	75.6%	81.7%	85.0%	28.9%	35.0%	-

[1] There is a small discrepancy with numbers reported in our previous blog post, as those were run on a former version of HLE.

Multimodal

	GPT-5 (high)	GPT-5 mini (high)	GPT-5 nano (high)	OpenAI o3 (high)	OpenAI o4-mini (high)	GPT-4.1	GPT-4.1 mini	GPT-4.1 nano
MMMU	84.2%	81.6%	75.6%	82.9%	81.6%	74.8%	72.7%	55.4%
MMMU-Pro (avg across standard and vision sets)	78.4%	74.1%	62.6%	76.4%	73.4%	60.3%	58.9%	33.0%
CharXiv reasoning (python enabled)	81.1%	75.5%	62.7%	78.6%	72.0%	56.7%	56.8%	40.5%
VideoMMMU, max frame 256	84.6%	82.5%	66.8%	83.3%	79.4%	60.9%	55.1%	30.2%
ERQA	65.7%	62.9%	50.1%	64.0%	56.5%	44.3%	42.3%	26.5%

Coding

	GPT-5 (high)	GPT-5 mini (high)	GPT-5 nano (high)	OpenAI o3 (high)	OpenAI o4-mini (high)	GPT-4.1	GPT-4.1 mini	GPT-4.1 nano
SWE-Lancer: IC SWE Diamond Freelance Coding Tasks	$112K	$75K	$49K	$86K	$66K	$34K	$31K	$9K
SWE-bench Verified[2]	74.9%	71.0%	54.7%	69.1%	68.1%	54.6%	23.6%	-
Aider polyglot (diff)	88.0%	71.6%	48.4%	79.6%	58.2%	52.9%	31.6%	6.2%

[2] We omit 23/500 problems that could not run on our infrastructure. The full list of 23 tasks omitted are 'astropy__astropy-7606', 'astropy__astropy-8707', 'astropy__astropy-8872', 'tjango__django-10097', 'django__django-7530', 'matplotlib__matplotlib-20488', 'matplotlib__matplotlib-20676', 'matplotlib__matplotlib-20826', 'matplotlib__matplotlib-23299', 'matplotlib__matplotlib-24970', 'matplotlib__matplotlib-25479', 'matplotlib__matplotlib-26342', 'psf__requests-6028', 'pylint-dev__pylint-6528', 'pylint-dev__pylint-7080', 'pylint-dev__pylint-7277', 'pytest-dev__pytest-5262', 'pytest-dev__pytest-7521', 'scikit-learn__scikit-learn-12973', 'sphinx-doc__sphinx-10466', 'sphinx-doc__sphinx-7462', 'sphinx-doc__sphinx-8265', and 'sphinx-doc__sphinx-9367'.

Source: OpenAI

Exhibit 2: Performance comparison of GPT-5 and predecessor models on various benchmarks (contd.)

Instruction Following

	GPT-5 (high)	GPT-5 mini (high)	GPT-5 nano (high)	OpenAI o3 (high)	OpenAI o4-mini (high)	GPT-4.1	GPT-4.1 mini	GPT-4.1 nano
Scale multichallenge[3] (o3-mini grader)	69.6%	62.3%	54.9%	60.4%	57.5%	46.2%	42.2%	31.1%
Internal API instruction following eval (hard)	64.0%	65.8%	56.1%	47.4%	44.7%	49.1%	45.1%	31.6%
COLLIE	99.0%	98.5%	96.9%	98.4%	96.1%	65.8%	54.6%	42.5%

[3] Note: we find that the default grader in MultiChallenge (GPT-4o) frequently mis-scores model responses. We find that swapping the grader to a reasoning model, like o3-mini, improves accuracy on grading significantly on samples we've inspected.

Function Calling

	GPT-5 (high)	GPT-5 mini (high)	GPT-5 nano (high)	OpenAI o3 (high)	OpenAI o4-mini (high)	GPT-4.1	GPT-4.1 mini	GPT-4.1 nano
Tau^2-bench airline	62.6%	60.0%	41.0%	64.8%	60.2%	56.0%	51.0%	14.0%
Tau^2-bench retail	81.1%	78.3%	62.3%	80.2%	70.5%	74.0%	66.0%	21.5%
Tau^2-bench telecom	96.7%	74.1%	35.5%	58.2%	40.5%	34.0%	44.0%	12.1%

