PEGASUS ONE

A GUIDE TO ARTIFICIAL INTELLIGENCE DEVELOPMENT BEST PRACTICES

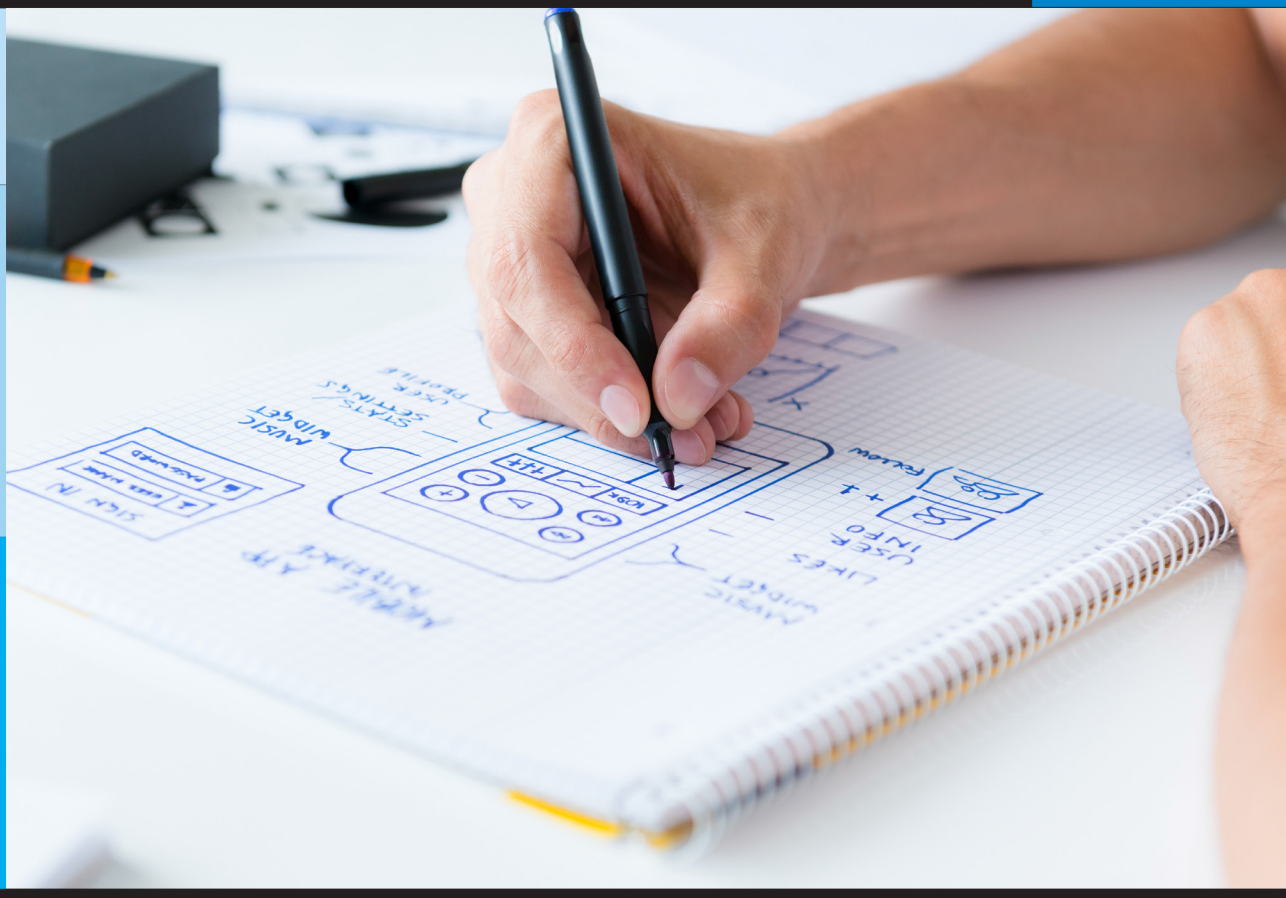

Pegasus One

1440 N HARBOR BLVD #900, FULLERTON, CA, 92835 | 10880 WILSHIRE BLVD., # 1101, LOS ANGELES, CA, 90024

CONTENTS

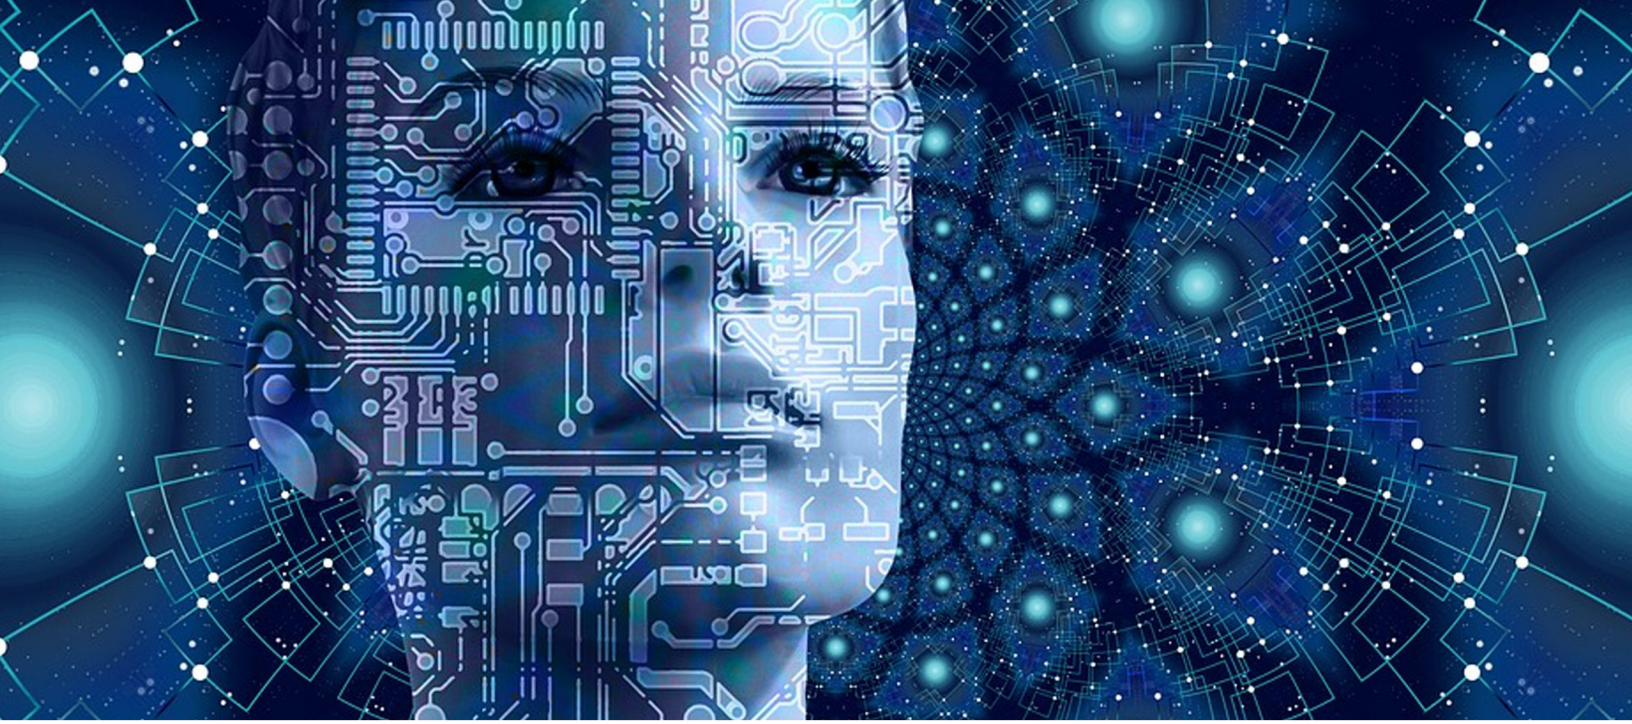

AI Development Best Practices

For digital transformers looking to take maximum advantage from AI, it is essential to learn, adopt, adapt, translate and add to the successful strategies of the digital natives. This represents a real challenge; it can be hard to judge which lessons are directly applicable to your business and your digital needs, and which ones need to be adopted, adapted, translated or ignored. Drawing from our breadth of expertise and hard-won experience we have found that the following rules increase the likelihood that digital transformers adopt and benefit from AI

