

FOUNDATION OF DIGITAL BUSINESS

Surojit Mookherjee

Vinod Gupta School of Management

Indian Institute of Technology Kharagpur

Week 07

Lecture 34

Lecture 34 : Orchestrating Agentic AI for Intelligent Business Operations

Good morning. Now, in the last session of this module, Tools and Techniques for AI, I will be covering Orchestrating Agentic AI for Business Operations. So, we have been talking about prompt engineering, large language models, etc. Now we will move to the next step, where we will discuss in more detail how to orchestrate agentic AI for intelligent business operations.

So now, I want the tool to perform tasks for me—not just provide information or make predictions or classifications, but actually complete the work, finish the task, and perform duties typically done by a human. Now, all such tasks will be handled by the AI tool, which is known as agentic AI. So, I will refer to a document an IBM study for this purpose to just to explain the things how the business organizations are looking the applications.

AI agents are transforming the playing field, as their widespread introduction into commercial establishments is reshaping industries. Organizations are turning to agentic AI—systems that can achieve specific goals autonomously—to gain a competitive edge. With agentic AI, operations autonomously learn, adapt, and optimize in real time. It is not just about being faster and more efficient; it is about intelligent agents that proactively anticipate challenges, personalize experiences, and drive innovation. It understands what you want, determines what needs to be done, and executes it.

The shift is from automating tasks to orchestrating ever-adapting processes. So, it is just not a plain simple automation. It is understand you just give it instruction that I need this in a plain language, simple language, no coding and then it totally finds out what needs to be done, but also executes that. Successful deployment agents in just about the tech it

depends on creating synergy between people and AI across virtually every operational transaction and communications. For the technology itself, I am trying to interact with a tool, with a computer and augment my performance, my productivity.

So, that many of the things I can offload to the tool, some of the tasks which I need not do or not my assistant need to do or my whatever my team members need to do. Human oversight and corrections, decision making and most critical. Creativity are more important than ever in this next generation of intelligent operations. So, creativity has taken a big exposure for identification tools. So, create something, create a video, create an image, create at this based on text prompts.

So, only give some text descriptions and then outcome is a video, a human like video. AI automation generally begins with rule based systems that handle repetitive tasks evolving to it is growing it is evolving progression to AI assistance that help with tasks through natural language interactions and finally to AI agents. So, you have automation then AI assistance then you have AI agents which are broader systems that carry out tasks and make decisions autonomously with little human input. So, it will carry out the task.

So, automation AI assistant and AI agents play different roles and here is what is setting them apart. And assistants they perform tasks as requested through prompts is an AI assistant, use a conversational AI interface and base insights on machine learning or foundation models. But if I talk about AI agents, they work autonomously to achieve specific goals and KPIs. They act independently deciding which tools to use and when. Use persistent memory and adapting learning to refine their approach.

So, I will talk more about this agent part, this act independently deciding which tools to use and when, I will give you some examples. Now if you take overall an agentic AI operating model, the agents and people work in tandem. So, the top is represented by the people that is the user and bottom here you have the agents and they are working in tandem. Using the inputs on the left in a continuous flow of interaction, reaction decisions, reaction decision and action to enhance the business operations. So, what you see on the left are basically the business operations, the enterprise applications, these are the systems which you already have in our enterprise organization.

So, enterprise ERP, CRM, other integrated planning etcetera various softwares we are using in an organization. Your ecosystem you have a partner system, all these your vendors, your customer suppliers you linked up through your again IT systems and softwares and other ecosystem. And then you have of course, external data which is there

in the internet for you. So, economic indices, the market trends, the geopolitical events, the supplier production etcetera. So, all of these things which are available in the web.