Long Context

	GPT-5 (high)	GPT-5 mini (high)	GPT-5 nano (high)	OpenAI o3 (high)	OpenAI o4-mini (high)	GPT-4.1	GPT-4.1 mini	GPT-4.1 nano
OpenAI-MRCR: 2 needle 128k	95.2%	84.3%	43.2%	55.0%	56.4%	57.2%	47.2%	36.6%
OpenAI-MRCR: 2 needle 256k	86.8%	58.8%	34.9%	-	-	56.2%	45.5%	22.6%
Graphwalks bfs <128k	78.3%	73.4%	64.0%	77.3%	62.3%	61.7%	61.7%	25.0%
Graphwalks parents <128k	73.3%	64.3%	43.8%	72.9%	51.1%	58.0%	60.5%	9.4%
BrowseComp Long Context 128k	90.0%	89.4%	80.4%	88.3%	80.0%	85.9%	89.0%	89.4%
BrowseComp Long Context 256k	88.8%	86.0%	68.4%	-	-	75.5%	81.6%	19.1%
VideoMME (long, with subtitle category)	86.7%	78.5%	65.7%	84.9%	79.5%	78.7%	68.4%	55.2%

Source: OpenAI

Exhibit 3: Performance comparison of GPT-5 and predecessor models on various benchmarks (contd.)

Hallucinations

	GPT-5 (high)	GPT-5 mini (high)	GPT-5 nano (high)	OpenAI o3 (high)	OpenAI o4-mini (high)	GPT-4.1	GPT-4.1 mini	GPT-4.1 nano
LongFact-Concepts hallucination rate (no tools) [lower is better]	1.0%	0.7%	1.0%	5.2%	3.0%	0.7%	1.1%	-
LongFact-Objects hallucination rate (no tools) [lower is better]	1.2%	1.3%	2.8%	6.8%	8.9%	1.1%	1.8%	-
FActScore hallucination rate (no tools) [lower is better]	2.8%	3.5%	7.3%	23.5%	38.7%	6.7%	10.9%	-

Source: OpenAI

Exhibit 4: Comments of prominent AI code editors on GPT-5's capabilities

Cursor	"GPT-5 is the smartest coding model we've used. Our team has found GPT-5 to be remarkably intelligent, easy to steer, and even to have a personality we haven't seen in any other model. It not only catches tricky, deeply-hidden bugs but can also run long, multi-turn background agents to see complex tasks through to the finish—the kinds of problems that used to leave other models stuck. It's become our daily driver for everything from scoping and planning PRs to completing end-to-end builds"
Windsurf	Indicated that GPT-5 is SOTA on their evals and "has half the tool calling error rate over other frontier models."
Vercel	"It's the best frontend AI model, hitting top performance across both the aesthetic sense and the code quality, putting it in a category of its own'

Source: OpenAI, Kotak Institutional Equities

Exhibit 5: API pricing per million tokens of GPT-5, Google Gemini 2.5 Pro and Claude Opus 4.1

	GPT-5	Gemini 2.5 Pro	Claude Opus 4.1
Input tokens	US$1.25	US$1.25, <=200k tokens US$2.50, >200k tokens	US$15
Output tokens	US$10	US$10, <=200k tokens US$15, >200k tokens	US$75

Source: Companies, Kotak Institutional Equities

Exhibit 6: Kotak Institutional Equities: valuation summary of key Indian technology companies

