Beyond the Hype: Artificial Intelligence in Retail

An Aptos White Paper

aptos

Engaging Customers Differently.

Table of Contents

Artificial Intelligence Hype vs. Reality 3

Defining “AI” 4

Not All AIs Are Alike 6

AI Challenges in Retail 8

- The Data 8
- The Algorithm 9
- The Platform 11

The Aptos Approach to AI 12

- Identifying Opportunities 12
- Prioritizing the Opportunities 12
- Solutions Approach 13
- Timeline for Aptos 15

An AI Journey 16

Artificial Intelligence Hype vs. Reality

Judging by the number of times the term "artificial intelligence" gets used in the press and by technology influencers, one could believe that AI is a mature capability being deployed more and more every day, well on its way towards taking over the way all software works as we know it.

The reality is quite different. Gartner reported that as recently as 2018, only 2% of all retailers had deployed AI and only another 24% were "experimenting" in the short term.[1] When retailers have embarked on AI initiatives, they have found that the algorithms take longer than expected to "learn" the business, the data they have is too sparse to be valuable, and the users who would be on the receiving end of AI-driven recommendations refuse to engage (possibly fearing their ultimate replacement). Success stories are few and very narrowly focused.

Artificial intelligence has immense potential. Early tests show that the capability can drive step-level improvements in business outcomes, when applied correctly and managed well. As a technology that has the potential to disrupt the retail industry, AI cannot be ignored.

At Aptos, we have watched the hype and the progress with equal interest. We do not engage in the pursuit of technology for technology's sake, but neither do we ignore new disruptions. Rather, we follow an innovation process for bringing new technologies into our solution suite. This process monitors developments in data generation (things like IoT, voice, user-generated content, etc.), computing power (the rise of cloud-native SaaS, quantum computing, etc.), as well as developments in new capabilities, like voice commerce, augmented reality, robotics, and artificial intelligence.

When technologies have reached a maturity level where both benefits and early pitfalls are visible, we embark on an evaluation process to determine how Aptos can best take advantage of a new capability, and how best to bring it to our clients.

This document discusses what we have learned about artificial intelligence through this process to date, as well as describes where we are in the process of adding AI to our solution portfolio.

Defining “AI”

Artificial Intelligence. Machine Learning. Deep Learning. Neural Networks. There is a lot of vocabulary associated with AI, and sometimes terms that do not mean the same thing are mistakenly used interchangeably.

To level set, Aptos focuses on the following definitions:

- **Artificial Intelligence (AI)** – Systems that automate decisions at scale, driven by quantitative modeling or classification of complex text, image, video and/or audio information.
- **Machine Learning (ML)** – An application of AI that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. Machine learning focuses on the development of computer programs that can access data and use it to learn for themselves.
- **Deep Learning (DL)** – One of many approaches to machine learning that can potentially accelerate a machine’s ability to learn, and thus require a smaller data set. Neural network machine learning is a specific type of deep learning that identifies and classifies patterns in rich datasets, such as voice, video, images and sounds.

Thus, when exploring AI, we would characterize Deep Learning as an approach to Machine Learning, which itself is one type of artificial intelligence.

One other aspect of AI that is important to distinguish is what actually goes into an AI capability. Within the retail industry, many vendors have described their capabilities as “AI applications”, but unless the solution provider is a genuine AI platform, on the level of Google DeepMind or IBM Watson, what most solution providers have are AI capabilities that are embedded within a traditional application, whether within a system of insight or a system of execution.

This is a critical distinction. AI capabilities, by their nature, are very “small.” Baidu, a Chinese technology company at the forefront of AI, recently released over 100 different AI capabilities that could be accessed through an API call. The list of services gets very specific very quickly: things like license plate identification, passport identification, foul language recognition, plant identification, food identification, and many more.

An assortment optimization AI might sound like a highly sophisticated AI solution, but instead it is a cover for one simple (small) AI capability – the ability to identify the class of garment and its color, for example, “yellow short-sleeved shirt” vs. “blue long-sleeved sweater.” That may be enough to give a boost to a choice count analysis, for example, but it is hardly a comprehensive, AI-driven approach to assortment. It is exactly this challenge of taking one small AI capability and using it to paint an entire solution as AI that has driven much of the hype that exists today.

