ikigai
AI Council

Developing and implementing time series forecasting functionalities responsibly

Quarterly Report – Q2 2024

Council Notes

Brief update since last quarter:

Last quarter's AI Council meeting centered on AI 'explainability' with a focus on prediction use cases. You can find the report from that meeting on the website. Since then, Ikigai has begun implementing recommendations from the AI Council in its platform. Case-based reasoning is now a default feature of time series forecasting in Ikigai: platform users can understand predictions based on both temporally- and spatially-related situations. For example, sales forecasts for SKU A can be explained based on sales for SKU A one year prior or by sales of a similar item, SKU B. The other two recommendations, calibrated confidence intervals and data provenance, are on Ikigai's roadmap.

Contact Ikigai to learn more.

Welcome to Dr. Christia:

The Council welcomes Dr. Fotini Christia as a new member. Dr. Christia is the Ford International Professor of the Social Sciences in the Department of Political Science. Her research has focused on issues of conflict and cooperation in the Muslim world, and more recently on societal impacts of advanced computing technologies. She will be the Director of MIT's Institute for Data, Systems, and Society starting July 1, 2024.

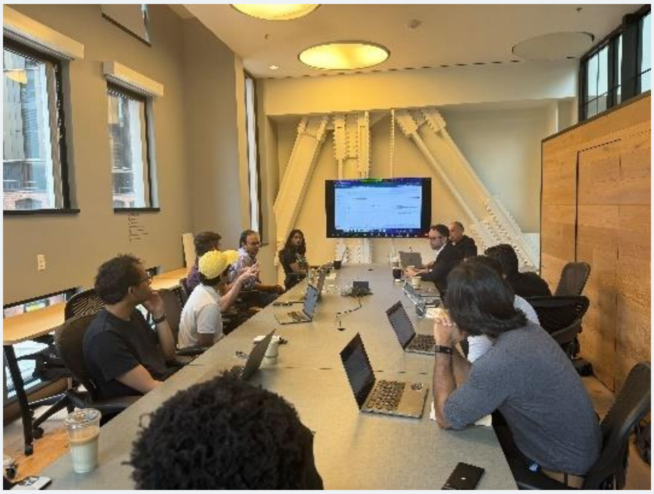

This quarter, the AI Council members gathered in-person in Cambridge, MA, where they were joined by members of the Ikigai team.

Developing and implementing time series forecasting functionalities responsibly

As businesses increasingly adopt AI for decision making, both AI companies and their enterprise customers must consider the implications of their actions on the world at large. The mission of AI Council is to identify and provide guidelines for addressing these societal implications surrounding AI development and usage.

In the previous AI Council meeting in Q1 2024, academic and industry experts discussed the issue of AI 'explainability': how can AI results be explained to ensure businesses and individuals can understand them, which is critical to test fairness? For example, if someone is denied a loan due to an AI produced recommendation, the bank in question should be able to explain the result or correct it. You can find the report from that meeting on the AI Council website.

The Q2 2024 AI Council topic was time series forecasting: how can enterprises use time series forecasting functionalities responsibly? Time series forecasting has implications across industries and use cases: businesses can use it for use cases from predicting sales or cash outflows to optimizing stock or shift schedules. While time series forecasting is not a new concept, advanced AI techniques are being developed and leveraged that can enable more accurate and actionable predictions. Used effectively, these time series forecasting techniques can lead to greater equity (e.g., fewer biased, better predictions for banking customer risk) and greater sustainability (e.g., fewer SKUs manufactured given lower expected demand); however, used ineffectively, time series forecasting can provide poor results and lead to worse outcomes for both enterprises and their customers. During the Q2 2024 quarterly meeting, AI Council members shared their thoughts on guidelines to use time series forecasting responsibly.

Time series forecasting functionalities and their importance

Time series forecasting is the process of leveraging tabular datasets, including at least one that is timestamped, to predict future outcomes. More than a single number, forecasts are about capturing overall uncertainty by utilizing all structural properties of it. Effective time series forecasting involves multiple types of time series data and several functionalities.

