Responsible AI for telecom

The next step

White paper

Anne Lee

In the last couple of years, adoption of artificial intelligence in telecom has increased substantially. We see standards defining architectures to support the use of AI. We see integration of AI into tools for network planning and trouble ticket root cause analysis. AI is being explored and implemented for operations, troubleshooting, predictive maintenance, self-organization, sustainability, and security. Additionally, telecom research is investigating using AI for new state-of-the-art real-time services.

As our industry ramps up deployment of AI in telecom products and tools, we must also ensure that responsible, trustworthy, and ethical principles are met, including defining support and new frameworks into telecom architectures in standards.

Contents

Introduction

Ethics has been a growing concern of scientists as AI technology becomes more and more pervasive in everyday life. To address this concern, principles have been defined for Responsible AI – also called Trustworthy AI or Ethical AI.

Without consideration for responsibility in AI, the field has encountered some serious issues. Here are three examples:

1. The carbon footprint for both training and inferencing has been growing out of control
2. Training AI models using only historical data have resulted in further reinforcement of unfair social biases or subpar resolutions of problems
3. Allowing AI to make predictions as black boxes based mainly on correlation can result in potential users of AI resisting adoption; without transparency and explainability, these users do not trust the conclusions from the AI models.

We will see in this paper that solutions to address these and other ethical issues can actually provide benefits – at times rather significant – that have nothing to do with responsibility. Being a good citizen has many rewards.

Governments take action

Ethics is a concern of governments around the world as AI has become more commonplace. In recent years, laws and regulations around artificial intelligence have begun to appear. These began as early as 2019 in Australia with their Artificial Intelligence Ethics Framework which guides businesses and local governments to design, develop and implement AI responsibly.

In North America, the US developed the 2020 USA Memorandum M-21-06 or “Guidance for Regulation of AI Applications” to give guidance to federal agencies in the development of regulatory approaches to the use of AI in the private sector. In 2022, the US unveiled an AI Bill of Rights outlining five protections Americans should have in the AI age. Also in 2022, Canada passed bill BLL C-27, the “Artificial Intelligence and Data Act”.

In Europe, the EU Commission created the 2021 Regulatory Framework for AI/ML Systems while the UK National AI Strategy intends to establish “the most trusted and pro-innovation system for AI governance in the world”. It is a10-year strategy.

In Asia, leading think tanks in China and India released white papers in 2021 on Trustworthy or Responsible AI.

And in Africa, three nations, Rwanda, Tunisia, and Ghana, are developing national AI strategies that are tailored towards developing AI governance frameworks, policies, and AI ethical guidelines.

The solution

Nokia has identified six pillars of Responsible AI. They are:

1. Transparency
2. Sustainability
3. Reliability, safety, and security
4. Privacy
5. Fairness
6. Accountability

For telecom, there are specific approaches that can be taken to address each pillar. In addition, Nokia is defining a standards framework to support responsible AI principles.

Figure 1 - The six pillars of Responsible AI for telecom

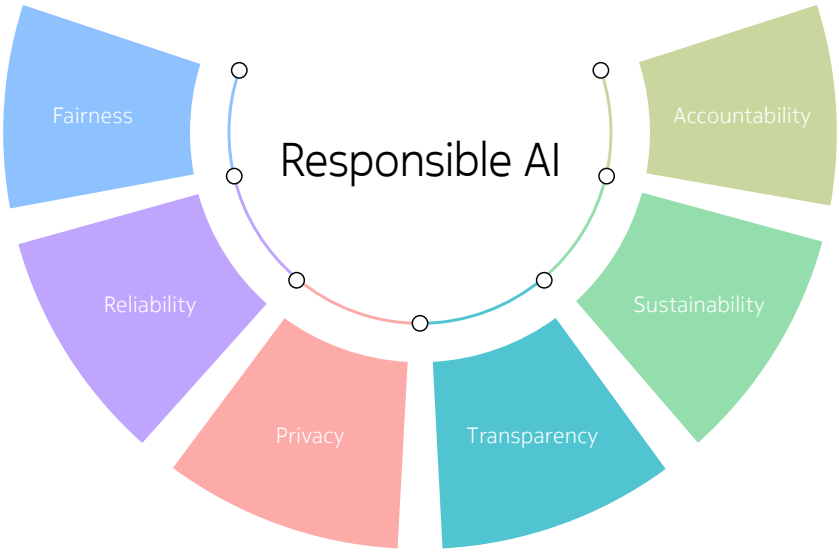

Transparency

There are multiple reasons and benefits for ensuring the transparency of an AI agent. First and foremost is to ensure the trustworthiness of the results. The more transparent an AI agent is, the more likely people will use it, and overall adoption of AI will increase.