As a recent example, in November of 2018, Yoox released 8 by Yoox, a high-fashion assortment that was credited as being "designed by AI".[2] When digging into how it actually worked, the solution fell far short of an AI actually generating fashion designs (though this is a capability currently available in a limited, highly-controlled environment). Yoox used an AI algorithm to create a "dynamic mood board". The AI was computer-vision oriented, and was pointed to sources of user-generated content. Yoox chose where to focus the AI's attention, on specific influencers and fashion content in specific cities.

The AI then used the images it aggregated to identify fashion trends – the rise of tan, double-breasted trench coats, as an illustrative example. As a trend was identified, it was added to the mood board with representative examples from the user-generated content. An additional predictive AI algorithm looked at the attributes identified for the trends, and made a prediction (a forecast) for how well the trend would resonate with typical Yoox shoppers. Finally, humans stepped in and decided which trends to pay attention to, and how to turn those trends into new product designs. The 8 by Yoox collection was designed using AI to assist certain steps in the process, but it was hardly "designed by AI."

Not All AIs Are Alike

Over the last few years, McKinsey, the consultancy, published several research articles on the topic of artificial intelligence, resulting in possibly one of the best resources on general AI and its business value across industries today. In one analysis, the company identified three types of AI:[3]

- **Classification** – analyzing various inputs to identify logical groupings of content. This could include text, audio, images, or video. Classification can be used to assign attributes to objects, and to group objects into likely segments. In the Yoox example, AI was used to identify defining characteristics of an image and group images into trends.
- **Prediction** – analyzing content in order to predict the next action. This could be examining buyer behavior to predict their next purchase step, predicting when and where the next unit of a SKU might be sold, or predicting an overall outcome, as pollsters increasingly use AI when making election predictions. Again, in the Yoox example, predictive AI was used to assess whether the trends identified would resonate with customers.
- **Generation** – identifying key characteristics of an object in order to create a new, unique representation of the object. For example, turning from identifying pictures of dogs (a classification exercise) to using what the machine learned to generate new pictures of dogs. In the Yoox example, the one thing AI did not do was generate actual designs.

While these are all excellent characterizations of what AI can potentially do, they do not recognize that different kinds of AI have different capabilities. An AI that is good at processing text is not good at processing images or video – the signal processing necessary to parse video is far more complex than that used to analyze text and has generally lagged behind text analysis capabilities as a result.

At Aptos, we focus on a different type of AI classification:

- **Natural Language Processing (NLP)** – the AI that is specifically targeted towards processing text. This can be applied directly to text, or assisted by an AI that is tuned to translate speech to text, before handing off to the NLP AI to assign meaning to the text identified.
- **Computer Vision** – recognizing that signal processing in video is a different capability altogether from understanding text, we regard computer vision as another important kind of AI. Again, computer vision can be combined with other types of AI. In the Yoox example, computer vision categorized user-generated images and then grouped them together according to perceived common attributes – tan, double breasted, and coat, for example.
- **Prediction** – the type of AI that takes information and tries to predict or project a next step or outcome, similar to how McKinsey defines it.

McKinsey was right to include prediction as a type of AI, whether you look at what an AI can do or what an AI is intrinsically good at. When it comes down to it, prediction is really the only thing that AI is for. Classification is just a prediction that the AI has identified an object properly. As well for generation: the plague of deep fake videos showing celebrities doing outrageous things is really just a machine's prediction of what kind of fake video will fool a human.

In Aptos' view, prediction has by far the most potential in the retail industry. It is the foundation for nearly every optimization capability that can be applied to retail opportunities.

AI Challenges in Retail

Before opportunities can be addressed, however, it is important to consider some of the current challenges in applying AI to the retail industry. There are three main areas of focus when looking at potential shortcomings of current AI approaches: the data, the algorithm, and the platform.

The Data

With regards to the data needed for AI to be effective, the most important question to ask is whether you will have access to a large enough data set to train the model. When data is limited or lacking, AI engineers can inadvertently fall into the problem of selection bias, where the data selected hides a pattern that humans do not see, which then leads the AI to an incorrect conclusion about the data, because the sample used was not truly representative.