Now, the agents what they are doing is data integration, they have impact evaluation, they have predictive analytics, multi agent collaboration, risk analyzer, feedback loop and finally, what we want is an output. Like some of the examples here given mentioned here like in the finance, human resources, order to cash flow, automation procurement, sales support and customer service. I will touch upon some of these later in more details, but what it overall tells you this model this study from used IBM as a reference to this model is that so people and agents they are interacting, the inputs to the system are existing whatever corporate level your IT systems are there, the softwares you are using, various tools, various databases

in addition to your linked whatever links you have with your vendors and suppliers. And thirdly of course, the overall web, the internet which is the big source of other relevant pieces of information which can be searched through the latest developments happenings etcetera. So, combining all of these as inputs based on whatever prompts you are giving, the agent is supposed to deliver certain outputs which are based on different functions, you want different things, your predictive financial planning, you might want a visualization model to inform financial performance, etc. There are n number of use cases, business use cases mentioned on the right hand side.

Optimizing outcomes. So, it is less about rules and more about results. To understand how AI agents work think self-driving car an analogy self-driving car just as the vehicle ingests a destination and navigates through environmental interpretations and decision making via sensors and algorithms. So, what you have told the car I want to go there.

from Calcutta to somewhere or maybe within the city somewhere that is the only instruction I have given them. After that the vehicle has to take care or find out which road it can utilize a Google map which is the best way. So, the Google map will tell it traffic conditions then it has continuously monitored the traffic on the road it has to monitor the signals stop start whatever. So, all the actions you are just given one instruction I want to reach this place that is all. But then the agent here in the automobiles vehicle it is taking help of so many other tools, GPS, Google map,

the sensors to see the red lights, traffic signals, nearby cars, pedestrians, crossings whatever. Similarly, an AI agent will perceive their surroundings, set objectives. And

then dynamically modify actions to meet these goals. So, the keyword here is dynamically. So, it is not a static thing.

So, it has to, at that point in time, find out what is happening and then take a decision. If you just compare that with a self-driving car. Three steps to unlock the potential of agentic AI in a business approach. Number one is to center your operating model around outcomes, not the tasks. So, what is the overall outcome you want?

The business goal, so the objectives that can be broken up into multiple tasks. So, tasks are not your focus; the focus should be the outcome. The tasks can be an output of those outcomes. So, embrace a startup mentality—that is what I am saying. So, some experiments will fail, and learning from failure is an important part of the transformation. So, agentification is not really a very well-matured tool as of today.

So, you have to keep trying; sometimes you will not get the desired result, but that should not set you back. So, that is why I say you start with a startup mentality because 70-80 percent of startups actually fail. hardly 10, 15, 20 percent startups succeed. So, that does not mean nobody is going for new startups because the failure rate is high. Consider management of a digital labor a new profession.

So, your workforce talent is key to ensuring transparency and quality in AI driven operations. So, anticipate and redefine roles and skills to manage digital labor. Non-skilling programs equip staff with digital literacy, strategic skills and ethics guidelines. So, today's labor you have to reskill them to adopt or to adapt digital techniques. So, you have to equip them with digital literacy, strategic skills and ethics guidelines.

So, we have been talking about ethics guidelines and we will talk about responsible AI more in the following sessions. Establish a digital labor operating team to build guardrails and drive ethical outcomes. So, bring together developers, ethicists, policymakers and domain experts to cultivate enterprise wide collaboration. And this goto squad serves as an advisory voice to set the guidelines needed for safe exploration and development. that is what can lead to scale.

So, this is very important step for scaling up. So, you must bring a team of different expertise, build on the guardrails, build the ethical part of the whole thing. So, that you can use successfully such tools without much problem. Step 2, prepare to scale securely. Enterprise business operations are prime candidates for capitalizing on agentic AI because they rely on data driven processes and are focused on measurable outcomes.

So, exactly what we want we do in a business operations database and a measurable outcome. By embedding strong data management governance and security from the outside is non-negotiable for responsible and impactful agentic AI adoptions. So, mark these words. strong data management, strong governance and strong security right from the beginning. Do not wait that we will do security at the end or this at the end etcetera.

Invest in data management, evaluate your data lineage, quality, privacy the whole anything in AI starts with data. I have been repeatedly saying this over and over again and the quality of data you will decide what will be the quality of your output. So, ensure that your data can deliver high quality unbiased results. do not leave AI governance behind. This autonomous AI scales opportunity as well as risk, because any output you are delivering your customers are using public are using etcetera.