How do we ensure transparency? We need explainable AI. An AI cannot be a black box. Explainable AI can be achieved in one of two ways. The algorithms can be innately explainable, which is true of decision trees, for example. There is also research work on the merger of causality with AI which has the potential to make the model and its results explainable. The other method is through the use of explainable AI toolkits. These toolkits can determine what primary features were used by a model for prediction.

As we all know, correlation does not mean causation, but it is possible for an AI to jump to entirely the wrong prediction based purely on correlation. Use of causality and counterfactual “what if” analysis may help explain an AI’s predictions or at least identify when it has come to an obviously wrong conclusion. Explainable AI is used to describe an AI model, its expected impact and potential biases.

The benefits of AI transparency are not only ensuring trustworthiness; it can also help AI engineers tune their models for higher accuracy. By knowing what features the AI primarily uses to make predictions, an engineer could make model adjustments such as changing hyperparameter values or feature set adjustments. This is especially needed when there is data drift due to changes in the environment, as in RF conditions, and re-training or re-architecting of the model is required.

And finally, explainability and transparency can help with root cause analysis. This is important in telecom, especially for troubleshooting.

Figure 2 - Multiple paths to realizing sustainability

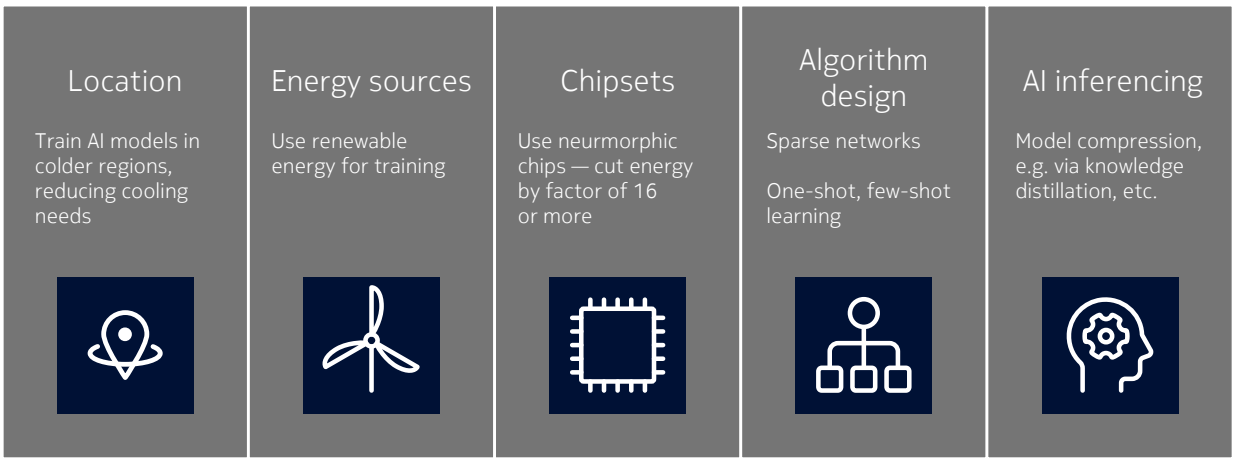

Sustainability

The sustainability pillar is concerned with both developing AI models sustainably and using AI to make systems, such as telecom systems, more sustainable. We have identified four areas of sustainability:

1. Training. AI model training sometimes requires a great deal of compute power that requires cooling, which requires energy. To reduce energy usage, put data centers for AI training in cooler geographic regions if possible.
2. Energy source. Renewable energy should be used whenever and wherever possible.
3. Chipsets. Chipsets that are designed specifically to run AI models are much more energy efficient than general purpose processors or GPUs. Earlier in 2022, Intel Labs demonstrated experimentally that a large neural network can process sequences such as sentences, consuming four to sixteen times less energy while running on neuromorphic hardware than non-neuromorphic hardware. And in March, MIT announced its new MIT AI Hardware Program, which will be a collaboration between academics and tech companies to develop energy-optimized machine-learning and quantum-computing systems.
4. Algorithm Design. Energy-efficient algorithm design for training and inferencing is key. The simpler the algorithm, the less processing is needed for training. Also, the less data required to train a model, the less energy is used. Approaches for using less data include one-shot, few-shot and zero-shot learning. Approaches to compress or minimize the number of nodes and layers in a model include a method called knowledge distillation where a “traditionally” trained large model is used as a teacher to train a smaller “student” model with fewer nodes and layers.