Company	7-Aug-25 Price (Rs)	Rating	Mkt cap. (Rs m)	Mkt cap. (US$ m)	EPS (Rs) 2025	EPS (Rs) 2026E	EPS (Rs) 2027E	P/E (X) 2025	P/E (X) 2026E	P/E (X) 2027E	EV/EBITDA (X) 2025	EV/EBITDA (X) 2026E	EV/EBITDA (X) 2027E	RoE (%) 2025	RoE (%) 2026E	RoE (%) 2027E
Coforge	1,704	BUY	570,063	6,501	29.0	41.0	53.3	58.8	41.6	32.0	28.5	20.1	16.7	19.5	21.1	24.5
Cyient	1,187	REDUCE	131,855	1,504	55.5	58.3	68.1	21.4	20.4	17.4	10.4	10.9	9.0	11.9	10.9	11.9
HCL Technologies	1,476	REDUCE	4,004,013	45,660	63.6	63.3	69.9	23.2	23.3	21.1	14.6	14.2	13.1	25.3	24.5	25.8
Hexaware Technologies	725	BUY	441,465	5,034	19.3	23.2	27.2	37.5	31.2	26.6	23.3	20.7	16.7	23.6	25.0	26.2
Indegene	576	BUY	138,232	1,576	17.0	19.6	22.9	33.8	29.4	25.1	22.1	18.7	15.7	20.9	16.8	17.0
Infosys	1,437	BUY	5,969,274	68,071	63.3	66.7	73.2	22.7	21.5	19.6	14.5	13.7	12.6	28.6	27.9	28.9
KPIT Technologies	1,210	SELL	331,577	3,781	27.6	29.1	35.0	43.9	41.5	34.5	25.7	23.5	19.9	29.7	25.1	25.5
L&T Technology Services	4,207	REDUCE	445,811	5,084	119.4	123.5	144.5	35.2	34.1	29.1	22.0	20.9	18.1	22.2	20.3	21.2
LTIMindtree	5,038	REDUCE	1,493,199	17,028	155.7	170.1	196.2	32.4	29.6	25.7	20.9	19.6	17.0	21.5	21.1	22.0
Mphasis	2,703	REDUCE	514,194	5,864	89.9	99.1	110.8	30.1	27.3	24.4	18.7	17.2	15.6	18.4	18.9	19.9
Persistent Systems	5,183	SELL	810,543	9,243	90.0	110.5	129.2	57.6	46.9	40.1	38.9	31.7	26.7	24.8	25.1	24.9
Rategain	439	REDUCE	51,865	591	17.6	15.9	19.4	24.9	27.6	22.7	20.0	22.3	16.6	13.3	10.6	11.5
Tata Elxsi	5,865	SELL	365,302	4,166	126.0	102.8	130.5	46.5	57.0	45.0	34.8	41.2	32.1	30.0	22.1	26.9
Tata Technologies	668	SELL	271,075	3,091	16.7	17.6	19.8	40.0	38.0	33.7	27.1	28.1	25.1	19.9	19.0	19.4
TCS	3,047	BUY	11,024,313	125,717	134.2	140.5	151.8	22.7	21.7	20.1	15.7	15.3	13.8	51.2	50.2	49.8
Tech Mahindra	1,482	BUY	1,311,211	14,953	48.1	58.4	77.8	30.8	25.4	19.1	17.8	14.4	11.3	15.8	18.6	23.7
Wipro	242	SELL	2,540,237	28,968	12.5	12.7	13.3	19.3	19.1	18.2	12.0	12.0	11.2	16.5	15.5	15.7

Company	Fair Value (Rs)	O/S shares (mn)	EPS CAGR (%) 2025-27E	EPS growth (%) 2025	EPS growth (%) 2026E	EPS growth (%) 2027E	Net Profit (Rs mn) 2025	Net Profit (Rs mn) 2026E	Net Profit (Rs mn) 2027E	EBITDA (Rs mn) 2025	EBITDA (Rs mn) 2026E	EBITDA (Rs mn) 2027E	Sales (Rs mn) 2025	Sales (Rs mn) 2026E	Sales (Rs mn) 2027E
Coforge	2,000	351	35.5	9.1	41.3	30.0	9,773	14,397	18,814	19,866	28,255	33,829	121,912	163,733	194,761
Cyient	1,180	111	10.8	(16.3)	5.1	16.8	6,157	6,470	7,554	11,433	10,664	12,783	73,604	73,427	82,335
HCL Technologies	1,500	2,718	4.9	9.9	(0.4)	10.4	172,704	172,158	190,204	254,929	261,179	282,005	1,170,567	1,261,704	1,367,409
Hexaware Technologies	910	611	18.7	17.5	20.2	17.2	11,764	14,188	16,627	18,302	20,824	25,833	119,744	132,838	150,960
Indegene	680	241	16.0	13.8	15.3	16.8	4,225	4,737	5,532	5,501	6,296	7,199	28,393	31,797	36,575
Infosys	1,850	4,151	7.5	8.2	5.3	9.7	262,930	276,940	303,886	392,350	409,559	445,468	1,629,900	1,720,320	1,871,006
KPIT Technologies	1,000	274	12.7	29.5	5.6	20.4	7,544	7,971	9,594	12,298	13,238	15,341	58,423	62,853	72,926
L&T Technology Services	4,100	106	10.0	(2.9)	3.4	17.0	12,668	13,091	15,323	18,925	19,765	22,535	106,702	119,148	133,159
LTIMindtree	5,100	296	12.3	0.4	9.3	15.4	46,020	50,284	58,007	64,949	68,125	77,682	380,081	405,998	447,064
Mphasis	2,600	189	11.0	9.1	10.3	11.8	16,991	18,737	20,940	26,467	28,503	31,093	142,299	154,615	171,416
Persistent Systems	4,800	156	19.8	22.2	22.8	16.9	14,001	17,263	20,181	20,582	24,982	29,294	119,387	141,572	165,669
Rategain	540	119	4.9	36.6	(9.5)	21.6	2,089	1,891	2,298	2,321	1,965	2,484	10,767	11,566	12,982
Tata Elxsi	4,000	62	1.8	(0.9)	(18.4)	26.9	8,047	6,403	8,126	9,927	6,403	10,684	37,290	37,096	42,560
Tata Technologies	450	406	9.0	(2.1)	5.4	12.7	6,770	7,136	8,046	9,341	8,935	9,853	51,685	51,569	56,154
TCS	3,800	3,619	6.3	3.5	4.7	8.0	485,530	508,242	549,122	674,070	687,271	756,064	2,553,240	2,581,947	2,803,253
Tech Mahindra	1,830	890	27.2	51.9	21.4	33.3	42,788	51,951	69,226	69,911	86,208	109,385	529,883	548,631	601,211
Wipro	240	10,347	3.0	22.2	1.4	4.6	130,225	131,523	139,602	180,851	177,891	186,981	890,917	904,649	958,044

