AI Augmented: a family of modern decision-making approaches

Thoughtworks' AI Augmented approach includes a family of novel decision-making approaches that go beyond the common computer-centric paradigm to solve organizations' most difficult problems.

Barton Friedland & Jarno Kartela

/thoughtworks

Table of contents

Where we left off 3

Decision-making under uncertainty 4

The AI Augmented family of approaches 5

The role of learning 6

Automation Augmented 7

Generalizing the Action 8

Marimekko Decision Factory 9

Operations Augmented 11

Case Study: Kittilä Airport 12

Creativity Augmented 14

Case Study: Mondelēz 15

Strategy Augmented 17

Case Study: Supply Chain Sustainability 18

Welcome to the family 20

About the authors 21

About Thoughtworks 22

Where we left off

In our foundational article about AI Augmented, we outlined an approach which goes beyond the common computer-centric implementations we see customers undertaking to solve new kinds of problems, and explained how the AI Augmented approach radically improves decision-making. We also discussed how common misunderstandings around using AI techniques to create customer value can limit an organization's ability to solve a wider range of problems using these tools.

In this longer paper, we will outline more specifically how the AI Augmented approach straddles a broad spectrum of decision-making activities, ranging from tactical and operational decisions, to creative and strategic decisions. Crucially, we argue that an AI Augmented approach produces better outcomes because it enables stakeholders to more effectively align their decision-making with a specific situation. We believe significant untapped value for our clients exists in the solution space which AI Augmented describes. To support this, we've also included client case studies for each to share the extraordinary impacts this approach has achieved for our clients.

Decision-making under uncertainty

We often hear stakeholders assert that increasing levels of complexity and uncertainty confound their decision making. If your organization applies AI using only a computer-centric approach, this will certainly be the case. However, as many strategy[1], leadership[2], planning[3] and complexity science[4] experts have noted, the methodology used for decision-making must support stakeholders to better orient to their environment; otherwise, less than optimal results should be expected.

From an AI Augmented perspective, we see AI as running across two key axes. These are depicted in Figure 1:

- Tactical and operational problems to creative and strategic problems;
- Low uncertainty to high uncertainty.

AI Augmented approaches are particularly useful in problem spaces such as the top right quadrant where greater uncertainty and the need for a strategic or creative approach is required. In such situations, there is often little data available. This is one feature that makes AI Augmented fundamentally different to dominant approaches to AI which rely exclusively on large datasets.

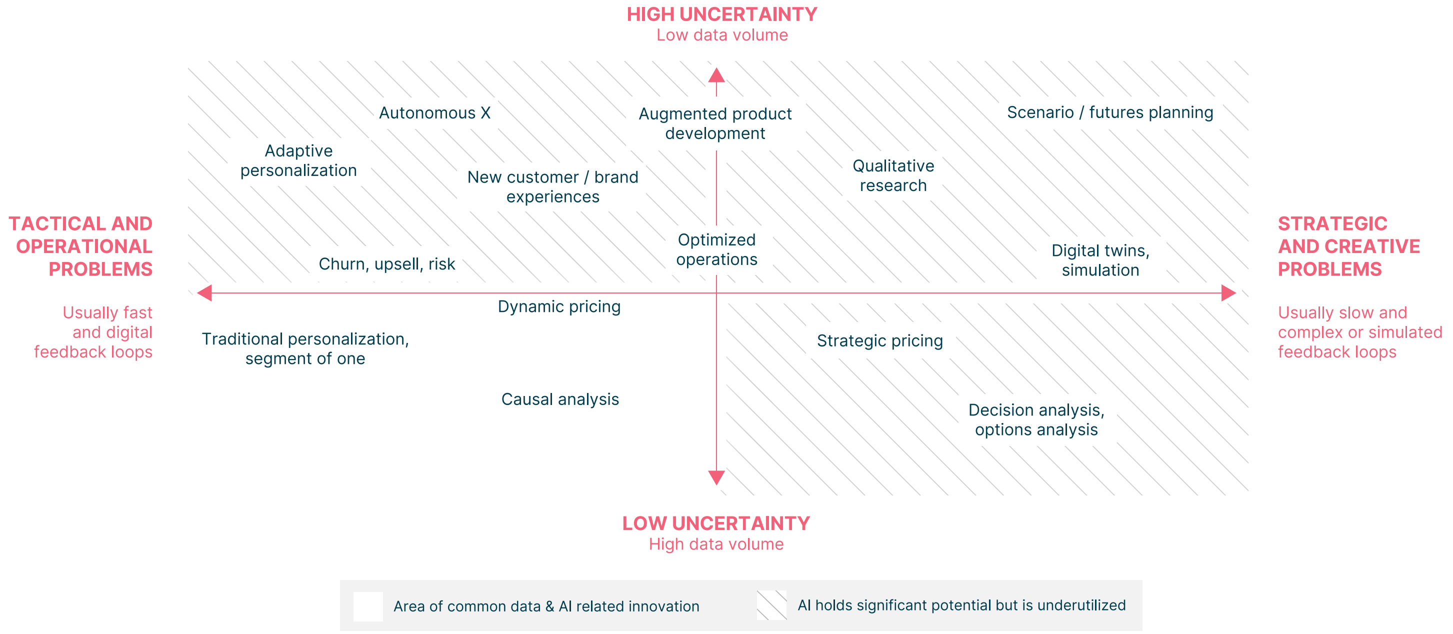

Figure 1 - AI solutions mapped across two axes - first, from tactical to strategic problems and second, from low- to high-uncertainty

[1] Courtney H, Kirkland J and Viguerie SP (1997) Strategy Under Uncertainty. Harvard Business Review 75(6): 67–79.

