

Building Intel's AI Center of Excellence



By Itay Yogev, Intel IT VP of AI

USD 1 billion in business value. That is what Intel IT's artificial intelligence (AI) group—quickly becoming an AI center of excellence—achieved in 2020. We are on track to double our business impact in three years by extending the depth and breadth of our transformational AI capabilities. We developed a set of best practices that optimizes the business value the AI group can bring to Intel. These best practices enable us to outperform industry benchmarks. (Industry-wide, 80 percent of AI projects never reach deployment, and of those projects that do reach deployment, only 60 percent are profitable.^[1])

Why Is an AI Center of Excellence Critical for Intel's Success?

Before we delve into the best practices of Intel IT's AI group, let's explore why using AI to its fullest extent is a key factor for Intel's future.

Intel's business complexity is growing exponentially. Intel is entering new and rapidly changing markets. For example, Intel recently announced the [IDM 2.0 initiative](#) and is expanding Intel product offerings to include discrete graphics, [infrastructure processing units \(IPUs\)](#) and more. With more products to design, develop and retain, manufacturing complexity increases correspondingly. If the number of knowledge workers^[2] was unlimited, this growth wouldn't be a problem. But there is a general shortage of knowledge workers today as the nature of work becomes more technically oriented and requires critical

thinking.^[3] AI is Intel's response to the knowledge worker bottleneck.

Intel uses AI for two primary objectives:

- To *augment* a human's decision-making by equipping employees with machines' intelligent "brute-force." At its best and at scale, AI and humans work in perfect synergy, where humans can take actions that are based on machine insights derived from enormous amounts of data that no human could process. For example, machines can provide a salesperson with actionable and timely insights about a non-trivial sales opportunity; these insights are derived from combining commercial and social media trends.
- Injecting human-like judgment into automated processes. By embedding AI into business processes, an algorithm can make instantaneous and complex decisions, that are personalized at scale and imitate a human's expert judgment. For example, an algorithm can actuate a different test scheme for every unit being manufactured based on previous test outcomes.

It's important to understand that AI does NOT replace humans. Instead, machines (AI) execute what they do best, and humans perform more complex intellectual tasks, such as forming strategies or applying creativity. Thus, the combination of AI and existing knowledge workers essentially creates a massive number of new, "virtual" knowledge workers that can develop more products and shorten time to market, with an increase in quality and improved features.

AI Center of Excellence Best Practices

I've been managing the AI group since its founding in 2009. The group consists of more than 200 AI professionals, and we have over 500 AI solutions in production, each of which uses several separate algorithms. We are hard at work today to further increase our velocity and scale to enable the next leap in our journey to transform Intel with AI. This will result in thousands of AI solutions applied to a diverse range of business challenges and data types.

According to research firm Gartner, only 14 percent of companies deployed AI in production in 2019. That figure grew to 19 percent in 2020 and is predicted to increase to 24 percent by the end of 2021.^[4] But at Intel, we've been deploying AI solutions for more than a decade. In fact, we've executed over 1,000 AI projects since 2009—and while we have had our fair share of failures, our success rate far exceeds the industry AI success rate. We used the collective wisdom gained during those projects to develop a set of best practices for success with AI across the enterprise. Our AI group has won Intel's Achievement Award,

the company's highest internal award for high impact and innovative projects, four times in the last three years.

The rest of this blog summarizes some of Intel IT's best AI practices.

Transform Critical Work. To make the biggest business impact, we focus our AI efforts on processes or areas that are most critical. These areas may differ from company to company. For Intel, the most critical focuses are threefold: anything involved with new product development, manufacturing and things that have a direct impact on revenue. We have aggressive goals for our AI projects, and we focus on developing algorithms that can produce significant improvements to the most critical business indicators at scale, which provides the best return on investment. We've made a conscious decision to focus our AI investment and precious skilled resources on these critical areas at the expense of AI maturity in other domains.

Focus AI Efforts on Specific Tasks. The algorithms we create are good at solving a specific business problem.^[5] Focusing on specific problems increases the likelihood of success and helps build momentum for an AI program in a given domain. Here are some examples:

- [Finding hidden insights that can assist Intel sales representatives](#) (16,000 assists provided in 2020)
- Adapting to changes (applied to thermal budget management)
- Automating actions to personalize unit testing
- Optimizing strategies to fine-tune weekly regression runs
- Synchronizing information, such as aligning supply and demand

You can find more information about many of these AI use cases in the most recent [Intel IT Annual Performance Report](#).

Define Specific Key Performance Indicators (KPIs). While KPIs are a well-known concept in managing any business process, a proper definition and diligent use of them are even more critical when working on AI capabilities. First, in most cases the KPIs will be part of the objective function that an AI algorithm will optimize. Second, we use KPIs to ensure properly set expectations between all stakeholders and to ensure success—actual and perceived over time. For Intel's business, we have prioritized the following AI KPIs: product power and performance (PnP)^[6], product quality, product and manufacturing cost, manufacturing yield, product time to market and revenue.