The above summarizes how to make AI training and inferencing sustainable. Now, let's talk about how AI can be used to make telecom systems sustainable. Here are four examples.

1. Use AI to learn usage patterns so that the telecom system can turn off power for components in the radio, transport network and core when not needed to optimize for the actual capacity required.
2. Optimize the amount of output power for RF transmission. There are already strategies to weigh and balance the amount of output power needed to maximize quality for an individual user while minimizing interference to other users. Sustainability should be another factor or feature in the final decision making.
3. Algorithms for load balancing and geographic redundancy must also factor in sustainability in deciding the minimal number of components to keep powered up.
4. AI-enabled network planning and deployment tools should include sustainability features such as the ability to use renewable energy at a location, minimizing the number of required sites balanced against minimizing the amount of RF power required.

Reliability, safety, and security

AI models must be safe, performing as intended. In other words, they must be reliable, resistant and resilient to being compromised by unauthorized parties – meaning they must be secure.

If you are a veteran in the telecom industry, you may start to notice an emerging pattern. Some of the goals of Responsible AI are similar to goals that the telecom industry has been addressing for years. The underlying drivers for the goals might not be the same. But, the goals themselves are the same. For instance, initiatives for cost reductions may also have as a side effect, power or energy consumption reduction.

Building secure and reliable networks are also key goals of the telecom industry before the rise of AI. So, when we integrate AI into our telecom networks, we must carry forward our requirements for security and reliability to also cover AI, which will have its own flavor of issues and solutions.

For example, when training an AI, it is vital that the training data be both secure and private. We must implement methods to protect against the poisoning of the data, which would then directly affect the performance of the AI. Once a model is trained, there must be assurances when we use the AI that the input data for inferencing is not compromised to the point of confusing the AI. We are all familiar with examples where seemingly innocuous information is added to the input data resulting in nonsensical predictions by the AI.

Privacy

As we all know, data is the fuel of artificial intelligence. And the privacy of that data is not just critical but protected by government regulations. Two well-known laws are GDPR and HIPAA. Europe's GDPR rules protect the privacy of people's data in general. While HIPAA is a set of data privacy regulations for the healthcare industry. These rules are put in place to protect people's privacy and to give them agency over their own data.

It is especially important to comply to these regulations in AI solutions from all relevant perspectives. This includes ensuring that the collection and storage of data for training and inference is secure and private. It also means that the transmission of the data from one location to another must be secure and private. There are a number of methods available and being developed to address data privacy for AI.

The most basic method is data anonymization. Removing personal identifiable information (PII) features that can identify someone, such as their name, address or phone number, is a way to anonymize data records. However, research has shown that data anonymization alone is insecure. The research has found that it is easy to reverse engineer the data by combining it with other data to determine the original source of the data. In fact, it has been shown that a person in the US can be identified with just three pieces of information: zip code, gender and birthday.

In response, researchers have come up with two classes of privacy solutions. One class protects the data before it enters the model. A popular method in this class is called differential privacy. The other class of solutions builds protection into the modelling process itself such as with federated learning.