[2] Grint K (2022) Critical Essay: Wicked problems in the Age of Uncertainty. Human Relations: 001872672110707. DOI: 10.1177/00187267211070770.

[3] Rittel HWJ and Webber MM (1973) Dilemmas in a General Theory of Planning. Policy sciences 4(2): 155–169. DOI: 10.1007/bf01405730.

[4] Snowden DJ and Boone ME (2007) A Leader's Framework for Decision Making. Harvard Business Review 85(11): 68–76.

The AI Augmented family of approaches

This section describes the four distinctive approaches of AI Augmented that enable more effective decision-making. Figure 2 below shows how these sit along a spectrum from, on the one end, tactical and operational decisions, which tend to be highly automated. On the other end, we find more strategic and creative decisions which tend to involve increased human interaction. However, what binds all of these together is an augmented approach that emphasizes human learning at every level. We discuss each of these in more detail below, and describe how this gives rise to the four distinctive approaches within the AI Augmented family, listed below in figure 2.

Better decisions

Four AI Augmented approaches to intelligently experiment and enhance decisions

1. Automation Augmented

A scalable system that interacts with it's environment to optimize / personalize virtually any decision.

2. Operations Augmented

A flexible optimization toolkit bringing mathematically optimal results to resource constrained situations.

3. Creativity Augmented

An interactive R&D process between AI and experts to enable speed and quality in product development.

4. Strategy Augmented

An approach to codifying & simulating scenarios, building robust and transparent strategies.

TACTICAL AND OPERATIONAL DECISIONS
Usually fast and digital feedback loops

STRATEGIC AND CREATIVE DECISIONS
Usually slow and complex or simulated feedback loops

Figure 2 - The four AI Augmented approaches

The role of learning

All the members of the AI Augmented family follow the same underlying model; **human learning** is always positioned as a desired outcome of the system. Human learning is the central reason we call the approach **augmentation** – it's built on the idea that computers can help people think more effectively.[5]

We view effective thinking as a function of human learning and a basis for continuous improvement.[6] However, systems that are designed to help people learn must incorporate a different design strategy than systems that are designed to predict answers.[7] To effectively support humans in learning, a system must allow us to test a hypothesis and measure its outcome.[8] Finally, if such AI systems are designed to additionally capture human learning, these amplify the overall learning and enable both learning in humans, but also enhancing effective machine learning.

[5] Bush V (1945) As We May Think. Atlantic Monthly (July): 47–61

[6] Poppendieck M (2002) Principles of Lean Thinking. In: 2002, Proceedings of 17th Annual ACM Conference on Object-Oriented Programming, Systems, Languages and Applications.

[7] Wenger E (2000) Communities of Practice and Social Learning Systems. 7(2): 225–246. DOI: 10.1177/135050840072002.

[8] BHighsmith JA, Luu L and Robinson D (2019) EDGE: Value-Driven Digital Transformation. Pearson Education.

1. Automation Augmented

Our first approach deals with the ways we can use automation to augment human thinking and decision-making capability. While personalization technologies or next-best-action solutions are not new to most industries, these are generally designed to work within the space (next-best-action) and not depart from it. The automated agents that are created do as they are told, and they are generally built only to do the thing that is asked of them. Under these conditions, it would be questionable to argue that they augment human thinking capability. Automation Augmented turns this idea on its head.

In our foundational article on AI augmented, we argued that AI is built on problematic assumptions about what AI can and cannot do. Our practice shows us that it is possible to design a 21st century human-centric AI-based system using an augmented approach that can outperform machine-centric approaches with far less computing resources and less data.

For example, in our client practice, we've learned that for automation to bring any substantial value beyond automating processes only, we should design for systems that:

1. Are designed to work under uncertainty, as past data only does not tell us optimal answers nor can predict customer behavior;
2. Are designed to take full advantage of available data assets but can work even without any past data if required;
3. Are designed to directly optimize against a business goal by interacting in its environment without any added supervision, adapting its behavior independently;
4. Are designed to create new understanding of customers, process, context and actions taken, producing new, strategically relevant data assets instead of only consuming data for incremental improvement;
5. Are designed to limit the need for point solutions.

