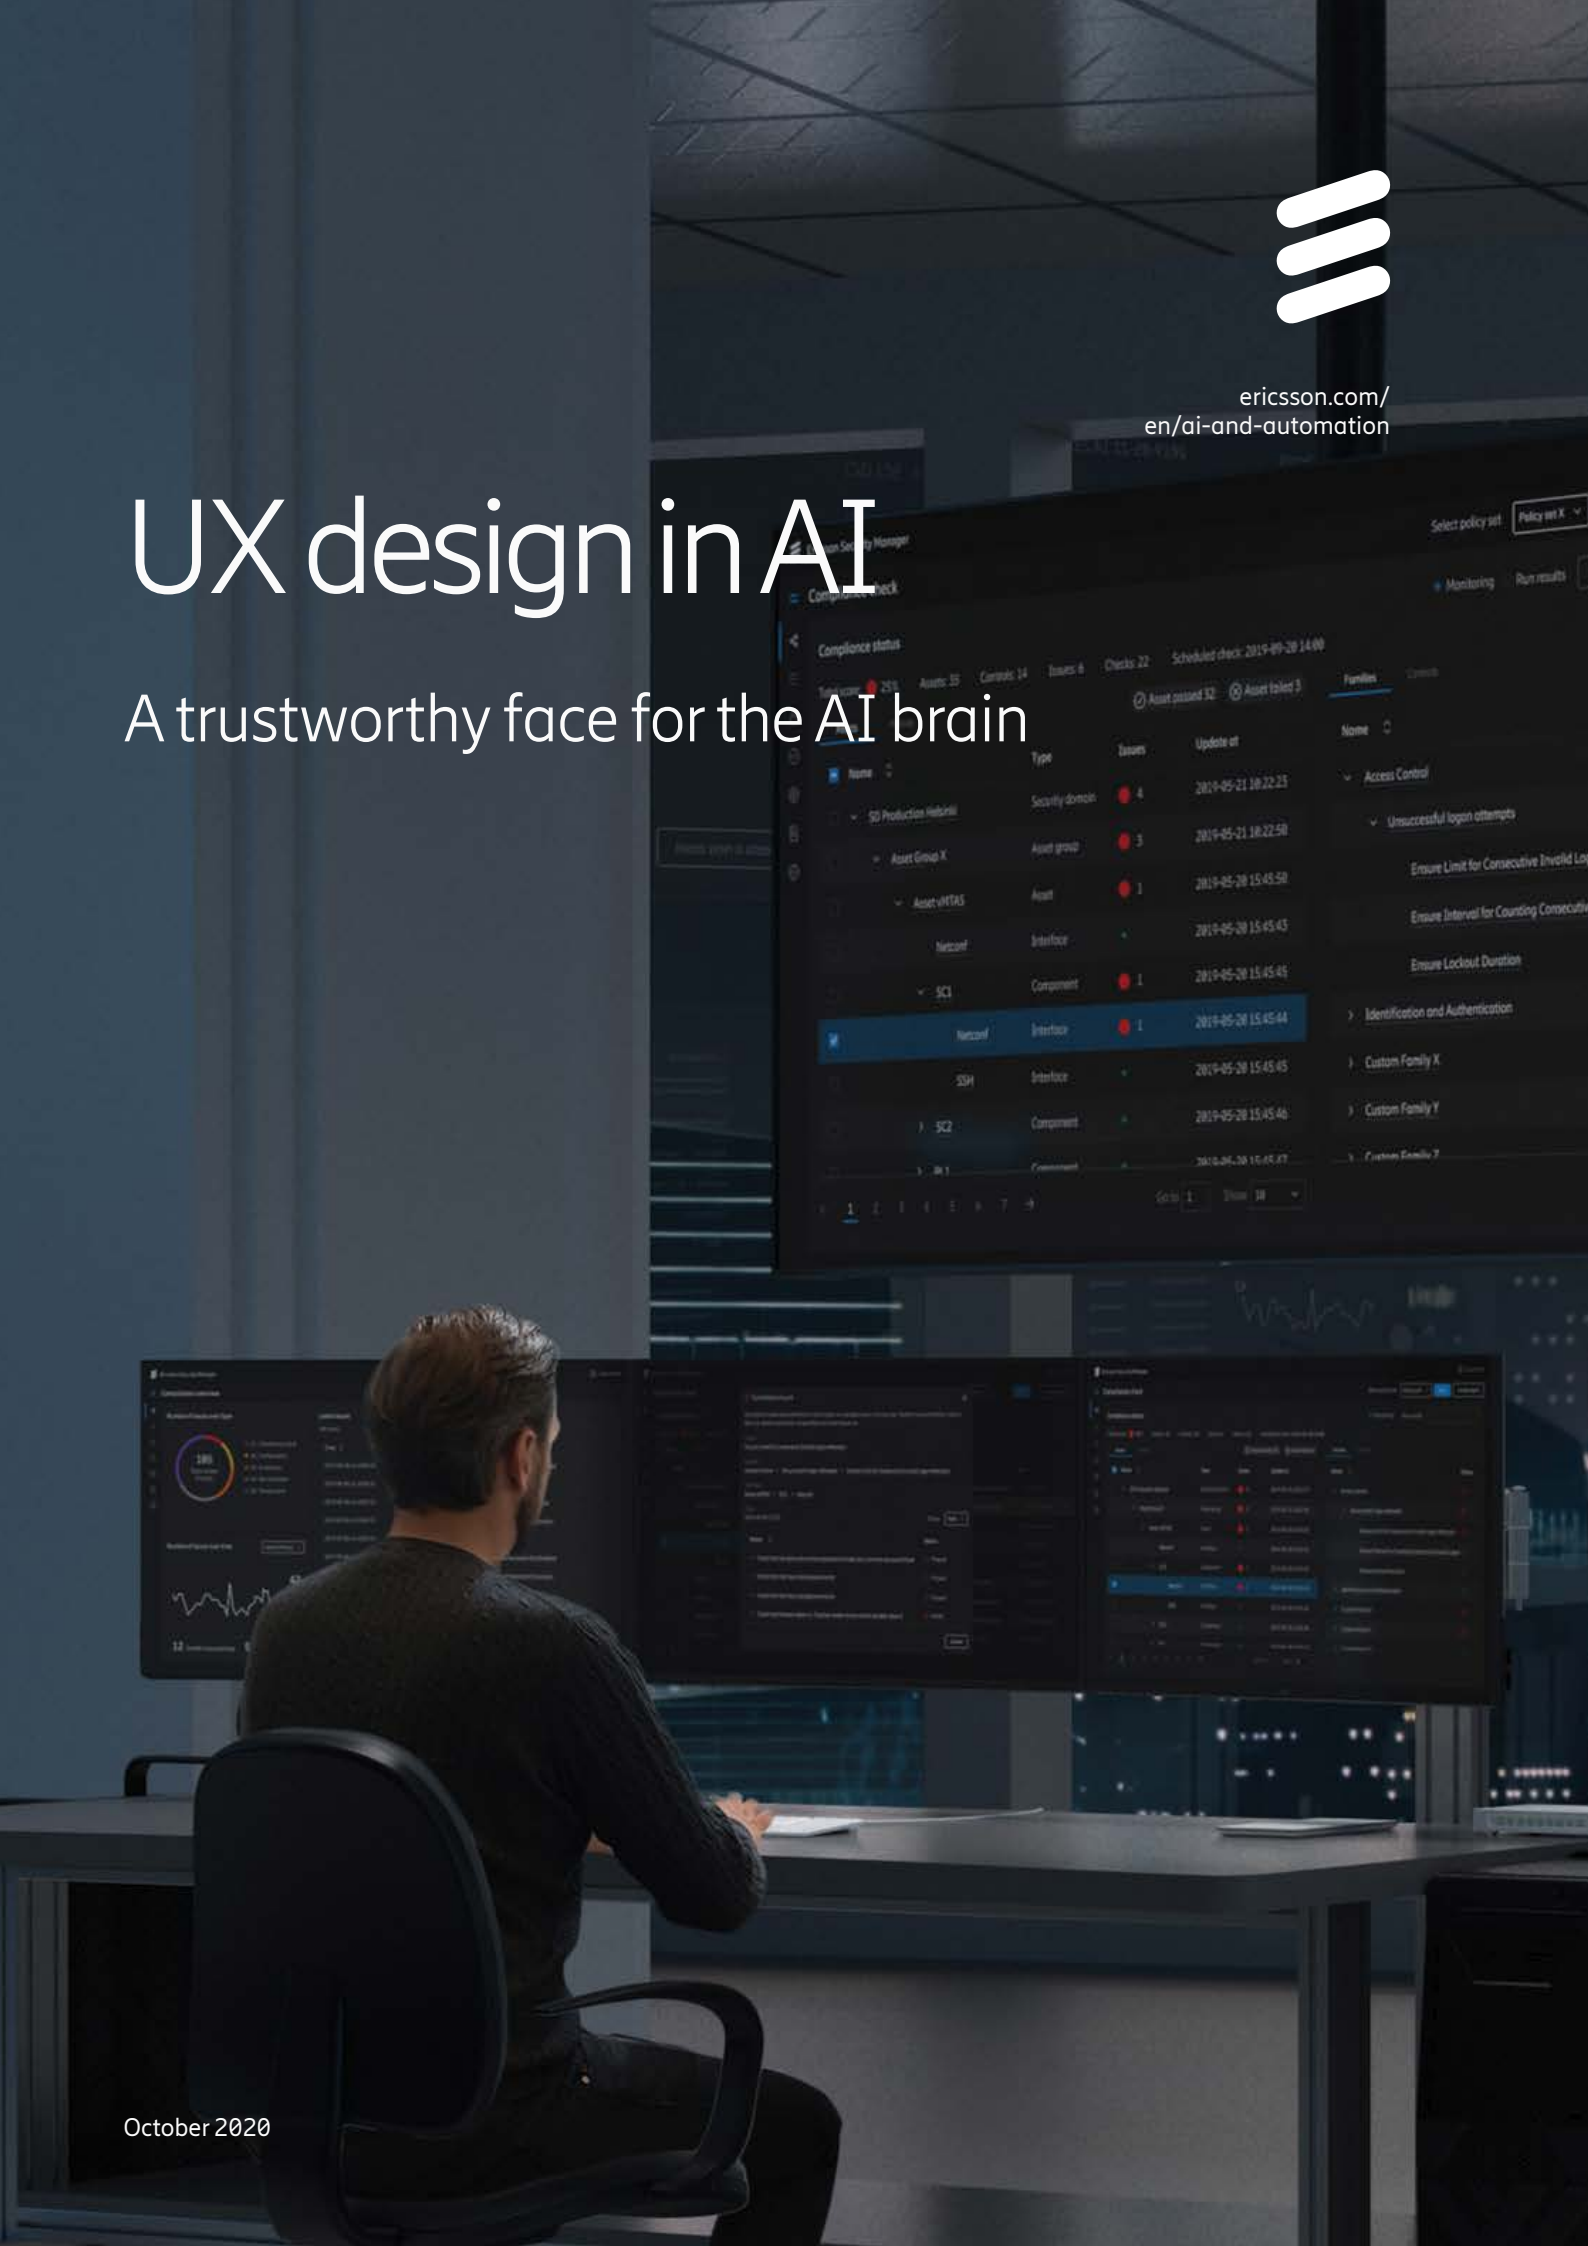

ericsson.com/
en/ai-and-automation

UX design in AI

A trustworthy face for the AI brain

October 2020

Introduction

Contents

Currently, computational capacity is doubling roughly every 18 months.

The pace of this development, amplified by rapid improvements in software, has resulted in artificial intelligence (AI) and advanced algorithms that are quickly evolving to understand and interpret some of our most complex natural processes.

At the same time, the ability to access this capacity is multiplying due to sharp increases in bandwidth, improvements in latency and other quality of service parameters with technologies such as 5G. Interfaces are also becoming more seamless due to advances in cloud computing as well as visual, tactile and verbal interface technologies.

These exponential improvements have brought what, just over a decade ago, were considered industrial-strength processing and communication capabilities into the homes and hands of individuals everywhere. As industries adopt these technologies to modernize and automate their business processes to increase value chain efficiency and effectiveness, a new service-based concept for the technology has emerged.

The self-driving or autonomous car is an example of this new concept. Eventually cars will no longer have drivers, a fundamental change in the concept of a car. The passenger of such a vehicle will interact with it on a much higher and abstract level as a service.