Figure 3. Two classes of privacy solutions

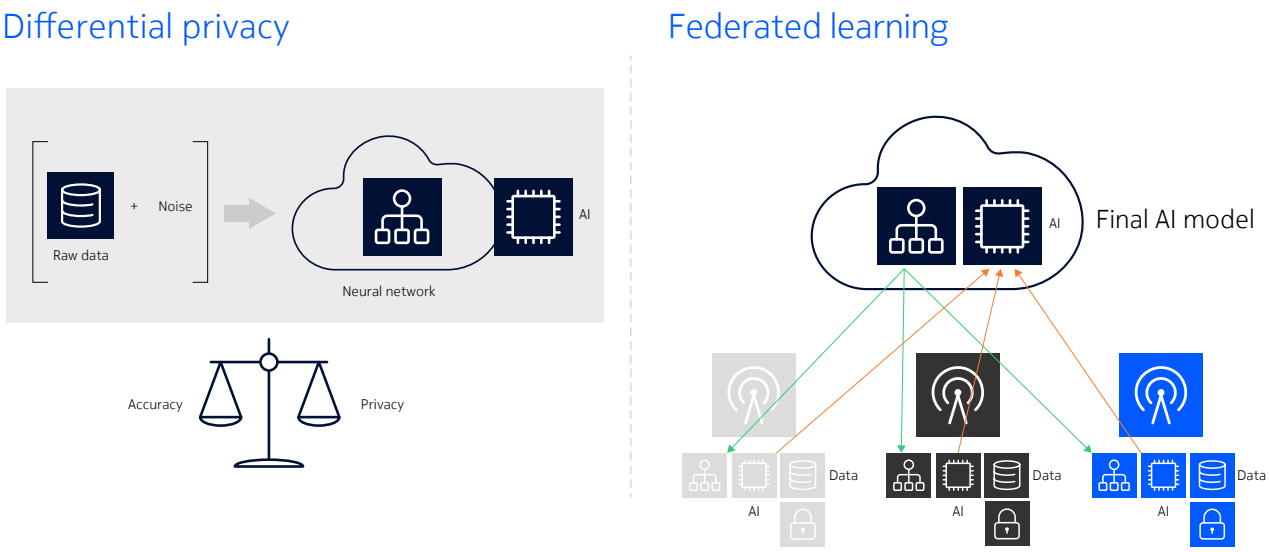

Differential privacy - at a high level, differential privacy methodologies add noise to the dataset. If the right amount of noise is used, then the source identifying features for each data sample is protected while not compromising the essence of the actual information that is needed by the AI. For example, a telecom AI might want to know if there is a pattern of dropped calls in a region, along with associated attributes of the call, without needing to know who the calling and called parties were that got dropped. These latter features could be modified by the added noise to mask them. A challenge for this method is that too much added noise can affect model accuracy and too little noise may not provide enough protection. It can be challenging to find the sweet spot.

Federated learning - in this case, the data is not moved from its source device. Instead, if the AI requires training using data from different sources, for example from multiple base stations, a copy of the AI algorithm is provided to each source. Then at each source, a version of the AI model is created by training the algorithm with the local data. Each version of the AI model is then sent from each location back to a central point where the model parameters are “averaged” across the multiple versions of the model to create a single model. This ensemble AI model is the final model used for inferencing.

Fairness

Fairness in Responsible AI usually brings to mind the goal of ensuring that AI models don't perpetuate historical, societal or institutionalized prejudices and biases, especially in the context of decisions for hiring, loans, housing, healthcare, and criminal justice.

For telecom networks, there are two main areas to address and an emerging third.

Resource allocation is one. This includes ensuring that schedulers are fairly assigning resources to subscribers or devices. Even prior to the rise of Responsible AI, this was a goal of telecom networks, sometimes due to government regulations. For example, telecom networks use proportional-fair scheduling, which is a compromise-based scheduling algorithm for resource allocation. It is based upon maintaining a balance between two competing interests: trying to maximize total throughput of the network (wired or wireless), while at the same time allowing all users at least a minimal level of service. As traditional algorithms are replaced with AI, fair resource allocation must be maintained if not improved. This goal extends beyond a single network to spectrum sharing. Spectrum sharing must also ensure fairness.

Network planning tools are being upgraded today to use AI, which must take care to ensure that sustainability is optimized and that there is no bias against geographic regions or neighborhoods.

The second area in telecom to be addressed for fairness is ensuring that all users, including users with different needs, are supported. If AI chatbots are deployed, especially for customer service, then those chatbots must support multiple modes of communications such as audio for the vision-impaired and video/text for the hearing-impaired.

Related to the second area, with the seemingly rapid emergence of large language models (LLMs) and especially the superset of generative AI technologies expected to touch aspects of all industries, telecommunications will not be an exception. It is therefore important that data ownership, content rights and copyrights be clearly defined and honored. This falls into both the Fairness and Accountability pillars.

Accountability

The final component of Responsible AI for telecom is accountability. Without consciously deciding and determining who the stakeholders are for deployed AI agents right up front, there can potentially be issues down the road if the AI does not work as expected. In the case of self-driving cars, it is easy to understand how this can be a problem. If an accident occurs, then who's accountable? The AI model vendor? The organization providing the training data or the vendor for the device that collected the input data for inferencing? The automaker?