Generalizing the action

Generally speaking, a large majority of automation problems, like robotic process automation, personalization and dynamic pricing, can be generalized to i) a set of actions or choices (the action set available), ii) an environment where the actions are taken into use, iii) options about the environment and customers, iv) a real-time feedback structure and signaling - how can we infer what constitutes a good action and what is a bad one, and v) an insight environment that builds new datasets which can be used for strategically important decisions to better understand the dynamics of our actions against their respective environment. We have sometimes referred to these five elements we list above together as a **decision-factory** approach, where 'factory' highlights the scalability and reusability of such an approach.

The decison-factory approach embedded within Automation Augmented enables scalability for the organization by opening up a possibility to **organize multiple problems through a generalized pattern and solve that pattern effectively with a factory architecture.** This is further explained through our work with Marimekko, where this approach achieved a 41% lift in front page clickthrough and a 24% increase in average revenue per user.

Case Study:
Marimekko Decision Factory

Marimekko invests heavily in growth through digital channels, which account for a large portion of the company's revenue (total revenue at 152 million, over 20% YoY growth). The company has invested in technology supporting customer centricity and data-driven decisioning for some time now — they have had product recommendations and a proprietary digital analytics platform since 2018, and it includes marketing mix modeling since 2019. They also added new computer vision-backed immersive services in 2021. In 2022, Marimekko moved to an Automation Augmented approach to fuel better customer experiences, increase average sales per user and drive additional growth.

At Marimekko, personalization and recommendations can be roughly split into two categories; i) recommending from a (potentially very) large action set, like products (SKUs), and ii) recommending from a finite action set, like frontpage banners, artwork items, discount amounts, banner artwork, product categories, and so on. Within the digital experience automation and personalization space, it is common to have problems where we need to pick the best action from a finite action set. Thus, solving with point solutions is infeasible, and for Marimekko to provide the best possible customer experience, they needed to build both.

The goals for Automation Augmented project therefore go beyond obvious core business metrics to creating new, future-proof capabilities for the organization, including:

1. Improving the customer experience for Marimekko customers across contexts;
2. Increasing overall sales and average revenue per user (A/B tested);
3. Designing new mechanisms to understand customers and customer behavior through digital platforms; and
4. Designing new next-best-action and automated decisioning capabilities for future omnichannel personalization, engagement, optimization and insight across channels, markets and contexts.

In order to do so, our approach was to build a decision factory abstraction for making this a reality at scale, visualized in Figure 3 below.

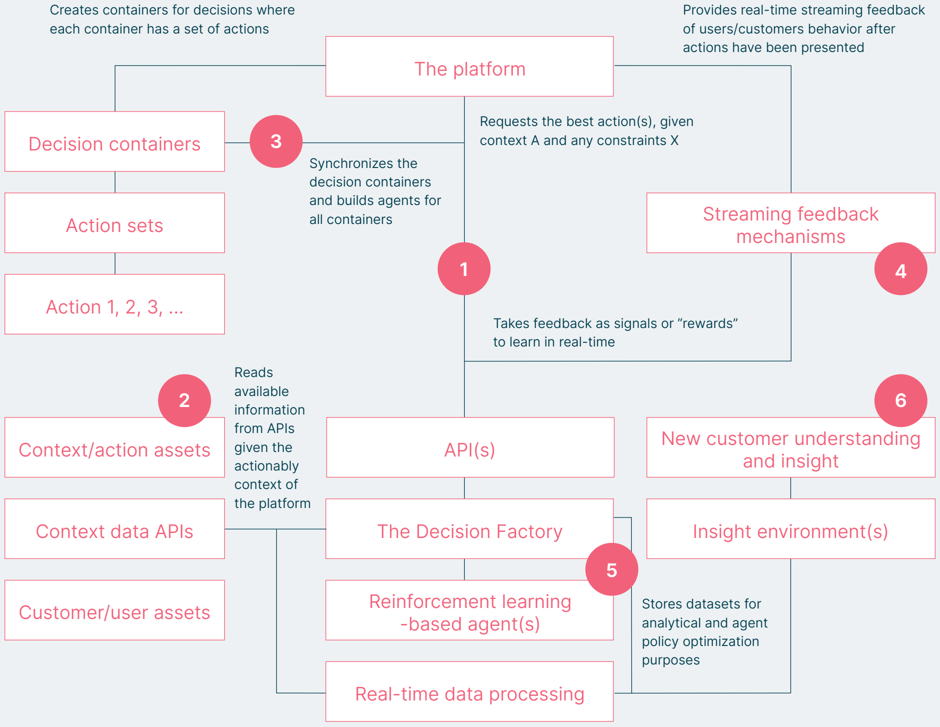

Figure 3: Decision Factory

Technically speaking, we are abstracting decisions to decision containers which hold a set of possible actions. Therefore, we can treat any problem as a problem of picking the best action, (1) given a customer and a context and (2) so that it maximizes our business outcome. Once in place, we don't need to use additional expert help to scale it for additional use-cases. We can just use normal e-commerce processes (3) to start optimizing and personalization virtually anything with choices.

In order to learn something, the decision factory ingests a real-time event stream (4) from marimekko.com, which triggers a learning process every 10 minutes to train every model inside every container that has new events from which to learn (5). Learning happens against a context — behavioral and contextual segment-level understanding about customers (which include above/below average interest in their new sub-brands like Kioski, overall fashion, home, accessories and above/below behavioral interest in sustainability, new collections,new items, and discounts).

