Augmented AI

Take advantage of the powerful combination of human and machine intelligence

KPMG

Myth: You've built and deployed your AI solution; now you can sit back, let it run on its own, and reap the benefits

Reality: Successful AI deployment to continuously gain ROI requires ongoing human engagement

There's no doubt that AI is a game changing technology that can drive new capabilities and incredible benefits to all types of businesses. However, when deploying these technologies companies need to be aware that they can't rely solely on the "artificial" part.

That's Augmented AI: the powerful combination of human and machine intelligence.

It's true. You can build the greatest model in the world for your AI solution, but if you have not built in periodic human engagement for validation or handling change management, or have not incorporated model re-training, you will slowly lose the value of what you developed. Over time, all AI models experience drift and degradation of performance. Often, a model is already stale and out of date by the time you first deploy it. The fix is to build human monitoring and remediation processes into your solution so that you can maintain value and improve results proactively. The details of the human aspect will vary depending on the solution, use case, and time sensitivity of the application. For example, monitoring chatbot conversations, extraction of information from documents, or daily forecasting applications will have differing timescales and need for frequency of human validation. Regulatory applications will likely have significantly less tolerance for data drift and prediction error and require more frequent attention from humans. But no matter the model, it will require validation and human interaction at regular intervals to retain benefits.

How can the human touch best add value to AI solutions?

Monitor Input

The team should assess whether the incoming data looks different from the data used to train the model.

In an ideal world, your model would be trained on a dataset with enough variety that it will continue to perform well and generalize to new data in the future. However, if the distribution of real-world data shifts too far from the composition of training data it could have a negative impact on the effectiveness of the model and accuracy of results, depending on the use case and its tolerance for data change and drift. Measuring shifts, determining when they become a problem, and defining how to correct the input requires two skill sets. The subject matter expert identifies what shifts or changes in data are problematic. The data scientist determines when an adjustment to the model is needed to maintain high AI performance. Their collective intelligence augments the business value of AI.

Monitor Output

It's critical to regularly re-assess whether results are correct and whether the solution is doing what it's designed to do.

For example, are the AI's predictions accurately predicting subscription renewals? Are customer service inquiries being handled effectively by a bot? Humans can observe the real-world results of the AI solution to validate outcomes and pinpoint areas requiring re-training. Building this step into your monitoring process from the start allows you to find and correct potential problems before they become an issue for your users. For some applications — predictive AI for forecasting for example — measurement of accuracy can be automated. In the vast majority of cases, however, human intervention is required to identify, address and redress performance deficits. Once again, enter the subject matter experts and the data scientists.

Humans can observe the real-world results of the AI solution to validate outcomes and pinpoint areas requiring re-training.

Intelligent re-training

Another AI myth: new data can just be fed back in to automatically re-train the model. It's usually not that simple.

Data used for re-training needs to be evaluated and curated to ensure the best results. Simply putting unscrubbed data into service can lead to over-fitting or bias in the model, and lead to unintended outcomes. Usually a correction to the model will require introducing new data. This in turn requires subject matter expert review to ensure that the data is labeled correctly. Once the subject matter expert has identified these new potential data points, you will want the engagement of a data scientist to look at the data and choose which to use to re-train the model. A high-performing AI solution requires a good amount of variety in datasets used for training so that the solution is better able to handle new pieces of data that come in. The data scientist will shape the datasets in a way that will ensure the model delivers the solution most effectively and most efficiently for the business goal. Together, the subject matter expert and the data scientist add their intelligence to that of the AI solution to keep it at peak performance. That's augmented AI in action.

You can't just feed new data back in to re-train the model.

KPMG can help

If you are crafting a new AI solution or improving the performance of an existing application, we can help. Working in concert with your in-house teams, we can assist in determining business need, defining roles for humans and machines, selecting, creating and deploying a solution, and developing governance processes to monitor and re-train models over time — all appropriate to the specific use case in question. We have subject matter experts, business process specialists, AI application developers and data scientists to help you implement a high-performing augmented AI solution — one that can yield ongoing value over time. Call us.

Contact us

Rachel Wagner-Kaiser
Director, Data Scientist,
KPMG Lighthouse
+1 206 913 4498
rwagner-kaiser@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

kpmg.com/socialmedia