For AI used in telecom networks, the consequences may not be life and death – unless emergency services are impacted – but accountability should be clear. To make sure of that, AI agents should be developed and deployed by consulting, involving, collaborating, and empowering all stakeholders early in the AI development process.

Standards frameworks

Telecom standards bodies have begun to recognize the need to support responsible or trustworthy AI in their architectures. For example, in ETSI, the zero-touch network and service management (ZSM) AI/ML framework is adding a domain for "trust and security" that covers:

- **ML Data Trust Management Service** - enables the consumer to manage the trustworthiness of data (training, inference) required for ML models based on the desired level of ML data trustworthiness

ML Data Trust Evaluation Service – measures the trustworthiness of ML data (training, inference) and detects trustworthiness degradations

ML Model Trust Management Service – manages ML models based on the desired level of ML model trustworthiness

ML Model Trust Evaluation Service – measures the trustworthiness of deployed ML models and detects ML model trustworthiness degradation

ML Event Notification Service – provides an event notification service for ML. The ML-specific events notification service has the same set of capabilities as the generic event notification service but provides notifications related to ML events such as ML QoT (quality of trust) threshold crossing, for example, the fairness metric threshold, robustness metric threshold and explainability metric threshold.

In addition, Nokia is actively working on a Trustworthy AI framework for other telecom standards within 3GPP. These are examples of nascent standards frameworks that are actively evolving today.

Business benefits

Addressing the principles for Responsible or Trustworthy AI is not just necessary to comply with rising government regulations. It is also not just necessary to be a good citizen and a good steward of our Earth. As it turns out, it is also good for business.

Here are some examples of how.

1. Reducing overall energy consumption by turning off components when not needed may reduce OPEX costs in the long run
2. Reducing the amount of data needed to train a model can result in new methods that support additional use cases:
    a. Zero-shot learning is a transfer-learning approach that re-uses a pre-trained model to make predictions for new classes of data that it has never seen before, and for which it may be hard to collect training data
    b. One-shot or few-shot learning addresses a subset of use cases, such as creating a building-security AI face-recognition agent, where it is impossible to collect more than one photo sample of each employee
3. Creating smaller models, i.e., with fewer nodes and layers, can minimize the size of devices
4. Explainability to ensure transparency can also help the engineer better optimize and re-architect an AI model when needed.

Also, as noted earlier, some of the goals of Responsible AI are very similar to goals that the telecom industry has been addressing for years. We must ensure that as we upgrade our systems to use AI, that these goals continue to be met if not exceeded.

Conclusion

In summary, the key points for the support of responsibility, trustworthiness and ethics in AI for telecom networks are:

- Responsible AI is essential to the adoption and successful use of AI
  - In fact, governments are requiring it and standards are defining support for it within telecom systems
- To ensure that an AI-enabled telecom network is responsible, the following is needed:
  - Transparency and explainability of the AI
  - Reduction of the carbon footprint of the AI
  - Use of AI to reduce the carbon footprint of telecom systems
  - The AI must be made secure and reliable
  - The data used by the AI must be kept private
  - The AI must be fair
  - Assignment of accountability

Additionally, AI could be used to ensure the security of the telecom network. In conclusion, Responsible AI is good for business, as well as for being a good citizen and steward of the earth.

Abbreviations

3GPP	Third-generation partnership project
AI	Artificial intelligence
ALG	Application layer gateway
BNG	Broadband network gateway
ETSI	European Telecommunications Standards Institute
GDPR	General Data Protection Regulation
GPU	Graphics processing unit
HIPAA	Health Insurance Portability and Accountability Act
MIT	Massachusetts Institute of Technology
ML	Machine learning
OPEX	Operating expenses
PII	Personal identifiable information
QoT	Quality of trust
RF	Radio frequency
ZSM	Zero-touch network and service management

References

1. 3GPP Release 8. http://www.3gpp.org/specifications/releases/72-release-8
2. 3GPP Release 9. http://www.3gpp.org/specifications/releases/71-release-9
3. ETSI DGS/ZSM-012 https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=62010

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering the future where networks meet cloud to realize the full potential of digital in every industry.

Through networks that sense, think and act, we work with our customers and partners to create the digital services and applications of the future.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: CID212898 (July)