Results

After being live for only a few hours, the decision factory pulled in a whopping **41% lift** in front page clickthrough — without pre-training the model. Further, after five weeks in use, a **24% increase** in average revenue per user had been achieved. On top of these, rich new insights were created on customer behavior, along with the instantiation of a scalable system to optimize and personalize virtually any tactical decision with a finite set of choices.

2. Operations Augmented

Connected, end-to-end streamlined operations are both strategically important and a source of competitive advantage.[10] We see startups that are formed around operational optimization, the world's biggest companies building in-house research centers for it and a general recognition that progress on the topic can yield massive increases in operational performance.

Operations research is a foundational management research topic in itself dealing often with complex systems analysis and optimization. We share the view on the fundamental importance of the field, but Operations Augmented adds specific emphasis on **how to solve complex optimization under extreme complexity and under high uncertainty**. At the same time, the approach augments our human capabilities as professionals, enabling learning and increased understanding.

Generally speaking, optimization is about making the best possible decision in a given context, under some restrictions and rules that the decision has to follow. In an enterprise context this usually means thousands of parameters and constraints with an infinite number of possible decisions but only a very select few of clearly optimal ones. In order to solve such complex problems, we need to ask:

1. What parameters are we actually optimizing for?
2. How do we observe if a decision was good or not?
3. What restrictions do we have for a possible solution?
4. How can we find the optimal solutions?
5. What are the implications of these solutions, and how are we surfacing them?

The last point is key to enabling an Operations Augmented solution. Without surfacing the plan to stakeholders in real time, they cannot act on it locally and autonomously. When applied in this fashion, an operations augmented augmented approach brings local, autonomous, and optimised decision-making to life.

Moreover, managing complex operations, supply chains and mission critical elements of a business involves a multitude of stakeholders. This includes both domain experts and management. From our perspective, it would be a fundamental oversight to try to solve problems in this space solely with technology only without also improving people and process. Furthermore, we claim that for operations to be truly augmented and significantly improved, that augmentation needs to be surfaced across different decision-making levels of the organization. Therefore, in almost all complex systems we are not optimizing for automation only, but rather we are optimizing to understand (what is the best solution and why), to improve (how are going to do this differently) and to automate (when applicable, can we directly use the outputs of such a system to increase efficiency). In so doing, Operations Augmented provides a way to forecast/simulate/plan, and therefore brings the benefits closer to strategic decision- making as well.

One could argue that answering questions 1-5 is solely a mathematical and theoretical exercise and while that may be true in textbook problems, it is never the case in a real-world business environment. For real-world applications, an end-to-end, human-centric approach to the problem and the design of the solution have significantly greater emphasis. We will explain this through a case study example augmenting the operational processes of running an airport.

[10] Beer S (1966) Decision and Control: The Meaning of Operational Research and Management Cybernetics. New York, NY: John WIley & Sons.

[11] Friedland B and Yamauchi Y (2011) Reflexive design thinking: putting more human in human-centered practices. ACM Interactions 18(2): 66–71. DOI: 10.1145/1925820.1925835.

Case Study: Kittilä Airport

At Finland's northernmost airport, Kittilä, the number of passengers grew by 12% in 2018. This growth put heavy pressure on infrastructure and resources for the airport. Even though Kittilä is a small airport with twelve airplane parking spots, 58 flights arrive and depart on the busiest days of the Christmas season. With this volume, the number of possible airplane parking alternatives at the airport is practically impossible for a human to calculate and optimize.

58 flights can be parked 10^31 different ways — a tough nut to crack even for a supercomputer. With 70-80% of those flights landing within a four-hour time frame, calling peak times busy would be an understatement. Parking spots aren't the only limited resource. The airport was constantly struggling to provide enough buses, staff and check-in counters to cater to the increase in air traffic. Simply adding more resources is ineffective, since outside the busiest season the airport only receives one or two flights a day. So. how do you flexibly operate an airport with limited resources?

In cooperation with the parking experts and the management of Kittilä airport, we created an optimization model that solves just that. Previously, daily parking plans at Kittilä were created manually the night before, using pen and paper.

As Kittilä's challenges were related to schedule changes and lack of resources, the criteria behind the perfect parking plan were defined as:

- The plan's ability to withstand schedule changes (robustness)
- Bus resources required to execute the plan

In creating an operations augmented plan for Kittilä Airport, our model allows the user to adjust the criteria freely. With these criteria set, the optimization model then uses flight data to build a mathematically-aligned parking plan based on all the data available. Further, the plan takes into consideration all rules (e.g. not all airplane types are allowed at all parking spots) and preferences (e.g. Non-Schengen flights parked near passport control) regarding the parking. The solution further employs machine learning to predict arrival times and passenger numbers, which are then incorporated into the optimization process. It is used daily in building the parking plan and in optimizing the bus routes and resources. It's also used to update the plan if delays occur and enables long-term resource planning (which was previously unheard of).