Every enterprise workflow involves forecasting and planning with complex data under uncertainty:

- For Product & Sales, how do they forecast demand and plan supply for Sales and Operations Planning (S&OP)?
- For Talent, how do they forecast work and plan for skills for workforce planning & optimization?
- For Customer Relationship Management, how do they forecast customer interest and plan for sales resources?
- For Finance, how do they forecast cash burn and plan for growth for Financial Planning and Allocation (FP&A)?

The data for each of these enterprise workflows is tabular and timestamped. In fact, enterprise data can be thought of as a collection of spreadsheets, each cell containing either a numeric value or a short text and associated with a timestamp. Forecasting on this tabular, timestamped data is essential to addressing the balancing act at the core of all enterprise workflows. Businesses have traditionally used deterministic models on disparate data to make educated guesses, but the real world is stochastic/probabilistic.

Ikigai and other time series AI solutions can help companies effectively capture uncertainty and achieve these enterprise balancing acts. As their capabilities are used, though, they must consider the societal and business impacts of their technology.

For more information on time series forecasting functionalities, see [this whitepaper.]

Key Insights from the Session

All time series forecasting workflows start with data. There are two broad types: time series data, which has variables that have different numerical values over time, and spatial data, which includes values that may or may not vary over time. There are five sub-types total:

- **[Time series data] Outcome data:** data that captures how one or more values of interest vary over time (e.g., product sales over time; capital expenditures over time)
- **[Time series data] Auxiliary data:** time series data that may impact the outcome variable of interest (e.g., holidays in the case of product sales; market credit rate in the case of capital expenditures)
- **[Time series data] Intervention data:** time series data that the workflow user has control over and may impact the outcome time series data (e.g., promotions in the case of product sales; pricing in the case of capital expenditures)

- **[Spatial data] Meta information:** additional information about time series data or variables (e.g., product meta-information like color in the case of product sales; types of revenue in the case of capital expenditures)
- **[Spatial data] Domain properties:** relationships between outcome time series data (e.g., geography of stores for product sales; collection information for capital expenditures)

An effective time series forecasting tool enables the user to collect and leverage all these types of data to best represent the real world and draw conclusions. The outcome data, meta information, and domain properties are natural choices for the user to incorporate when forecasting. If the user wants to forecast product sales, then product sales over time, the product catalog, and details about the distribution network are the datasets to bring in, respectively. By looking at temporal and spatial data, enterprises can use AI to understand their customers. For example, as Dr. Dahleh pointed out, banks can use AI to "provide 360 view that helps from mitigating risk to identifying cross- and up-selling opportunities." However, the user has more discretion when considering auxiliary data or intervention data. As AI Council Member and GW Law scholar Aram Gavoor pointed out, "you can do [time series forecasting] with all kinds of data, like real estate or population density" – but which information helps, and which information confuses the problem?

AI Council Member and Penn Computer Science Professor Michael Kearns noted the importance of guardrails around auxiliary and intervention: is there "explicit guidance or controls for customers around [auxiliary data or intervention data]? Or guidance internally?" When scientists experiment with data, they try to be rigorous with what data they use and why; if customers are "doing experiments themselves, it's very hard for [companies] to control" data use, Kearns said. Time series forecasting tools – and AI companies broadly – should consider their **data strategy** when enabling customers to incorporate exogenous data.

A well-defined data strategy can prevent customers from obtaining spurious correlations or drawing misleading conclusions. Dr. Kearns pointed out that "spurious correlations with exogenous data sources" seems "morally related to … p-hacking," a practice in research where data can be manipulated or unintentionally selected to obtain statistically significant results even if there is no real pattern. In the case of spurious correlations in time series forecasting, certain exogenous data choices can lead to assumptions of causation where none exists.