Source: Kotak Institutional Equities estimates

"Each of the analysts named below hereby certifies that, with respect to each subject company and its securities for which the analyst is responsible in this report, (1) all of the views expressed in this report accurately reflect his or her personal views about the subject companies and securities, and (2) no part of his or her compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views expressed in this report: Kawaljeet Saluja, Sathishkumar S, Vamshi Krishna."

Ratings and other definitions/identifiers

Definitions of ratings

BUY. We expect this stock to deliver more than 15% returns over the next 12 months.

ADD. We expect this stock to deliver 5-15% returns over the next 12 months.

REDUCE. We expect this stock to deliver -5-+5% returns over the next 12 months.

SELL. We expect this stock to deliver <-5% returns over the next 12 months.

Our Fair Value estimates are also on a 12-month horizon basis.

Our Ratings System does not take into account short-term volatility in stock prices related to movements in the market. Hence, a particular Rating may not strictly be in accordance with the Rating System at all times.

Distribution of ratings/investment banking relationships

Kotak Institutional Equities Research coverage universe

Percentage of companies covered by Kotak Institutional Equities, within the specified category.

Percentage of companies within each category for which Kotak Institutional Equities and or its affiliates has provided investment banking services within the previous 12 months.

* The above categories are defined as follows: Buy = We expect this stock to deliver more than 15% returns over the next 12 months; Add = We expect this stock to deliver 5-15% returns over the next 12 months; Reduce = We expect this stock to deliver -5-+5% returns over the next 12 months; Sell = We expect this stock to deliver less than -5% returns over the next 12 months. Our target prices are also on a 12-month horizon basis. These ratings are used illustratively to comply with applicable regulations. As of 30/06/2025 Kotak Institutional Equities Investment Research had investment ratings on 293 equity securities.

Source: Kotak Institutional Equities

As of June 30, 2025

Coverage view

The coverage view represents each analyst's overall fundamental outlook on the Sector. The coverage view will consist of one of the following designations: **Attractive, Neutral, Cautious.**

Other ratings/identifiers

NR = Not Rated. The investment rating and fair value, if any, have been suspended temporarily. Such suspension is in compliance with applicable regulation(s) and/or Kotak Securities policies in circumstances when Kotak Securities or its affiliates is acting in an advisory capacity in a merger or strategic transaction involving this company and in certain other circumstances.

CS = Coverage Suspended. Kotak Securities has suspended coverage of this company.

NC = Not Covered. Kotak Securities does not cover this company.

RS = Rating Suspended. Kotak Securities Research has suspended the investment rating and fair value, if any, for this stock, because there is not a sufficient fundamental basis for determining an investment rating or fair value. The previous investment rating and fair value, if any, are no longer in effect for this stock and should not be relied upon.

NA = Not Available or Not Applicable. The information is not available for display or is not applicable.

NM = Not Meaningful. The information is not meaningful and is therefore excluded.

Corporate Office

Kotak Securities Ltd.
27 BKC, Plot No. C-27, "G Block" Bandra Kurla
Complex, Bandra (E) Mumbai 400 051, India
Tel: +91-22-43360000

Overseas Affiliates

Kotak Mahindra (UK) Ltd
8th Floor, Portsoken House
155-157 Minories, London EC3N 1LS
Tel: +44-20-7977-6900

Kotak Mahindra Inc
PENN 1,1 Pennsylvania Plaza,
Suite 1720, New York, NY 10119, USA
Tel: +1-212-600-8858

Investments in securities market are subject to market risks. Read all the related documents carefully before investing.

Registration granted by SEBI and certification from NISM in no way guarantee performance of the intermediary or provide any assurance of returns to investors.