When we apply this concept to the telecom sector; that is, creating a "self-driving network", AI technology will be the brains behind this change. This presents two main challenges for those developing the concept and service:

- The conceptual shift from today's understanding of what a network is, becoming something more abstract than what it is today, operating on new parameters
- The fact that a user of such a service will interact with the system on a much higher, more abstract level

Therefore, the understanding of the business goals and the user of the system is key to success.

With the role of users shifting from drivers to passengers and from operators to managers, designers will need to create highly collaborative solutions allowing tangible and reliable interaction between AI technology and the user.

In light of this, our Experience Design team has been researching and developing how to design trustworthy, AI-powered services for telecom operators. Through designing the Cognitive Operation Support System service concept, we have identified four components of human trust that can be applied to AI-powered systems.

These four pillars – competence, benevolence, integrity and charisma – are the key areas designers and business owners need to address to be successful when it comes to the adoption of AI.

In this paper, we will share our experience of designing a trustworthy, AI-powered Cognitive Operation Support System (OSS) service.

AI today

The current face of AI

AI is an umbrella term encompassing many different methodologies and concepts, referring to any machine developed to perform tasks that would require intelligence if done by a human.[1] Although the media commonly portrays AI capabilities as superior to human capabilities – that is, as an artificial super intelligence (ASI) – the truth is quite different. Since the earliest explorations into the AI field, scientists and practitioners have sought to create a computer with a level of intelligence similar to a human. Known as an artificial general intelligence (AGI), these would be machines with a reasonable degree of self-understanding and autonomous self-control, able to solve a variety of complex problems in a variety of contexts.[2] Despite the huge advancements of AI, especially in the last decade, we are still far away from being able to create an AGI, let alone an ASI.

The current form of AI we are working with is known as an Artificial Narrow Intelligence (ANI), or "weak AI". ANI systems are created to carry out specific tasks showing specific aspects of intelligence in a specific context.[3] All current applications of AI, whether it is an autonomous car, a chatting app camera filter, or an intelligent OSS, are all considered narrow or "weak" by this definition.

An easier way to describe the role of current AI applications is to call them "agentive technology"[4] – whereby we can think of them as our assistants or agents, handling a discreet task and not the entire job.

In this current context, humans still need to have a view of the bigger picture and are still required to supervise, evaluate and orchestrate the work of these AI systems.

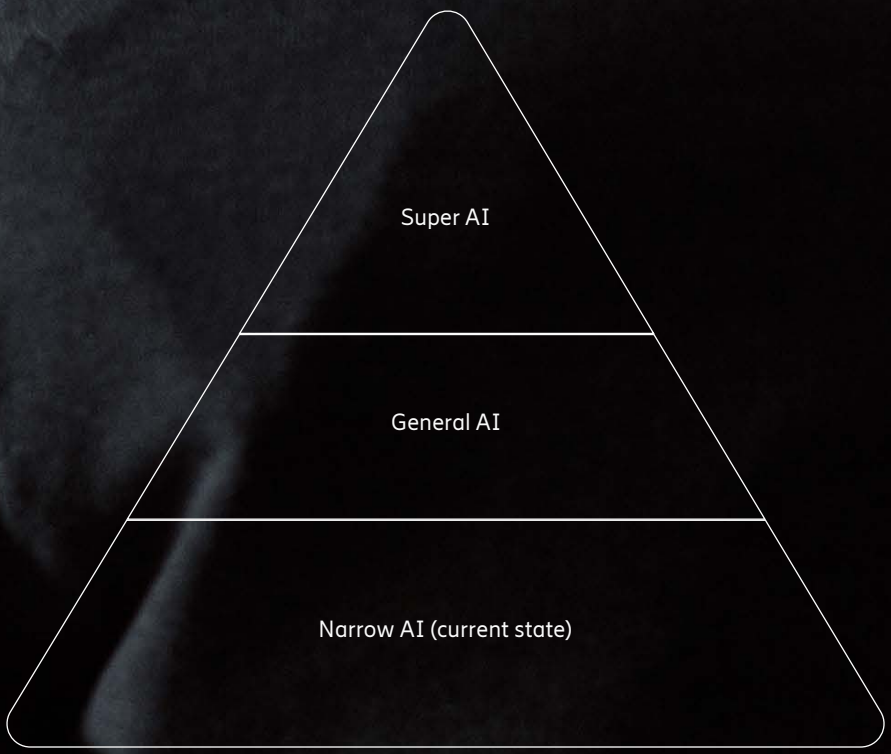

Figure 1: Levels of AI

[1] Minsky, 1982
[2] Goertzel and Pennachin, 2005
[3] Gobble M. M., 2019
[4] Noessel, 2017