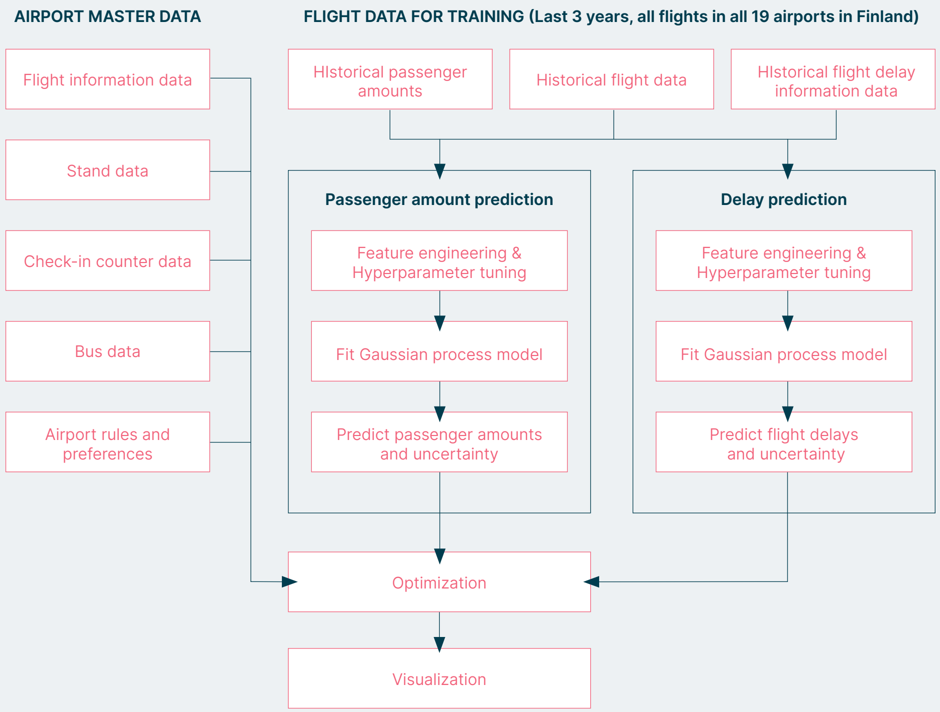

Figure 4: Operations Augmented, airport optimization example

Results

Planning time from 3 hours to 30 seconds — a 99.7% reduction — along with over 60% overall improvement in performance. The solution was deployed in the peak season of December, and comparing to the previous December:

- The share of airport-related flight delays decreased by 61%, even though the number of flights increased by 12% from previous year
- Duration of average airport-related flight delay decreased by 66 %
- The decrease in delays resulted in an estimated monthly 500 000 € cost savings
- The airport's Net Promoter Score (NPS) score increased by 20 points

Moreover, the Operations Augmented approach proved to be much more resilient to schedule changes. The new parking plan created by the model is optimal in the sense that idle time between successive flights at the same stand is as long as possible while minimizing resources needed and maximizing timeliness. Therefore, when delays or exceptions inevitably occur, it is less likely that the delays will impact other flights. As the plan is more robust, delays do not snowball to other flights.

It also created a lot more transparency and alignment across stakeholder groups. Using the Operations Augmented system, it's now possible to instantly get an overview of the current status of key resources across the entire airport. Parking spots, planes and buses are now shown on a single screen, which makes it easy for humans to understand what is happening at the airport. And a single click creates a new plan, if problems appear on the horizon.

Finally, it significantly improved collaboration and augmented the overall process of running the entire airport for the people involved. As the parking plan is now easily presentable and familiar to everyone at the airport, collaboration and coordination between different parties at the airport is simplified and the optimal plans can be distributed across the value chain.

3. Creativity Augmented

What if you are trying to produce a new version of an existing product for the market, but are not sure if the changes you are making add any value? In generative modeling with deep learning, it is hard to ask for a certain output: the algorithm will produce an endless amount of prototypes that can all be new and very inspirational. But what if the creative task you are working on is constrained by materials, resources? What if legislation or regulatory requirements affect the outcome, and many of the prototypes your generative algorithm provides are actually not usable in the real world or do not fit into the market? These are some difficult questions we have answered in recent years, while working on projects that can be labeled under a Creativity Augmented approach.

In our work, we have identified the following challenges as the major hurdles in a wide variety of creative and R&D fields:

Main challenges:

- Hard to experiment with lots of factors
- Lack of data
- Expensive to experiment (in terms of both time and/or money)
- High level of training and expertise is needed
    - Need to minimize extra training
    - Account for uneven performance

These challenges are significant for a person who is trying to use an algorithm to produce value, and then on top of these, we need also to inject creativity into the loop.

Creativity is something that intuitively does not feel like something a machine could handle, but with certain restrictions and constraints (which are essential to specify) we can produce useful and creative results with the help of algorithms. Since creative processes rely heavily on experimentation, this is one aspect we capitalize on in Creativity Augmented.