There have been two very prominent cases of selection bias in AI over the last few years. More recently, Amazon debuted a hiring AI that was supposed to review resumes and flag certain resumes for a closer look by a human.[4] Unfortunately, as soon as it was put into service (in a provisional way), it almost immediately started discriminating against women. How did this happen? Amazon was using the AI to screen technical resumes, and there is already a selection bias inherent in the pool of candidates – far more men than women get technical degrees, and even fewer women with these degrees apparently apply to technology jobs at Amazon. The company estimated that only 20% of its resumes were from women.

In training the algorithm, the AI watched human HR representatives not select women's resumes, and concluded that being female was a negative trait. It started rejecting resumes that explicitly contained words indicating that the applicant was female – terms like "women's studies" in the activities portion of a resume, for example. Had Amazon tested the resume AI against all resumes out there seeking any kind of technical job, the company might have had a more representative data set of resumes to start from.

A no less disturbing example of selection bias came out of a Google facial recognition AI. The AI was supposed to distinguish people from animals, but suddenly started labeling dark-skinned people as gorillas.[5] How did this happen? The AI had been trained using only light-skinned people, so that when it encountered a wide variety of skin tones, it did not have the training to know that people came in such a variety. When AI ethicists talk about the challenge of a lack of diversity in AI engineers, it's exactly this kind of oversight they are referring to. The AI engineers at Google were not deliberately racist, they just didn't think about making sure they had a truly representative sample of people when training the AI to recognize them.

Selection bias can be managed, if the people selecting the data are aware of the dangers. However, there is still an open question about how much data is needed in order to protect against selection bias. When IBM tasked Watson to interpret radiology scans, one of the major failings of the initiative was lack of access to a large enough data set to build confidence in Watson's predictions.[6]

In retail, companies often feel awash in data and may discount the importance of having a broad enough data set to protect against selection bias. But that perspective has not been fully tested. Aptos views this challenge as an opportunity to play an aggregator role, to ensure that algorithms start from a fully representative data set before turning them over to start making business decisions. This role can be delivered both in terms of looking across a large customer base when training an algorithm and in terms of providing access to third party data sources that an individual retailer may not be able to afford.

The Algorithm

The next question to ask is, who is maintaining the model over time. Another way to frame this challenge is to look at the dangers of "black box" AI.

Most AIs do not provide a ready, easy-to-read map of the decisions the machine is making or the lessons it learns over time. When a machine comes to an erroneous conclusion, as in the cases of Google and Amazon described above, in order to find out where the machine went wrong, AI engineers have to go in and look at the algorithms and interpret them into something that makes sense to people. In Amazon's case, "Women are undesirable tech job candidates" probably read more like a line of math from a Calculus student's nightmare, rather than the sentence you just read.

This is one of the biggest challenges for AI today: most algorithms have not been designed to share the assumptions and lessons learned with people. Without such sharing, it becomes far more difficult to monitor the AI to make sure it does not learn something it should not.

Additionally, when a machine learns the wrong thing, "fixing" the algorithm apparently is very difficult. When Google found that its image identification had made a mistake, the solution was not to correct the wrong, but to make the algorithm "forget" an entire species of animal.[7]

This is not the first example where forgetting turned out to be preferable to fixing – when Microsoft released Tay into the Twittersphere (the AI chatbot was supposed to use Twitter to learn how to talk like a teenage girl), it took the worst trolls of Twitter less than 48 hours to teach Tay how to be a racist, fascist bigot.[8] For Microsoft, the fix was not to step her algorithm back to an earlier state. It was to make her forget and refuse to engage at all on the topics of religion or race.[9]

That's all well and good when you are running experiments to prove a technology. It is something else again when you're using it to run your business. In 2018, the British government began launching an investigation into retailers' online personalized pricing practices, to make sure that the personalization is not resulting in discriminatory practices against "vulnerable" populations.[10]

Whether they intend discrimination or not, retailers who do not have a way to monitor what their algorithms learn risk running afoul of an outcome that is effectively discrimination of a kind. Here's how easily that could happen: The machine driving online promotions learns that people coming to your eCommerce site from certain IP addresses or from certain search terms are long-term high value customers, and provides lots of promotions and perks early in the customer journey in order to capture those customers for the long term. Unfortunately, all of those people happen to live in a rich, white neighborhood, while meanwhile all of the people who are not offered discounts or perks live in disadvantaged neighborhoods with a high ethnic minority population.