Inside a network operations center

The key to AI success

Trust as a vital component in AI adoption

There is an increasing trend of digital assistants appearing in many different aspects of our lives. Powered by machine learning (ML) models, they analyze data to come up with statistical probabilities that can be used to offer recommendations and make predictions and decisions, from suggesting the optimum route to take on a commute, to adjudicating whether we are viable for a loan or not.

Although they sound less impressive than the idea of the ASI's superior artificial brain, the recommendations, predictions and decisions taken by the AI systems can be considered a fundamental change in the way humans are using tools – a paradigm shift in the human-tool relationship.

Since the beginning of this relationship, humans have always been in full control not only of what the tool should do, but also exactly how it will work, at least in the design and creation phase. The progression to the current status of AI is an evolution of this relationship in two ways.

First, it is an upgrade of the tool's status from the role of a "slave" to that of "agent", giving it agency by having a degree of autonomy with regards to "what" it should be doing. And second, it is a change in that with AI, we no longer entirely decide "how" the tool executes its function. In fact, in many cases, the creators of an AI system cannot entirely describe the criteria that the ML model has used to reach the output. This is known as the "Black Box problem".[5]

Taking these points into consideration, the following can be said about the current state of AI systems:

1. Rather than just executing what the human user wants, AI systems will autonomously come up with predictions, recommendations and decisions.
2. We are not always able to fully understand or explain why an AI/ML system has reached its output.
3. AI/ML output is based on statistical probabilities just like human decision-making – it judges low or high probability of outcomes, it's not some kind of ultimate truth or absolute objective correctness.

We can therefore reach the conclusion that a degree of trust is needed, before the user can hand responsibility over to the AI – and give the autonomous car the steering wheel.[6]

The requirements in building this trust-based relationship vary according to the specific task the AI is supposed to handle. Accepting an AI's recommendation on which movie to watch is much "easier" with a lower trust threshold than, for example, the recommendation on which medicine a doctor should prescribe to their patient.

In a recent survey[7] asking owners of smart voice assistant devices to list the tasks they perform using the device, 84.9 percent reported they use it to set a timer, while only 3.5 percent reported using it to call a cab.

A recent study[8] has shown that when it comes to the application of AI in a business context, 94 percent of business executives understand that AI is essential to business strategy. However, a separate study by MIT Solan[9] found that only 18 percent of companies are widely adopting and understanding AI.

In designing an OSS AI solution that takes critical decisions affecting the performance of an entire network, we discovered that the success of the system depended on more than building more efficient and accurate models and algorithms. We came to the realization that trust is an essential factor in the human-AI interaction, and if we want the users of our AI solutions to accept handing over more critical tasks and decisions, we need to design them to be trustworthy.

[5] Bathaee, Y., 2017
[6] P Anders et al 2018, Muir, B. M. 1987, Siau, K., & Wang, W., 2018
[7] www.experian.com/blogs/insights/2016/09/commerce-is-a-conversation-a-survey-on-amazon-echo-and-voice-assistants/
[8] www.perspectives.eiu.com/sites/default/files/EIU_Microsoft%20-%20Intelligent%20Economies_AI%27s%20transformation%20of%20industries%20and%20society.pdf
[9] www.sloanreview.mit.edu/projects/artificial-intelligence-in-business-gets-real/

What is trust?

The four components of trust in human relationships

Although the concept of trust in human-AI relationships is a new field, we don't have to work from scratch.

The approach of our Experience Design team delete was human-centric. Humans have been trusting other humans since the beginning of our existence, and once human-human trust relationships were established, we started to put our trust in entities and organizations; religions, political parties, banks, schools, business and so on. Our approach was to draw on the formula of trust already functioning in these human-human and human-organization relationships, and use it as a base for building trust in human-AI interaction.