Carefully and Thoroughly Select Projects. Over the years we have created a set of feasibility

and value assessment questions we answer prior to investing in a project. In a nutshell, we make sure the project is tied to one of our defined critical work domains with direct impact on our main KPIs; we ensure we have a business partner willing to invest and own the required change management; and we ensure the project is feasible in terms of data and execution. By performing proper “due diligence” prior to investing heavily in a project, we can ensure we select projects with the greatest potential, or that at least fail faster.

Integrate AI Teams with the Business. Our AI teams are organized in a vertical manner. Each team is engaged in joint ventures with one of Intel’s various business units (BUs). Each team includes data scientists, subject matter experts, machine learning engineers, AI product managers, a “CEO” (the vertical team manager accountable for the vertical’s success), and a BU sponsor. Therefore, each team understands the business *and* AI, and they can determine which business problems are the best fit for AI. In essence, each team is treated like a startup company, built to disrupt with full autonomy. The AI competency-centered horizontal teams are built to support these vertical startups with MLOps’ reusable capabilities^[7], breakthrough AI R&D, business development for new verticals, and assistance with talent acquisition and retention. This structure empowers innovation, speed and the scalability to which we aspire.

Define Concrete Goals for Each Vertical Team. The teams set specific, measurable goals (relating to the relevant KPI), as well as aggressive scale plans. For example, the team responsible for building algorithms that help find product bugs may target finding 22 percent of the hidden bugs within three years, and 50 percent within five years. This helps foster accountability and transparency to senior management and for the team.

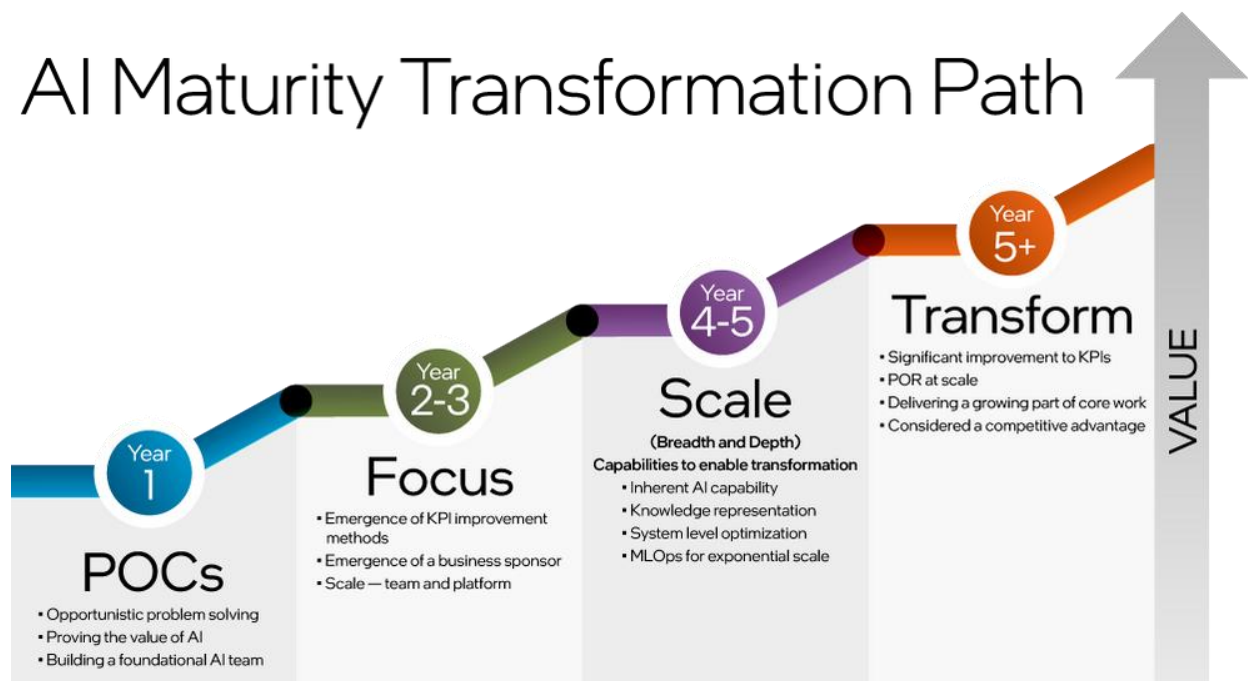
Help the Company Win with AI. Intel creates AI-enabled products, hardware as well as AI software. As a consumer of Intel’s products, the AI group as a whole helps Intel understand what people working with AI really need. Our feedback can include providing data scientist persona requirements and feedback on specific Intel AI products. We also help analyze AI workloads and the AI tools ecosystem.