When government regulators come looking, how can a retailer prove that they are not discriminating against poor people of color? In this case, intent (or lack thereof) is not going to matter. The effective outcome is discrimination. And if that discrimination is coming out of a black box AI, the retailer can easily be accused of not undertaking the due diligence required to monitor the algorithm and ensure that it does not come up with inappropriate – or illegal – conclusions.

Compound this with the current state of the technology: you can't "fix" it. Retailers would have to either start over and hope to train the algorithm in a way that prevents it from coming to the same discriminatory conclusion, or they would have to find a way to ensure that the machine is blinded to some of the customer attributes that have a high coincidence with wealth or ethnicity. In effect, they would have to cripple the model, just to make sure it doesn't do discriminatory things.

To make the problem more challenging, some algorithms are built in a way that makes it difficult to expose the assumptions or lessons learned – early versions of AI are very much designed this way. But it's a big risk. Not even the most advanced AI researchers know how, exactly, machines learn – certainly not enough to let them run off on their own. If experiments have demonstrated anything, it is that machines can easily learn "wrong" things because our own taboos and ethical prohibitions prevent us from considering those kinds of outcomes as in the realm of possibility.

When these realities meet up with employees in a retail enterprise who are supposed to give over their expertise to "a machine", the result is that even when there are improvements over typical end-user results, they are not attributed to the machine learning so much as just having a clearer view of the data to begin with – something that end users themselves could have benefitted from. In other words, AI may deliver an 80% accurate forecast, for example, but so can a human. That's not that much improvement on the business. And thus their resistance to using algorithms continues.

Aptos is looking closely at how AI companies are addressing the issues around algorithms. Some are taking the approach that people will "just get used to" the black box that is AI. Some are attempting to address the shortcomings of black box AI and do a better job of exposing what an AI learns and why. We feel an emphasis on the latter approach plays an extremely important role in driving adoption.

The Platform

In discussing AI types, we identified three: natural language processing, computer vision, and prediction. These are three different engines operating three different kinds of algorithms. An engine that specializes in NLP is not going to be good at computer vision or prediction. A forecasting engine won't recognize images, because it does not have computer vision. If a retailer wants to use all three capabilities, they will need three engines with three different types of training.

This leads to its own set of challenges and considerations, impacting how much of a platform you may take on when making an AI investment.

When you buy a packaged application that comes with AI included, you are typically buying a small data scientist group, with limited resources, packaging whatever engines they develop into a set workflow that you must use. They have identified the problem, and specified how it must be solved. Because of the limited resources available, there is typically little flexibility in the algorithm or how to use it, and it is highly dependent on the expertise available to that vendor only. This usually means strong industry expertise, with some more limited AI expertise, and it is more often focused on one or two micro AI capabilities (image identification or consumer sentiment analysis, for example), rather than a complete bundling of AI engines.

At the other end of the spectrum are AI platform players, like IBM, Google, and Amazon, but also independent platforms like Data Robot. These sometimes specialize only on one engine, like forecasting or computer vision. They are very deep in AI capabilities, but have limited knowledge about how to apply AI to specific industry problems. Platforms like these are more often implemented side-by-side to an execution process, so that it still requires a human to understand what the AI solution is recommending and to figure out how to take that recommendation into the executional system.

What is the best approach for a company looking to invest in AI? That is still an open question. Both have benefits and risks. Aptos' intent is to explore both kinds of approaches as we watch continued adoption in the industry

The Aptos Approach To AI

To get to the best strategy for AI for Aptos' customers, we followed these steps.

Identifying Opportunities

We started with a greenfield approach to AI, by not being constrained by existing capabilities. Basically, if we could start our solution suite over from the beginning, how would we want to embed AI into our solutions, and which use-cases would we want to focus on?

Across our solution areas of merchandise lifecycle, order lifecycle, and customer lifecycle, we identified over forty different opportunities where AI could be applied to processes. This reflected a mix of market opportunities as well as opportunities to further our own product capabilities.

We then grouped those use-cases into 5 themes:

- **Forecasting** – Statistical methodology to predict/project sales, with an array of inputs modeled, learning to maximize accuracy for the future.
- **Personalization** – Using a combination of characteristics, attributes, and behaviors, tailored to support and influence individual shopper journeys.
- **Price Optimization** – Suggesting the best price, before selling, during selling, and end of life.
- **Inventory Optimization** – Predicting inventory needs geographically by specific times, while identifying & utilizing multiple factors to prescribe inventory position.
- **Fulfillment Optimization** – Automatically analyzing available inventory/shipping providers and lowest cost method of fulfilling an order, against customer expectations for service.