The four main factors that contribute to building trust in another person or entity are: competence, benevolence and openness, integrity and charisma.[10]

A good example to illustrate these four components in action is the process of decision making when hiring an employee that will be responsible for a task in an office.

When dealing with "digital" assistants in the form of AI systems, exactly the same framework of trust applies.

The following pages examine each of these components in turn, from the perspective of human-AI interaction, along with relevant examples of design-related decision and focus areas that can contribute towards creating a trustworthy AI experience.

Figure 2: The components of trust

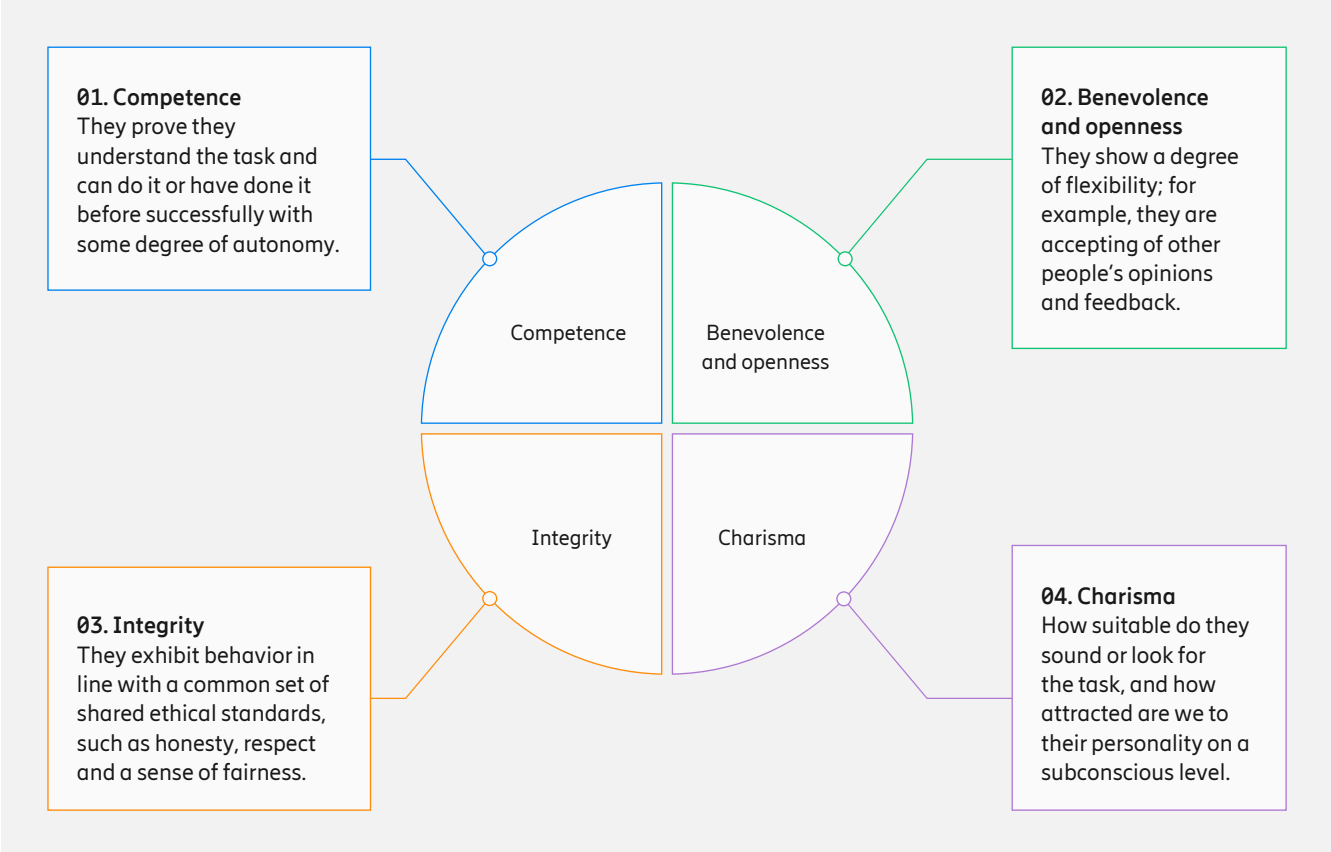

[10] Stern & Coleman, 2015, Dietz and Den Hartog, 2006, Schoorman et al., 2007

Competence

Can you do the job?