Scale AI Everywhere. In August 2021, we launched a program called “AI Everywhere.” This program provides AI self-service tools for solving a vast range of problems with AI and helping Intel employees grow their AI skills through a central platform for AI training and access to AI resources. As part of this program, we also mentor communities of practice, through which we help provide advice and expertise to less experienced employees. Through this program, we also offer AI consultation and coaching services to upskill individuals and teams. And when a BU requires it, we provide on-demand tailored AI deliverables.

Understand that AI Is a Journey—and Be Ready to Move up the Maturity Scale. It takes time for AI to mature in a particular BU. Although we strive to provide business value from the

outset, it can take as many as five years for a BU to progress from proofs of concept to specific focused projects, to scaling a solution and finally to true transformation. Proofs of concept help build understanding and trust. Then, the team chooses a theme and starts thinking about which KPIs we seek to improve and how we can achieve those goals. Eventually the efforts can scale from one project to four, to 10, to 20 applications. One way we measure success, although somewhat subjective, is to ask the BU manager, “If we removed the AI can the BU still deliver?” That is, does the AI provide a competitive advantage? The below illustration provides the main stages in the path towards transformation, as we experience them.

And while each vertical team matures, the overall team also matures. At first, when we were just beginning our AI efforts, our group had not yet established its reputation and the BUs did not yet perceive AI as a game-changing tool. At this stage, it was important to show business value quickly. We looked for projects that could generate USD 10M in six months with just five people. Today, our group has a solid reputation of success and we have gained the trust of the BUs. This enables us to tackle more transformative projects that can take more time but also generate at least 10 times the business impact.



Plan for Scale—both Depth and Breadth. Without scale, AI’s potential is severely limited. So, of all the best practices discussed here, this one is perhaps the most important. We are building AI platforms around reusable productization processes (implementing cutting-edge MLOps concepts). These platforms enable exponential AI productization at scale (stay tuned for a new IT@Intel white paper on this topic in the coming months). Our AI platforms are unique because they are built of reusable components that we can assemble

in different combinations for different vertical domains. Each platform is integrated into the business processes and with the relevant information systems, which enables our AI algorithms to run in real-time at high scale. Only 20 percent of our overall team is invested in maintenance while 80 percent work on new concepts and projects—illustrating the effectiveness of our AI platforms.

We also continually research state-of-the-art AI methods to enable depth of transformation. The R&D team understands what is happening with AI in the industry and academia, and works to connect recent AI breakthroughs to our projects to take the vertical efforts to the next level. And while our current seven vertical teams have formed the foundation of our AI efforts to date, we are actively defining several new teams that will help us reach our 2x business impact goal.

Concrete Example of AI's Impact on Intel's Products

Intel has products coming to market that will illustrate the broad power of AI for business transformation. Across the vertical teams, each has contributed to KPI improvements. Consider these examples (intentionally vague, considering we're discussing a future product):

- PnP: Up to 30 percent performance improvement through AI algorithms in various scenarios
- Hardware validation: 22 percent of hidden bugs identified automatically
- Product engineering: Test time reduced by more than 20 seconds
- Manufacturing: Auto detection of baseline and new issues to accelerate yield ramp
- Sales: Recommendations and autonomous sales actions for accounts that are most likely to be interested in a new product

Our journey with AI is exciting—we're learning, solving problems and achieving goals, and then doing it all over again. By sharing our best practices for an AI center of excellence, we hope to enable other companies to start or extend their own AI journeys. To learn more about how Intel IT is helping transform Intel's business, visit [intel.com/IT](https://www.intel.com/IT).

[1] <https://venturebeat.com/2020/12/16/the-future-of-ai-deployments-reaching-production-is-bright-in-2021/>

[2] The term "knowledge worker"—coined by Peter Drucker in 1959—refers to an employee who applies theoretical and analytical knowledge, acquired through formal training, to develop products and services. (source: <https://corporatefinanceinstitute.com/resources/knowledge/other/knowledge-workers/>)

[3] http://www.digitalcommons.www.na-businesspress.com/AJM/PobstGF_Web14_1-2_.pdf

[4] <https://www.ciodive.com/news/artificial-intelligence-AI-business-case/594090/>

[5] <https://bdtechtalks.com/2017/05/12/what-is-narrow-general-and-super-artificial-intelligence/>

[6] Achieve the maximum performance for the lower possible power to enable longer battery life and better customer satisfaction.

[7] Machine learning operations (MLOps) is the practice of efficiently developing, testing, deploying and maintaining machine learning (ML) in production. It automates and monitors the entire ML lifecycle and enables seamless collaboration across teams, resulting in faster time to production and reproducible results. For more information, visit <https://ml-ops.org/>.