By their nature, processes that require experimentation are exploratory and therefore, from a data science perspective, one can not simply leverage past data to predict future outcomes.

For this reason, Creativity Augmented solutions provide a creative scaffolding for domain experts that help them optimize the application of that domain knowledge (e.g. minimizing the overall number of steps to a solution) to solve the creative problems where they have specialized expertise with machine learning support. We can best explain this through a case study.

Case study: Mondelēz

Since 2019, we have worked in close collaboration with Mondelēz International, helping their product developers create snacks that consumers love. Our journey together started with a set of proof-of-concepts and has since expanded to be a Creativity Augmented driven product development platform for several Mondelēz brands, including Oreo cookies.

Coming up with a recipe that is tasty but also checks a vast amount of requirements is a critical step for getting the products to market and to this end, we aim to support the product developer to make better decisions. The challenges product developers are solving vary – they range from creating consistent samples to creating new winning flavor combinations, ingredient availability, and how to substitute a flavor after shortage in supply or a regulatory change.

With the Creativity Augmented product development platform, product developers first set up a project and specify the constraints, metrics and ingredients that are affecting the product. This process is shown below in Figure 5.

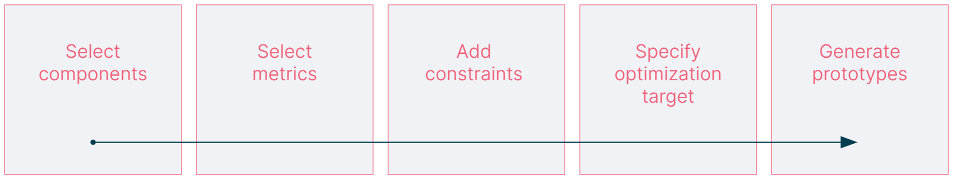

Figure 5: General process flow for setting up a product development project

After the setup of the product specifications, the machine learning model starts generating recipes that the product developers can work with in an iterative fashion, tasting samples and evaluating the generated recipes. After the feedback is provided back to the platform, it is used as training data for the next set of generated recipes. The recipes can be optimized for numerous criteria such as cost or overall liking.

Interestingly, before Creativity Augmented, the data gathered from this part of the workflow was unstructured — there was no systematic connection between the experiments and the test results that could be used for learning or knowledge distribution. This meant that the most critical decisions of what goes forward were made with no structured way of storing the data and no way of collectively learning from past experiments and using them for the benefit of new experiments.

We focused on closing this gap, and built a solution to help product developers in recipe generation by introducing an approach where AI creates the new recipes to experiment against the constraints and goals that are defined by the product developers themselves. When a product developer provides feedback of the experiment back to the model – learning from each one of the experiments, getting better sample by sample – not only the particular product developer obtains more knowledge, but the collective knowledge of all product developers using the Creativity Augmented approach is increased.