In practice within an AI system, the trust component of "competence" essentially means the system is designed to demonstrate that it is capable of fulfilling the user's needs and that it can deliver what it promises.

Here are some practical examples of how UX designers and practitioners can contribute to an AI system's ability to demonstrate competence:

Explainability
Ensuring the system can communicate the reason behind its decisions and its confidence in different results and recommendations in a way that users can easily understand.

Usefulness
Making sure the system is employing AI capabilities to fulfill an actual need or solve a real problem for the users in an effective way.

Trialability
Giving the users the ability to try the AI system or test out its recommendations in a quick, safe and controllable way before they decide to use or approve it.

Demonstration of results
Being able to show evidence that using the AI system has resulted in an improved outcome.

Figure 3: Components of trust – competence

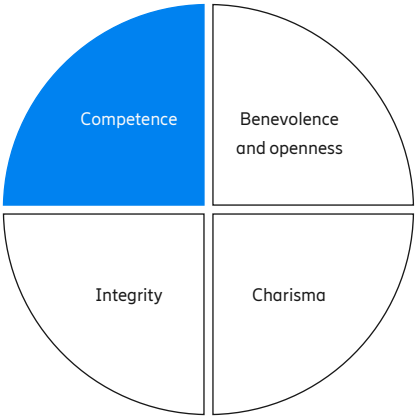

Adoption of AI-powered networks requires knowing they're up to the task.

Network performance diagnostics

Figure 4: Cognitive OSS prototype design demonstrating "Explainability"

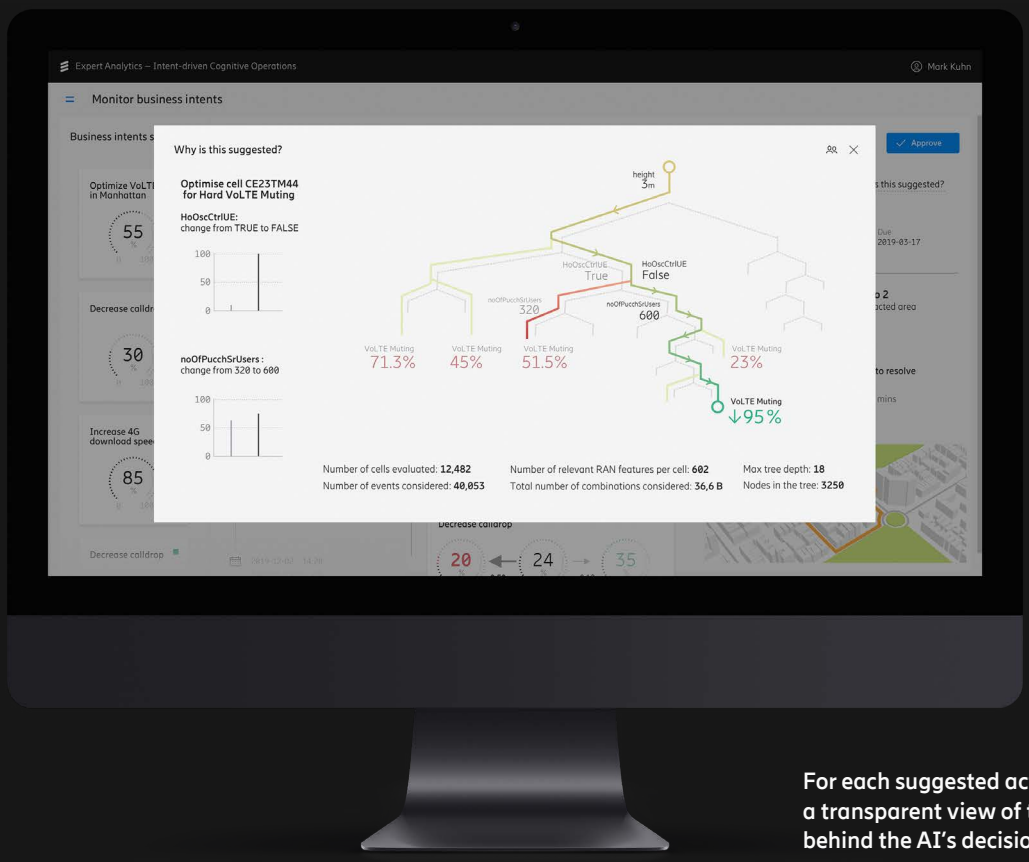

For each suggested action, users have a transparent view of the motivations behind the AI's decisions