So, governance becomes very important, because anything wrong anything bias anything which goes against the customers or the public etcetera will create lot of noise in the system. So, be prepared for that. Establish governance right at the beginning, not only to prevent but also to mitigate. If something goes wrong, then what to do, what actions to take. So, have that already pre-planned; do not wait for things to go wrong and then start thinking—it is too late.

Manage your AI agents' identities. Include identity access management for AI agents, such as non-human authentication, authorization, etc. To provide dynamic access management, policy-based authorization, and governance. So, what agents are being used should also be identified and known to people. So, if something goes wrong, you know which agent was the problem.

So, you have to remove it, improve it, clean it, or fine-tune it, etcetera. Step 3 is speed time to value. Blend internal domain expertise with external specialized skills to rapidly deploy tailored agent TKI solutions that deliver a competitive advantage. What it means is that basically, we may not have adequate internal talent to scale up agentic AI. So, there you can think of interacting with external skills vendors.

Evaluate organizational capabilities, identify existing skills, knowledge, and experience related to AI, machine learning, automation, and process orchestration. Review your budget, infrastructure, and tools needed to support AI agent development and integrations. If you are thinking of doing things inside you should check what is the skills available, do an audit. Find out the gap, skill gap, knowledge gap and experience gap in related to of course AI machine learning automation and this process orchestration.

Then of course, you have to give your budget that if you can want to hire can you afford to hire more talents new talents because these new area skills these are now in high demand. So, obviously, they will be charging higher compensation. So, they will be expensive also building new infrastructure that you also Expensive or you can use the cloud for a more economic solution and other tools whatever you may need. So, that is of course, everything has to be covered through your budget.

Benchmark external options, compare the managed service providers and outsourcing partners based on their capabilities reputation and alignment with the business needs and of course, the cost angle. And then do pilot projects, start with that test in-house and external solutions to gauge effectiveness cost and potential return on investments. This hands-on approach will help executives make informed decisions on building and buying agentic AI expertise. And you must always establish clear performance metrics and KPIs for agents, otherwise how will you compare how it is performing and find out measure it and with other tools and techniques. And be sure to align them with the business objectives and future strategies.

So, they are very powerful tools they will be very useful in future and more and more refinement will happen it will be of the quality will improve. Obviously, more and more of it will come into the regular usage, but you must have a very managed approach in the sense you have to step with caution start always with a pilot Use cases have it in a small scale, do in-house testing, get it reviewed by external agencies, get it audited by external agencies if required. So, bring in strong governance for that before you finally put it in a commercial use or give it to throw it open for your customers or to the public. So, be very careful about how you proceed with the whole thing for the scaling up and commercialization of such tools.

Some operation specific views, agentic AI use cases in finance. Predictive financial modeling, intelligent reconciliation, automated churning or anomaly detection. The KPIs and impacts, forecast accuracy, day sales, day sales outstanding, APR cycle time, cost reduction etcetera. So, these are all the KPIs and impacts for these agents. So, this financial modeling agents can analyze historical data and build predictive models to forecast outcomes such as cash flow projections or budget variances.

So, this is the output which is giving it to you straight away. Forecast, anomaly detection, fraud detection etcetera. So, this allows finance professionals to focus on assessing inherent uncertainty in the forecast and developing risk mitigation strategies. This AI

power virtual assistant streamline the inquiries from employees on payment processing, expense tracking, compliance reporting. So, all of these queries can be handled through agents.

And the other major area is the fraud prevention practice. This is one finance area which every finance professional would be interested in how to prevent financial frauds. In HR, the many use cases include predictive workforce planning, intelligent talent acquisition, automated benefits and compensation administration, employee sentiment analysis, virtual assistance, and you can use these KPIs for measuring performance, productivity, retention rate, training effectiveness, etc. So, they can spark a radical change in HR, the agent KPIs. They can use historical data on employee turnover, promotions, and performance to forecast future workforce needs.

These virtual assistants can provide employees a single point of contact for HR queries, for example. Assistants are enhanced by agents that deliver personalized responses and execute transactions. Such as travel bookings and compensation analysis. So, this is where HR, you know, the major applications you can have a virtual assistant to which any of the employees can interact and get, for example, travel booking done. So, in many sales organizations or in the service industry where a large number of employees incur a lot of travel.