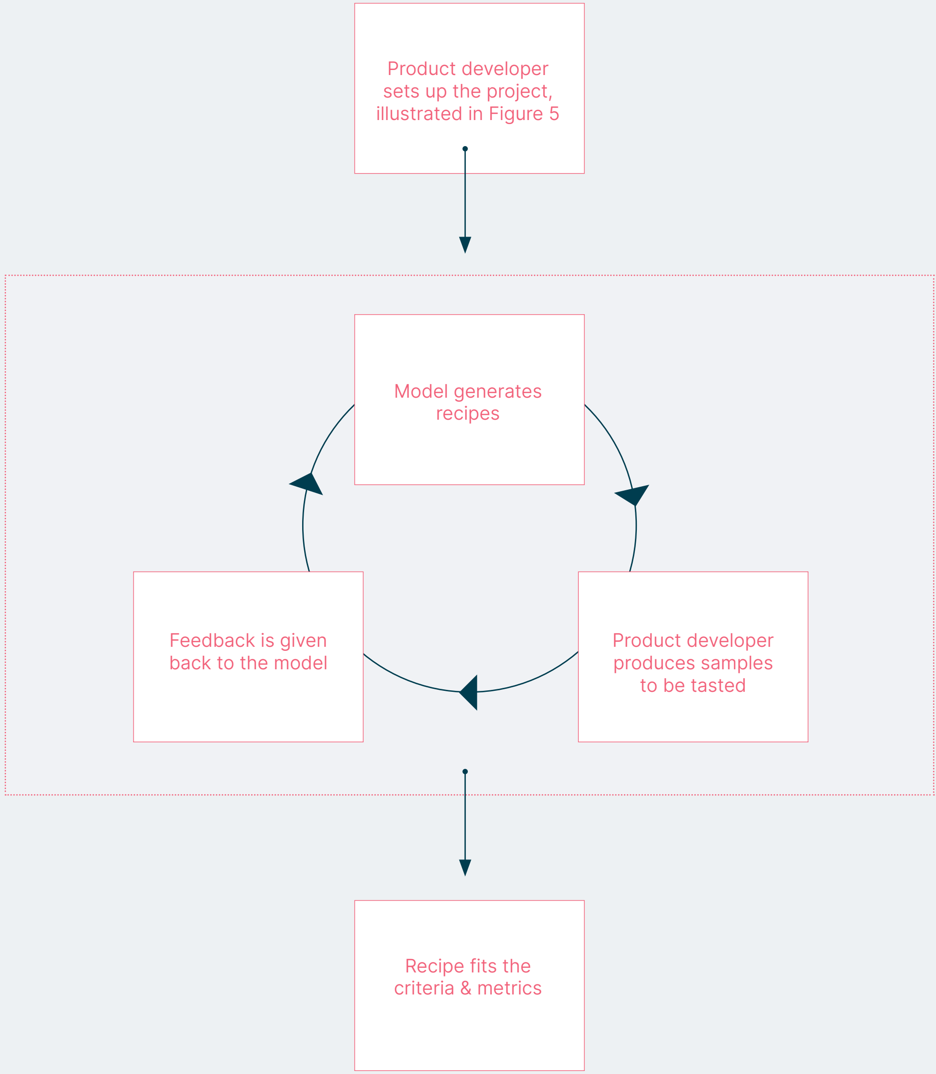

Figure 6: Creativity Augmented process for product development

Overall, Creativity Augmented is not just about creating AI models, but about leveraging an augmented learning capability to develop products that have been unachievable while at the same time helping people get results they have not been able to previously obtain.

Components of the solution:

- Faster, more efficient experimentation and speed to market
- Quick iterations
- Greater understanding for interdependency of factors
- Extent the scope and number of factors
- Abstract our expertise away

4. Strategy Augmented

And finally, we have Strategy Augmented, where we apply AI Augmented principles to human work of strategy development and decision-making.

Even without Strategy Augmented, our clients are better equipped for strategy work than ever before. Data allows them to better understand their customers, business and industry as a whole. On the other hand, the decision-making environment is becoming increasingly complex. Cycles are faster, the number of stakeholders is larger and ecosystems are more complex. Often our client's tools and coping mechanisms cannot adapt quickly enough — which calls for new and innovative approaches.

From our perspective, this can be achieved by enabling rigorous exploration of the possibilities for their business model, strategic differentiation, operational efficiency and performance.

To allow this, we combine a set of distinctive technological and human methodologies. These support a critical range of activities, including the robust development of scenarios, simulations, and causality analysis combined with a practical evidence-based, value-driven approach we call EDGE which tightly couples strategy development, execution, and measurement.

Computational models, when applied through Strategy Augmented, make it possible to iteratively optimise strategy work. Out of dozens of future scenarios, an optimization model can dynamically weigh each option in terms of impact, pointing towards the most favorable scenarios. This, in turn, enables the human decision-makers to have more meaningful discussions about the possible futures that lay ahead of them.

Case Study: Supply Chain Sustainability

Enterprises are under greater pressure than ever before to significantly reduce their carbon emissions and move to more sustainable business models. The demand for change is escalating rapidly — and it comes from citizens, governments, investors, consumers, and companies themselves. Inaction around sustainability is often caused by an assumption: that making changes will negatively affect the bottom line.

An international manufacturing and services company we worked with believes it doesn't have to be that way and started looking for bold opportunities through identifying actionable insights and future scenarios. Figure 7 below shows the sequence of work we undertook to support them in augmenting their strategy development and decision-making around their sustainability strategy.

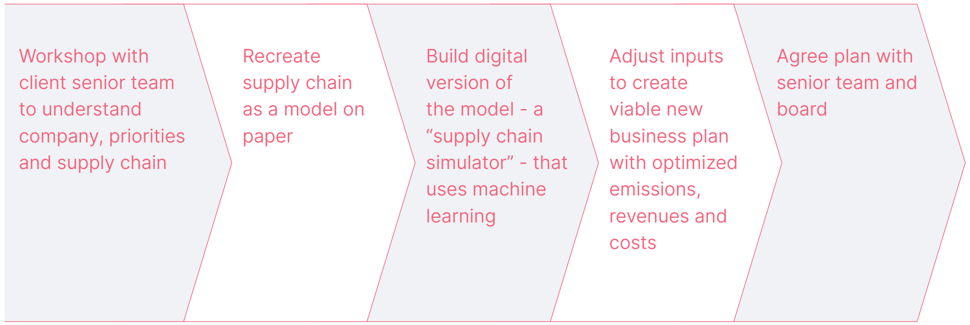

Figure 7: Strategy Augmented, example process from client work

Strategy Augmented allowed the firm to make adjustments to its plan as its circumstances change — for instance if new technology, or new sources of finance, become available to it. Strategy Augmented provides not only actionable insights to decision-makers within an enterprise, but also transparency, helping to align investors, employees, customers, partners and other stakeholders with the company's plan and priorities.

A particularly interesting aspect of our work together is that we created a bespoke simulator, which allows decision-makers to change the assumptions and variables to show how a range of future scenarios would affect carbon dioxide emissions and costs. Machine learning was also used to suggest and optimize choices available to management, adding significant depth to their group discussions and ultimate decision-making process.