BUILD TRUST IN AI: TAKE A RIGOROUS APPROACH TO DATA AND AI TRAINING METHODOLOGY

AI is going to make businesses more productive by automating repetitive operations, releasing resources to focus on higher-value tasks and remove the subjectivity from a human judgment call. However, organizations will only adopt AI if they can trust that the decisions made are within an acceptable level of risk. This is one of the areas where it is essential that digital transformers adapt the thinking of digital natives. The natives are able to accommodate some imprecision in AI decision making. For example, in consumer marketing, it is reasonable to run an AI on vast data sets, without the need to understand that data well. It might find that people who tweet a lot about shoes are more responsive to shoe adverts. If targeting all shoe tweeters consistently leads to 10% uptick in sales that is great. It does not matter who within that group clicked – just that shoe buyers are a clear group somewhere inside a broader group of shoe tweeters. On the other hand, if you are applying AI to spot when a plane engine might fail, you need much higher certainty. Knowing 1% of planes might fail isn't very good – you need to know which ones and when. Using AI to make precise predictions requires a rigorous approach to data and the AI training methodology. You cannot just let an AI loose on all the vast unstructured data you have collected from your engine measurements. You need to turn this data into rigorous training data tagged and curated by experts. This needs people who understand that data represents something in the real world – material strain, temperature readout, chemical reactions, and maintenance schedules – who can put together effective training regimes. AI should be designed not just by technical teams, but by people who understand the underlying data and what it represents within this business context. As a general rule: The greater the negative consequence of incorrect AI behaviour is to the business then the more rigorous the approach to data and training needs to be. From time to time we hear stories about AI making mistakes, such as the Amazon AI which designed phone covers from random stock photography, including pictures of heroin needles and men in adult diapers. Many of these AI bloopers are the consequence of not following a rigorous approach to training and data.

Ignoring the importance of trust is one of the highest risks to any investment in AI, undermining your strategy and leaving you vulnerable to being disrupted by the competition. Trust is invested in people, not the machines or the AI platforms.

MAINTAIN OVERSIGHT: AI NEEDS HUMAN GOVERNANCE

AI is a tool and can be easily misused if not fully understood – human oversight will always be an essential part of ensuring that an AI is delivering correct and trustworthy results. Regularly monitor AI decision making to identify when it does something wrong.

Taking random samples of AI outcomes and checking them against human experts for accuracy and appropriateness must be part of even the most basic governance process.

Establishing normal/acceptable operating parameters of an AI is equally important. AI is good at automating routine tasks, even complex ones like predicting drug chemistry or flying a plane. However, it cannot deal with situations outside its training where something rare or unexpected happens, such as chemical contamination or extreme weather. This is partly because of the complexity of non-routine events, and also because there is often not enough data on unpredictable events to train it. In this situation expert human intervention needs to be ready to step in and take over. Consequently and paradoxically the rise of AI will necessitate the maintenance of the expertise it is automating.

USERS FIRST: MAINTAIN AN UNWAVERING FOCUS ON USER EXPERIENCE

The balance between control and autonomy when applying AI technology is vital. For efficient collaboration between humans and machines, the appropriate level of automation must be carefully defined. This is even more important in intelligent applications that are designed to change human behaviors such as medical devices that incentivize humans to take their prescribed treatments on time.

The interaction should not make people feel like they are being monitored or controlled, but instead, assisted. This is an area where the digital natives are incredibly adept at embedding AI experience and expertise under the hood, while at the same time making sure that the complexity does not undermine the experience of users and customers. For example, Google Photos runs sophisticated neural networks, image analysis, and natural language understanding to give users the ability to search their photos using simple terms that describe what they are looking for. However, all the user needs to master is a search bar and a list of photos.

Your AI implementation needs to be the same; simple, intuitive and natural to interact with. Whether designing new materials with specific physical properties or scaling up production of a monoclonal antibody to treat a rare disease, any interaction with AI must be designed with the user in mind. The ultimate aim of AI is to act as a collaboration partner to the human, to inform decision making and generate trust in its data-driven conclusions. History shows us that without a relentless focus on user experience the reality is that it is more likely to confuse and distract.

TALENT: AI IS ABOUT TALENT AS MUCH AS TECHNOLOGY

The tendency to prioritize adoption of the most hyped and expensive technologies as the best solution to growing complexity needs to be continually challenged, whatever the given context. This applies to AI and digitalization technologies. The big disruptions have not followed purely from significant capital investments into existing shiny AI enabled platforms. It has come from the right mix of people, with a clear vision and the right blend of experience, problem-solving and technical expertise.

It is noticeable that successful digital natives have concentrated considerable effort on hiring and developing data scientists and AI engineers, and in evolving their ecosystems and processes for how these teams interact and work together. In the traditional sectors, those transforming into digital businesses are often likely to focus money and effort on platforms and infrastructure and neglect the importance of people and culture. These biases tend to be inherent in how they approach procurement and organize around information and knowledge.