And you have travel agents doing all this job: travel booking, ticket booking, etcetera. So, all of that can be automated in-house using such agents, virtual assistants. In a process of, say, order to cash, you raise a purchase order on a vendor, and then the material gets delivered, and then you pay against the invoice after checking the material input, etcetera. So, all of these can also be identified. So, automated order management, inventory optimization, invoice management, etcetera.

So, it can extract order data, check compliance with terms including credit risk, and initiate fulfillment instructions. So, there could be more things like a smart card system, which is now being introduced in supply chain management. So, if you have an agreement with the vendor and based on the smart card and this agent they can check whatever order has been released, check whether it has been met, satisfied, complied with in terms of quality, quantity, date, delivery etcetera and then instruct the payment. So, minimal or probably no human interaction is required to complete the whole cycle order to cash.

The agentic AI can also transform customer experience, it can respond immediately no matter the day or night to customers who order questions such as product pricing or delivery windows and it can predict customers needs resolving the issues before they ever occur. So, they can detect what is the likely problem and then can proactively suggest solutions. So, before the problem has occurred it can be resolved. Source to pay, so dynamic sourcing, contract negotiation, automated source to pay again related to your supply from vendors and supply chain issue here, where you can use agents for your purpose.

The agents can identify which suppliers based on the quality, price, location, capacity and reputation. So, they can anticipate disruptions such as weather or geopolitical. and proactively recommend alternative sourcing. So, this is one use case it is very interesting use case. So, since it is made to track geopolitical situation or the weather condition, it can predict that the delivery of this good from such and such location to such and

such location can get impacted because of say weather for example, likelihood there is going to be a cyclone. and it can proactively recommend alternative sourcing. So, it can suggest instead of that supplier from this location maybe change the supplier from another location to avoid this temporary problem. So, they are armed with insights on market conditions, supplier capabilities and organizational needs, they can even draft and negotiate contracts. So, the agent can draft a proper contract.

which of course, you will can edit and refine it, but the basic contract can be drafted. Well, AI agents analyze procurement professionals can focus on cultivating supplier relationship, driving innovation in sourcing strategies and navigating complex. This is since it is being so productive doing all these tasks, so it releasing capacity of the regular people who were earlier doing this work, so what they can do next is do something a higher level activities like innovation or developing new suppliers for example, supplier extension and developing new sourcing strategies making more efficient sourcing strategies etcetera. So, that is where you need human participation.

In sales and support system, sales forecasting, lead prioritization, virtual assistance, sentiment analysis, competitive intelligence. So, these are some of these areas we already discussed where agents can also make to play a big role. So, AI enabled automation in customer service is an untapped gold mine for pinpointing sales opportunities and tailoring customer interactions. So, you can from doing your sentiment analysis or following through your sentiment analysis from the feedbacks or going through social

media feeds given by your prospective customers or existing prospective customers, the agents AI agents can be able to will be able to pinpoint sales opportunity.

So, more than half of the customer service executives report minimal automation in customer communication and little use of self-service assistance for field service and customer service. This is today's situation, but this is going to change because many say they do leverage partial automation in customer feedback and customer support inquiries and which helps in customer retention and customer onboarding through all this use of tools for your automation. So, it can give you meaningful sales forecasting and lead generation by analyzing the various material which is available on social media platform for example. Organization-specific views, sales support and customer service with agentic AI organizations can drive breakthrough customer service and sales support.

The agents can automate lead scoring, sharpen forecasting, personalize customer engagement and curate sales content all while handling routine interactions. So, you can personalize your customer engagement every customer feels to be they feel when they like that when they given some importance. And they realize the importance when you specify that this is being developed or made just for you, for your liking—this is what you like. So, we are making it for you, etcetera. Analyzing historical sales data and market trends, they can zero in on promising customer leads, helping the sales team focus on high-value opportunities, etc.

