Bringing Responsibility to AI

The Green Glass Approach to Raise Transparency and Lower Environmental Impact

As the use of AI is growing in the mainstream, organizations are looking for solutions that can ensure business benefits, regulatory compliance, as well as internal accountability.

The problem is that AI-powered tools for natural language understanding and processing (NLU/NLP) are often exclusively based on machine learning (ML) techniques. ML-only language solutions are essentially black box systems and their inner workings cannot be easily identified, explained or fixed. In addition, these ML-only models are typically trained on massive data sets via large language models (LLMs) that require excessive amounts of compute power, resulting in massive carbon footprints and a negative environmental impact.

Expert.ai builds transparent, sustainable and human-centered technology that works efficiently and reliably. We call our industry-leading commitment to building Responsible AI the "Green Glass Approach."

Most AI platforms are built like black boxes.

Instead, expert.ai takes the Green Glass approach.

The 4 pillars of Green Glass Approach

Transparent

To build trust, it's critical that AI provides outcomes that are easily explainable and accountable.

Being a good corporate citizen implies that you avoid bias and make a positive contribution to the world of users, shareholders and the environment around you. Today, AI systems are making decisions that impact many people's lives (medical treatment, candidate recruiting, loan applications, etc.), so understanding how they make decisions, as well as the implications of those decisions, is crucial. Unfortunately, it's very common, for example, in chatbots that are trained on generic material to show a diversity bias, which can negatively impact reputation or brand.

A system that performs tasks and reaches results that are not auditable presents a level of operational risk that should not be acceptable by any enterprise. The 'black box' phenomenon where it cannot be determined how a model arrived at a particular decision and why something happened inside an Artificial Intelligence system is an internal problem, and increasingly a regulatory problem.

Symbolic techniques, however, are white boxes that leverage rules and, in the case of expert.ai, a rich knowledge graph — all elements that are fully auditable and understandable by humans. When paired with machine learning, these symbolic techniques introduce much needed transparency into the model, offering a clear view on how the system behaves in a certain way to identify potential performance issues, safety concerns, bias, etc.

Sustainable

Enterprises want AI systems that are both effective and less energy intensive.

As a growing number of enterprises seek to mitigate the environmental impact of their operations, energy consumption and carbon footprint tied to training and maintaining large language models (LLMs) — based on massive data sets with billions of parameters — has become a primary concern. LLMs can be very expensive both in terms of training (and retraining) and processing for CPU/GPU, memory and time. The computing power needed to run these immense models is massive, and research has estimated that some LLM processes generate nearly 626,155 pounds of carbon dioxide emissions — more than four times the average lifetime fuel emissions of a car.

Expert.ai R&D Labs tests demonstrated that high levels of accuracy can be obtained using alternative methods, simultaneously reducing energy consumption and pollution produced by 100 times in the training phase and about 25% less in the prediction phase.

The expert.ai Platform enables you to leverage less energy-intensive techniques, such as the symbolic approach that is based on commonsense complements like knowledge representation. By mixing and merging different techniques into each step of the workflow, expert.ai offers a concrete answer to these issues and an alternative to consuming disproportionate amounts of energy typical of an ML-only technique.

As organizations pledge to attain net-zero carbon emissions, and as initiatives like the European Union's 2030 Climate Target Plan and Paris Accords gain momentum, organizations will need to start acting where they can. Addressing power usage is an easy and logical starting point.

Less computing = less energy = lower carbon footprint.

Practical

AI-powered solutions must solve real-world problems and provide tangible value to actual users.

When it comes to transformational technology such as AI, it's easy for the hype to overshadow the actual tangible results you can achieve via practical and operational applications. These inflated expectations around AI generate a difficult challenge in establishing the strategic value AI can offer for your business. Companies want tools that can transform their ideas into real-world applications that solve business problems at scale, not experiments that never make it into production.

Expert.ai's a hybrid platform is designed to solve the practical language-based problems that businesses face.

By integrating symbolic AI with machine learning techniques, we're able to ensure maximum performance for a wide range of use cases and domain-specific problems. Leveraging our knowledge and experience in implementing hundreds of practical AI applications, we provide the most comprehensive set of AI tools and workflow to support practical process automation and knowledge discovery use cases in the enterprise.

Human-centered

Enterprises need a "human-in-the-loop" approach, where data and inputs can be monitored and refined by users.

AI must capture expertise and humanize the work people perform, allowing them to be more productive and efficient.

Only explainable-by-design and interpretable-by-design AI models offer humans full control during the development and training phases. Because it's based on an open, interpretable set of symbolic rules, hybrid AI offers a simpler way to correct insufficient performance. So, if an outcome is misleading, biased or wrong, users can intervene to prevent future mistakes and achieve the success metrics most valuable for each use case and improve accuracy, all by keeping a Human in the Loop (HITL) who is a subject matter expert.

Whereas black box machine learning models only offer the opportunity to add more data to the training set without an opportunity to interpret the results, a hybrid approach uses linguistic rules to tune the data in the model. Hybrid AI makes it possible for people to be in control of a model.

Conclusions

When it comes to Responsible AI, as a business owner you should look for platforms that:

- Know (and can demonstrate) how the AI is going to produce the results and without bias (Explainable)
- Provide effective results without excessive energy consumption (Sustainable)
- Ensures the AI approach is able to solve real-world problems (Practical)
- Ensures that humans are in control from the outset (Human-in-the-Loop)

With the Green Glass Approach, we want to confirm our commitment to help organizations put in practice Responsible AI initiatives, driving awareness to grow natural language AI technology in the right way, on the right path.