Before Strategy Augmented, alongside numerical datasets, crucial knowledge about markets, processes, dependencies and other factors was held in the minds of the company's leaders and subject matter experts. Strategy Augmented allowed us to crystallize this group's knowledge and built solid relationships that ensured robust executive engagement. We added to this growing body of data and knowledge by conducting additional market research, which, for example included industry experts outside the company, and researching more widely on harmful emissions.

Using all of this relevant information, we jointly created a model, on paper, of how the business worked, reflecting interdependencies around carbon emissions and costs. The purpose of this model was to ensure among the senior stakeholders that the model itself met with their acceptance and trust before it was encoded. This model laid out elements such as demand, choices around which raw materials are used, pricing and use of different types of trucks for transport, together with energy sources, water consumption and other factors.

We then re-created this model as a mathematical system using dynamical models and concepts, refining it in consultation with the client. The mathematical model allowed us to capture complex business processes in a way that a computer can understand. We then used the data to create a bespoke piece of simulation software, we were able to deliver a custom supply-chain sustainability simulator, developed for this enterprise.

The simulator uses a dashboard format to show predicted outcomes in terms of emissions and financial costs if relevant elements of the supply chain are changed. The inputs can be adjusted using simple user-interface control sliders, and the outcomes can be seen in user-friendly graphic charts that can easily be shared for further discussion.

Results: From suboptimization to large-scale impact

Options for updating the client's business model included using greener energy and transport, replacing inefficient manufacturing steps, whether to use centralized service locations or be hyperlocal, different customer pricing strategies and whether to vary the set of manufacturing materials.

Once the options had been sufficiently refined, a viable plan was presented to the board. This plan is now being put into action and was recently announced publicly, further enhancing the company's already solid reputation for taking action to protect the environment.

We believe the potential for supply chain sustainability simulators based on a Strategy Augmented approach is significant — using, as they do, machine learning to help suggest the best routes to specific emissions and cost outcomes over defined time periods. We can also see them being used in pursuit of other goals, such as consolidating supply chains or managing dependencies.

The transparency of this approach was particularly useful as modifying supply chains involve not just changes to processes, but also significant changes to culture. We find that the modeling in a Strategy Augmented solution provided employees an important opportunity to see for themselves how changes to their work can impact the larger sustainability goals of the business.

For mature companies in particular, this transparency is especially very helpful to communicate with employees, as environmentally-driven process change could mean disruption to jobs and long-established ways of doing things.

Welcome to the family

In this paper, we have provided an overview and case studies from the family of AI Augmented approaches we have developed at Thoughtworks to address problems our clients face. These exist on a spectrum of solution approaches that range from tactical & operational to those which are highly creative & strategic.

What ties them all together is a commitment to augmenting human capability by building in feedback loops for learning.

AI Augmented moves us in the direction of taking advantage of the benefits of artificial intelligence in a systematically different way than the data-intensive, predictive approaches. While there is certainly a place for these, our clients are benefiting immensely from the AI Augmented approach described here.

If you are interested in finding out more, please contact us.

About the authors

Barton Friedland
Principal Advisory Consultant / Thoughtworks

As a digital pioneer at both Apple and NeXT Computer, Barton has worked with and implemented leading-edge technologies for over 30 years for large-scale global organizations. His primary areas of focus are strategy, decision-making, and execution. His educational background brings together computer science and anthropology as a powerful toolset to drive technical and organizational change — and the alignment required to drive it. Over the last 15 years he has focused exclusively on strategy & digital transformation, with emphasis on the development of distinctive competencies within organizations. This includes the development and management of intellectual property, acquisitions, as well as integration and evolution of business units within an organization. Barton brings a refreshing and empowering perspective to his clients that enables them to foster sustainable innovation, evolve organizational practice, and to outperform through an evidence-based, value-driven approach to strategy, decision-making, and execution.

barton.friedland@thoughtworks.com

This paper included significant research and contribution from a number of authors, editors, and designers, including Nemanja Cvijanovic, David Colls, Oege Dijk, Richard Gall, Jim Highsmith, Emma Kämäräinen, Nagarjun Kandukuru, Jake Luscombe, Vladimir Metodiev, Mikko Pajulahti, Aappo Pulkkinen, Jojo Swords, Lasse Tammilehto, Jessie Wang, and Ran Xiao.

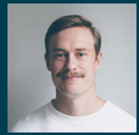

Jarno Kartela
Director, Digital Transformation / Marimekko
Previously Global Head of AI Advisory / Thoughtworks

Jarno is a strategist, advisor, speaker and occasional designer with an applied data science and emerging technology leadership background. His work includes designing and transforming how modern companies do decision making, creativity, strategic planning and sustainable growth as well as how they build immersive experiences and future-proof services. Areas of special interest include all types of emerging technology, creativity and innovation, circular economy business models and data-enabled decisioning.

jarno.kartela@marimekko.com

About Thoughtworks

Thoughtworks is a leading global technology consultancy that integrates strategy, design and software engineering to enable enterprises and technology disruptors across the globe to thrive as modern digital businesses.