Agentic AI can deliver 24/7 multilingual global support, as all of us know, with proactive customized responses. And it can even decipher customer sentiment. So, customer service representatives can fine-tune their in-person engagement. So, if ahead you know what is the sentiment of this particular customer whether he is happy or not happy let us say unhappy customer then with that knowledge will help you to start your conversation and the way

have the discussion because you already know that he is unhappy and what exactly his sentiment is. To accelerate transformation across customer service top areas they are likely to outsource, self-service digital assistance for field service, customer support inquiries, customer product and service training, and customer communications. So, you can see there is tremendous potential in self-support and customer service for agentic AI. Putting agentic AI into practice to understand how this works, think of a self-driving car.

Just as the vehicle ingests a destination and navigates through environment interpretation and decision via sensors and algorithms, AI agents perceive the surroundings, set

objectives, and dynamically modify actions to meet these goals. I have already talked about this—just wanted to bring it up in the concluding session because this comparison, if you think about it, will actually help you understand how the thing works. If you leave a car, it has to drive through, and nobody is driving the car; it is being driven by the software. Assume that is the agentic AI system—of course, it has a lot of AI built in.

This is getting all sorts of inputs, and then it has to make a decision—turn left, turn right, brake, stop, accelerate, etcetera. Various actions has to be taken when a car is going through traffic. So, same thing with agentic AI, you have given it a destination, this is what I want to get output, now do whatever you want to do, then you leave it to its own judgment, own intelligence. So, it can connect with whatever systems your company's data, history data, your transaction your customer data, your vendor data, supplier data plus what is the scenario situation global situation scenario and

like I said the environment, the weather condition etcetera or geopolitical situation what is happening between China, US etcetera. And then combining all of these things and extracting meaningful information from each of these systems it is reviewing and studying and then finally give you a decision. So, you want a decision because you want an action. So, agent TKI is supposed to do a task, just not information will not just tell you that this is where it might go wrong. So, you might get delayed etcetera.

No, difference here is that agent TKI will tell you that because of this reason, you change from this vendor transfer the order it will change the order purchase order from this vendor to another vendor who is located in a different place. So, all of that will be done by the AI system which is here is a agentic AI. So, that is why we are calling it as agentic AI. So, it not only finds out the problem, it finds out the solution and executes the solution.

like the car autonomous car it will finally, take you you are sitting in the back side seat it will take you to the destination. So, that is the whole purpose. So, agent API similarly the whole purpose is it will execute the task just not tell you and give you an information. The frameworks and guardrails coupled with strong digital identity management control the safe correct execution of core enterprise functions. So, the digital identity management is very important.

If you take the case of the car autonomous car, this is often very common question which is discussed that if a car makes an accident, who is responsible? There is no driver, is it the owner? The owner we say I depend on the software, I did not design the software,

who designed the software? So, is it the software programmer or is it the car manufacturer who chose that software and put it in the car or is it the internet service provider? because the sensor data response time was took much longer than what it was supposed to take a decision

in a fraction of a second, but because of latency it took 2 seconds for the response to come from the cloud. So, it got delayed action got delayed and then the accident happened etcetera. Similarly, with agent TKI you have to identify which are the agencies doing what. With agents operating on the clock and the globe automated operational functions never take a break. Touchless workflows and processes are also gaining traction by 2027 it is estimated 85 percent of executives forecast that their AI driven automation effort

will significantly facilitate touchless operations with varied applications across different functional areas. What it is saying here with touchless is you do not need the human intervention that human once you give it an instruction the job will be done. Finally, the job will be executed like that example I was giving you over the purchase order, if it finds out that this vendor will not be able to supply in the time because of whatever reason, I will change my vendor because I have an alternate vendor in a different location which is not impacted by this change of situation or condition and will be able to meet my requirement.

So, nobody changes or corrects the purchase order and sends it to the new vendors everything is done by the agent TKM. So, that is why it is saying it is touchless. So, the human being does not have to touch the keyboard or type something, etcetera. So, with that, I will end this session and also the module on tools and techniques in AI. I have covered roughly and there are n number of tools huge number of tools etcetera, but I have touched upon some of the important ones the common ones and

can, of course, always go on searching and learning. And, of course, new tools will keep evolving; that is the most important part of AI technology—this exponential growth. Every day, every month, etcetera, new things keep coming. Keep yourself updated with whatever new things are happening. Thank you very much.