Prioritizing the Opportunities

Aptos quickly realized that forecasting served an important foundational role: all of the other four themes depend on forecasting in some way, so it became easy to prioritize forecasting as the first focus area for our AI approach.

Forecasting with AI is not a "settled" field or capability. And it can be easy to focus so much on the platform or foundation that you end up creating a solution that cannot be leveraged across many different focus areas. So, to keep our focus broad, we looked for another area to evaluate alongside, to ensure that whatever forecasting decision we make, we make sure it is something that can be used as the basis of supporting all of the themes identified. Because there is a good tension between forecasting, which tends to be product-based, and personalization, which tends to be customer-focused, we quickly decided to add personalization to our priority list.

Solutions Approach

Aptos has identified six areas where AI impacts forecasting:

1. **Examining a nearly unlimited number of data inputs to identify impactful causal factors.** One advantage of AI is its ability to identify patterns out of very large data sets. This can be used in forecasting to bring both much more granular datasets (local weather or event data, for example) and to examine many more different kinds of data sets (unstructured data like social media activity) to be able to identify which data sets actually have the most impact on forecast accuracy.

2. **Applying advanced algorithms, like Neural Nets or evolutionary algorithms, to create new forecast methods.** One of the most exciting – but uncharted – opportunities for forecasting is to move beyond traditional time-series forecasts and into new territory. Neural Nets (NN's) decompose decision-making and spread it across multiple nodes, similar to how the human brain is structured and operates. Some NN's will create new decision nodes as they "learn", also approximating the operation of the human brain.

   Evolutionary algorithms approximate biological processes, much like Darwin's "evolution of the species". For forecasting, evolutionary algorithms create predictions by combining various possible factors into expected outcomes. When the algorithm compares the results to what was predicted, it "kills off" unsuccessful combinations of factors, leaving only the most successful combinations to start from when making a new prediction.

   While these new ways of forecasting are exciting, they are not well established in retail forecasting and require a lot of evaluation before they can be applied deeply in retail processes.

3. **Selecting the right forecast model to use for each specific situation.** Rather than inventing new ways of forecasting, one promising application of AI is in using traditional time-series methods of forecasting, but using AI to identify which models are the best to use in a specific situation – think, which SKUs in which locations, for example – and then weighting the application of those models in each circumstance, and adjusting which models are used and to what degree over time.

4. **Applying the forecast to every granular SKU, across every granular location where it will be sold or across every customer that will buy it (or both).** For a long time, the state of the art for retail forecasting was at an aggregated level. Some fashion retailers might forecast to an "item" level, but may not break that down into color forecasts, or size forecasts, or even more challenging, color/size combinations. Some may aggregate above the item level, and use some kind of business rule to assign actual forecast quantities to specific items with a line or class, for example.

   The same is true when predicting demand at specific locations. Many retailers cluster stores into groups, and apply forecasting at the group level, and then simply divide by the number of stores to determine how to distribute a total quantity across the stores in the group. This practice is only as good as the clustering method used to group stores, however. AI can help identify better groups of stores to use, and can also help identify how to best forecast the quantities needed at each store. In fact, some AI-driven forecasting methods always forecast at the store/SKU level (even when looking at long horizons of time), and then aggregate that forecast up to whatever level is needed by the end user.

This approach is not limited to items – it can also be used to predict consumer demand without looking at locations, and then predicting where each customer may be most likely to buy the item – in a store, online, etc. This avoids treating online as a “rounding error” or markup to total store demand.

5. **Identifying when specific causal factors deteriorate in their contribution to the forecast, and replacing them with new, more relevant causal factors.** Models can change over time, but one of the challenges for retailers is keeping up with those changes. If a retailer is depending on social sentiment to create a product forecast (this would be a very sophisticated and leading-edge retailer), say they were focused primarily on Facebook as the social channel providing this input. But if their customers suddenly moved away from Facebook and onto TikTok, the value and relevancy of Facebook input into the forecast would decline. An AI can identify this shift, and at the least can begin more heavily discounting Facebook inputs, or, depending on the data the AI has access to, can begin also more heavily weighting models in favor of using TikTok, recognizing that this data source is now more relevant.