Let's be clear, commercial-off-the-shelf platforms do play an important part in solving a well-defined task and are powerful when deployed on the right level of problem. Often, however, the business or organizational context is not given enough significance, particularly in off-the-shelf, AI enabled platforms. AI should be designed not just by technical teams, but by people who understand the problem, the underlying data and what it represents within this business context.

Organizations that are looking for genuine transformation using AI, and seeking to protect themselves from being disrupted out of business, must take a more comprehensive approach. Adopting a platform does not remove the need for embedding AI expertise in mixed teams, to lead the curation of data, the crafting of algorithms and the accurate training of data-driven models. Most high importance business questions are a unique combination of complexity and subtlety and accelerating their resolution using AI does not come with the turn of a handle.

BECOME AN AI EXTROVERT: ENGAGE EXTERNAL EXPERTISE TO BOOST YOUR INTERNAL CAPABILITY

Become an extrovert and look outside your organization for support in building a fast, effective AI capability. The skills required to make, train and implement AI are highly prized, and the demand for expertise has never been higher.

AI demands a rare blend of specialist skills; data management, machine learning, and data science may not all be readily available internally to the standard you need. Seek them out externally and form partnerships. Bring them into your organization and embed within your business teams. Seek an AI partner who has experience working across multiple industry sectors. Your problem is likely to have been solved elsewhere, and they can bring ready-made solutions to your organization. New thinking and fresh perspectives on your most challenging issues are drivers for game-changing innovation.

MIX PEOPLE: EMBED AI PROFESSIONALS WITHIN YOUR BUSINESS UNITS AND CUSTOMER TEAMS

The most effective AI implementation teams are multi-disciplinary, with representatives from IT, operations and business teams complemented by AI and data analytics experts. These teams should be led by those with a keen understanding of the business needs and the tangible outcomes that must be delivered by the AI. Their role is never to lose sight of how the programme is linked back to business value. They are the bridge between an abstract aspiration of what "better" is into discrete changes that deliver "better" in a way that respects real-world business constraints and the reality of people's working environments.

Critically, the AI specialists, who are not only skilled in the technical aspects of AI and modern data science, also understand the business domain into which any solution is to be deployed. They must be capable communicators in both the language of AI and the domain – they are the translators around which everything becomes joined up. The digital natives have shown how these translators accelerate transformation when embedded within core business and customer-facing teams, by balancing the various stakeholders and their differing expectations, demands, and requirements for AI. In our experience, these translators are the vital link between using a promising technology like AI and achieving a real competitive advantage.

BUILD MOMENTUM: PROGRESS VIA SMALL STEPS, MOVING RAPIDLY AND WITH PURPOSE

The temptation is to cover everything, launching into an all-encompassing initiative to build and implement a complex cross-enterprise AI platform. Our experience shows that starting too big too early undermines effectiveness and delays delivery of impact felt by the organization. When feasible, aim to cover the entire business in the fullness of time – but don't try and do it all from the beginning.

Start by laying out an AI roadmap, clearly identifying the decisions that AI can inform and the key business problems it should tackle. Quick wins are paramount, so the initial stages of the roadmap should focus on well understood, immediate opportunities with a clear statement of the outcomes required. Concentrate early AI projects on accessible and well-scoped problems then execute quickly and establish clear milestones to demonstrate success.

This offers several advantages. First, it rapidly increases the hands-on experience of working with AI. Internal resources are quickly skilled up, strengthening in-house expertise. The early delivery of real-world results with an immediate impact upon operations builds greater confidence compared to work done upon isolated proof of concepts which offer theoretical value. This generates vital momentum in your AI workstreams, where the demonstration of results increases interest and excitement in AI and creates new demand. As C-suite, executives, and leaders see early return upon investments, this increases corporate confidence and calms fears commonly associated with adoption of new technology.

Establishing this momentum builds lasting trust in the AI programme and its ability to deliver usable results. Without trust, the insights generated will be met with skepticism or dismissed entirely, condemning AI investment as another hyped technological flash in the pan.

PARALLELIZE: BECOME MORE AGILE, EXPLORING MULTIPLE AI PROJECTS IN PARALLEL

Enterprise-scale mo-mentum behind AI builds trust, skills, increases adoption, value delivery and the sheer buzz around AI in an organization. Accelerate this further by exploring multiple AI projects in parallel, in an agile delivery framework that can support rapid reallocation of effort and mix of skills. This flexibility is lacking in many pre-digital organizational structures, though it is the cornerstone of how the natives explore new ideas to deliver the next innovations.