A common example of **causation vs. correlation** is the strong positive relationship between chocolate consumption and Nobel prizes won per capita[1]: eating more chocolate was interpreted to lead to more Nobel prizes. It's easy to dismiss the link between chocolate and Nobel prizes as a spurious correlation, but business users need a methodology to determine causation in their own workflows. As AI Council Member and UC Berkeley Computer Science Professor Michael I. Jordan pointed out, a business user wants to know: "if I implemented a promotion, did it really affect my sales?" Dr. Jordan advocates for rigorous statistical testing: if feasible, "[companies have] to do a randomized experiment" to know whether an intervention like a sales promotion has an impact or if changes to sales are unrelated. Randomized Controlled Trials, or RCTs, are rigorous methodologies for testing whether a treatment or intervention works by randomly allocating participants to two or more groups, where at least one group does not receive the intervention. This design isolates the effect of the treatment by minimizing selection bias and other confounding factors. In a demand forecasting context, for example, an RCT could be used to measure the effectiveness of a certain promotion. There are also a range of methods for causal inference from observational data that can be viewed as approximations to RCTs. These are widely employed in econometrics and it's important to make them broadly available in a general time series context.

An effective forecasting tool, AI Council Member and MIT Professor Munther Dahleh noted, would enable the user to model potential outcomes: "What if [a user has] to run a promotion, but [they] don't know what sort of promotion [they] need to do... could the user simulate an RCT?"

Experimentation via simulation can enable customers that can't run RCTs to estimate causal effects of interventions like promotions. Tools like Ikigai's Time2Vec™ are the first step towards enabling this type of experimentation: Time2Vec™ enables users to find time series with similar structural characteristics.[2] For example, if SKU A and SKU B have similar sales patterns, their Time2Vec™ embedding would place them close to each other in space. The user could use the fact that SKUs A and B are similar to experiment and draw conclusions based on their questions: e.g., if promotion 1 works well on SKU A, then promotion 1 could work well on SKU B? The user should be able to then simulate sales for SKU B with promotion 1. Then, the user can leverage this type of experimentation to make informed operational decisions.

By implementing a well-defined data strategy, methodologies to determine causation vs. correlation, and experimentation enabled by simulation, time series AI companies can help their customers forecast responsibly and effectively.

1. See https://www.nejm.org/doi/abs/10.1056/NEJMon1211064
2. Time2VecTM embeds time series in lower dimensional space in the same way Word2Vec vectorizes words.

The AI Council concluded with the importance of further deliberation on data strategy and the role of experimentation in enterprises, which will form the basis of the next quarterly discussion. ■

See next page for Recommended Actions for AI companies to implement time series forecasting functionalities responsibly and effectively.

Recommended Actions

Three ways AI companies can implement time series functionality effectively and responsibly:

1. Define **a data strategy**

Define and implement guidelines to follow when incorporating auxiliary or intervention data in forecasting workflows to prevent spurious correlations or misleading conclusions.

2. Test correlation vs. causation

Implement guidelines or features to enable rigorous statistical testing and understand underlying mechanisms to drive effects of interest (e.g., via Randomized Controlled Trials or other approximations).

3. Experiment via **simulation**

Enable users to experiment and test causal effects by simulating historical and future results.

AI Council Members

Devavrat Shah

Viterbi Professor of AI and Data Science and MIT; Ikigai CEO and co-founder

Fotini Christia

Ford International Professor of the Social Sciences in the Department of Political Science at MIT; Director of the Institute for Data, Systems and Society at MIT

Munther Dahleh

William Coolidge Professor in Electrical Engineering and Computer Science, Founding Director of the MIT Institute for Data, Systems, and Society

Aram Gavoor

Associate Dean for Academic Affairs, GW Law School

Michael Kearns

Professor and National Center Chair, Department of Computer and Information Science

Michael I. Jordan

Professor in the Department of Electrical Engineering, Computer Science, and Statistics

Kamal Ahluwalia

President of Ikigai; ex-President of Eightfold AI, ex-CRO Apttus

Visit our website to learn more.

https://www.ikigailabs.io/

Report by Parvathi Narayan; please reach out to parvathi@ikigailabs.io for questions.

ikigai