6. **Reacting to changes in demand assumptions more quickly than a human could.** Finally, an AI can recognize that the actual results are beginning to deviate significantly from the plan, and trigger activities to respond to this shift. This is one area where, without AI, retailers have really struggled. When “the plan” is made, it is almost never revisited in light of changing assumptions – retailers tend to respond in reaction mode in execution, rather than triggering a reset of the plan, even when the time horizon is relatively short.

There are other places where AI can contribute, for example through image identification and classification that could be used to create a richer set of product attributes, or natural language processing that allows end users to interact with the data in a more “Google-like” experience, where natural language questions are asked, and the AI uses the data at its disposal in a natural language response. However, these move the focus into other areas of AI beyond prediction and are outside of our immediate focus.

As we evaluate use-cases for AI in retail forecasting, we look to identify where the AI capability that enables the use-case falls within this spectrum of six different ways that AI can be applied to forecasting. Some opportunities span multiple areas identified above, while some focus on one or two particular capabilities. What is most important to our approach is to ensure that we are leveraging AI in forecasting in a way that is not "ground-breaking" or "experimental". While forecasting has the opportunity to drive an enormous amount of value in the retail enterprise, getting it wrong also promises an enormous amount of harm. A conservative approach here should not be perceived as lagging, but rather ensuring the application of AI for direct business value with minimal risk.

Timeline for Aptos

Aptos is currently in the process of evaluating potential technology partners as well as specific use-cases. We are exploring a wide variety of use-cases at this time, driven in part by partner capabilities in the market, as well as our own analysis of opportunities. We are also engaging with clients who are interested in partnering on the process, and who can help guide our roadmap.

An AI Journey

While the hype about AI still exceeds the reality, the gap is starting to narrow in important ways. Aptos sees a significant opportunity in leapfrogging current AI efforts, by focusing on prediction as the most valuable type of AI in retail, and by looking at both product-oriented prediction (forecasting demand) and customer-oriented prediction (personalization).

Our objectives for our AI journey are to address some of the shortcomings in the current approaches to AI, including the black box approach and potential for selection bias, while also combining our own industry and process expertise with the AI expertise of a partner.

We anticipate moving quickly in 2019 and 2020, and are actively seeking retailer partners to help us shape this process. Even if your own company is not ready to join us on this journey, we look forward to sharing with you what we learn along the way.

---

1 https://www.gartner.com/doc/3873631/cool-vendors-ai-retail

2 https://fashionista.com/2018/11/yoox-private-label-collection

3 https://www.mckinsey.com/~/media/McKinsey/Business%20Functions/McKinsey%20Analytics/Our%20Insights/The%20age%20of%20analytics%20Competing%20in%20a%20data%20driven%20world/MGI-The-Age-of-Analytics-Full-report.ashx

4 https://slate.com/business/2018/10/amazon-artificial-intelligence-hiring-discrimination-women.html

5 https://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/

6 https://www.massdevice.com/report-ibm-watson-delivered-unsafe-and-inaccurate-cancer-recommendations/

7 https://www.wired.com/story/when-it-comes-to-gorillas-google-photos-remains-blind/

8 https://techcrunch.com/2016/03/24/microsoft-silences-its-new-a-i-bot-tay-after-twitter-users-teach-it-racism/

9 https://qz.com/1340990/microsofts-politically-correct-chat-bot-is-even-worse-than-its-racist-one/

10 https://techcrunch.com/2018/11/05/uk-to-probe-fairness-of-personalized-pricing-practices-in-online-retail/

About Aptos

Aptos: Engaging Customers Differently

Aptos is the largest provider of enterprise software focused exclusively on retail. Our cloud-based Singular Retail™ solutions are trusted by over 1,000 retail brands in over 65 countries. With industry-leading omnichannel commerce and merchandise lifecycle management solutions, we help retailers develop dynamic and responsive assortments, streamline operations and deliver integrated, seamless experiences…wherever shoppers choose to engage. More than 1,300 colleagues share our collective passion for engaging customers differently, and we are committed to developing relationships built on trust and tangible value by partnering with our clients to create agile retail enterprises that are built to thrive in an era of constant change.

Learn more at **info@aptos.com** and **www.aptos.com**.

aptos®

Engaging Customers Differently.