Parallel project execution increases end value delivery by ensuring that the best ideas and solutions are progressed rapidly, with bad ideas allowed to fail early. This parallelized, experimental approach to AI development, combined with the rapid, agile delivery of the best solutions is used by the natives to bring their innovations and digital products and services to the marketplace as rapidly as possible.

As we will explore below, early attrition reduces operational waste, in terms of time, costs, materials, and resource, while still giving plenty of room for innovation to be explored rather than always choosing the low-risk options.

Rapid prioritizing of resource into the delivery of the most successful ideas demonstrates an AI strategy focused on the realization of tangible value which, it is worth repeating, is crucial in building trust.

EARLY ATTRITION: FAIL FAST TO GIVE AI QUICK SUCCESSES AND RICHNESS OF EXPERIENCE

Digital and AI are accelerating shifts in the shape of markets, and in customer demands and expectations. Accurate navigation in this world of uncertainty requires the ability to assimilate an understanding and experience of these changes quickly and to

detect the direction in which they are pointing. Humans need to sample their surroundings, and AI is no different.

AI learns by experience; this experience is embedded in its structure which is built from the data you train it upon and the priorities you assign it. How rapidly your process can gain this experience, and then adapt to exploit it, is a critical factor. Agility and staying openminded is vital. The most successful strategy is to trial many AI prototypes, monitoring and checking relative performance, then proactively improving training regimes to reflect previous success.

Expertise comes only from having many different experiences. Sampling widely and failing fast results in a far better trained AI, dependent upon fewer misplaced human assumptions and more open to innovation. Agility in your approach to delivering AI is therefore crucial if you want your digital transformation programme to release business impact.

DON'T BLINDLY HOARD DATA: MASTER ITS MEANING FIRST, SO THAT AI IS APPLIED IN A CLEAR CONTEXT

Because if you are not clear at the start, then clarity in any AI output will also suffer. An AI framework will give its most precise and most reliable answers when fed consistent data. Trust in the conclusions is essential if you want to deliver change with AI. Lack of trust is the most significant challenge facing AI.

It is tempting to collect as much data as possible, from sources both internal and external to your operations. There is a common belief that any data has value, with no exception, and that throwing more data into a pot will always increase the business impact that AI can mine and expose. This is not true.

The winners focus first on the business problem they need to solve, and then work smartly with AI experts and internal teams to identify only the data best suited to solving that problem. This promotes an understanding of the problem's context and builds an informed data acquisition and management strategy for the capture, control, and exposure of the data. That is how you start an AI programme if you want to end up with reliable, trusted results.

QUANTIFY VALUE: USE DATA TO TRACK AND QUANTIFY DELIVERED BUSINESS OUTCOMES

Organizations are adept at using data to measure performance; productivity of business units, operational costs, and effectiveness measures are standard. Metrics are often used to measure the performance of production lines, marketing effectiveness, customer engagement and KPIs are then used to determine success, progress or failure against targets. You should do the same with your AI implementations.

Define measurable goals and KPIs for each release of a new AI capability into your organization: how it has increased customer engagements (number of clicks); how it has improved the quality of output from your production line milling machines; how it has reduced non-productive time on your drilling rigs. This further establishes confidence to act upon the results delivered by AI.

Do not over-complicate measurement, but do enough to give you and your stakeholders confidence that you are delivering valuable results. If a deployment is failing, missing defined performance targets, that is still valuable as the lessons learned can be reinvested into the next iteration. Placing AI within such a managed framework delivers many operational and commercial advantages.

Performance tracking makes it possible to dashboard and visualize the impact of AI, to manage the programme as a service and embed it into a broader DevOps environment. Service management frameworks such as ITIL will only work if you are tracking this performance data.

Pegasus One is a privately held, professional IT services company and a **Microsoft** and **Amazon Partner** with its U.S. headquarters in the Orange County area, and satellite offices in Mexico and India. We provide custom software development and consulting services in various technologies including Microsoft Solutions, IoT, Mobile and BI. We are a CMMI Level 5 company with high focus on quality and ROI.

email: info@pegasusone.com
phone: +1(714) 485-8104
corporate office: 1440 N Harbor Blvd #900, Fullerton, CA, 92835
sales office: 10880 Wilshire Blvd., # 1101, Los Angeles, CA, 90024