Figure 5: Demonstration of results

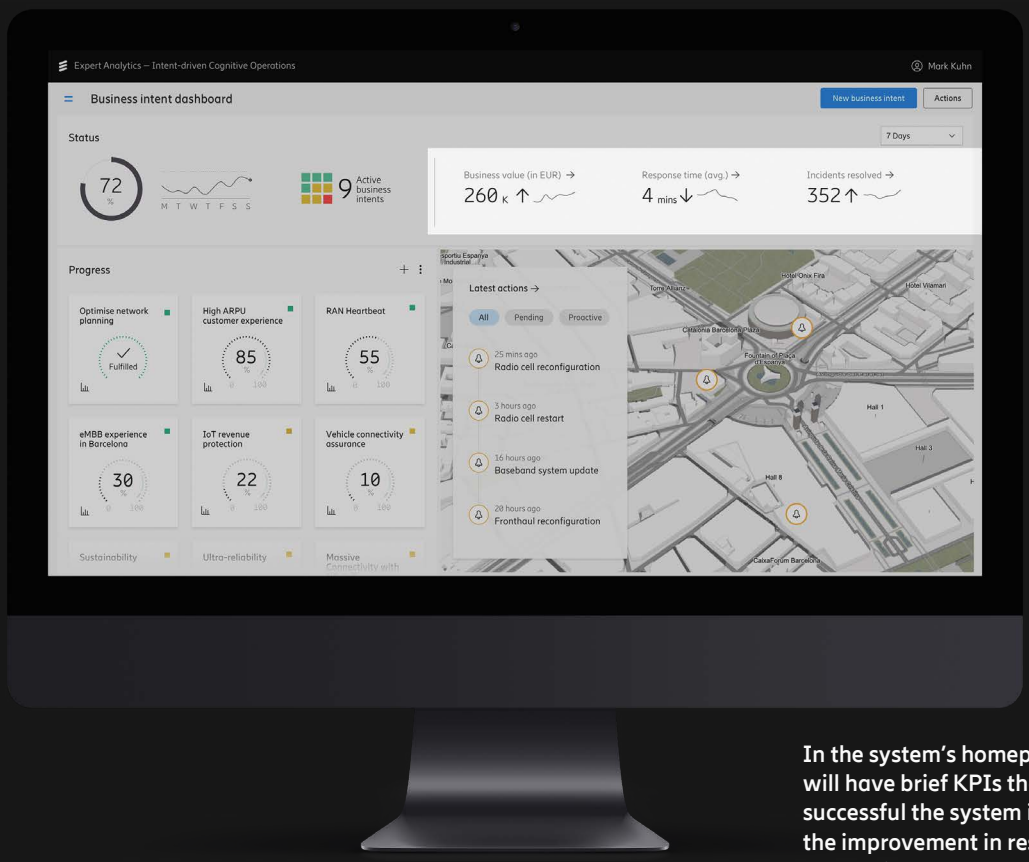

In the system's homepage, users will have brief KPIs that show how successful the system is, for example the improvement in resolved incidents and business value numbers

Benevolence and openness

Are you on my side?

An AI demonstrating "benevolence" can be defined as a system designed to make decisions in the user's best interest, and to communicate the intentions behind decisions to the human user. It should also show flexibility, acceptance of change and new input – exactly as you would expect from a new human colleague. Some practical examples of how UX designers can contribute to the benevolence and openness of an AI system are:

Controllability
Providing an easy way for the user to intervene and change, undo or dismiss an action or decision taken by the AI, as well as the ability to feed their own recommendations into the system.

Adaptability
Making the system flexible and dynamic enough to adapt to the user's explicit or implicit preferences and feedback.

Figure 6: Components of trust – benevolence and openness

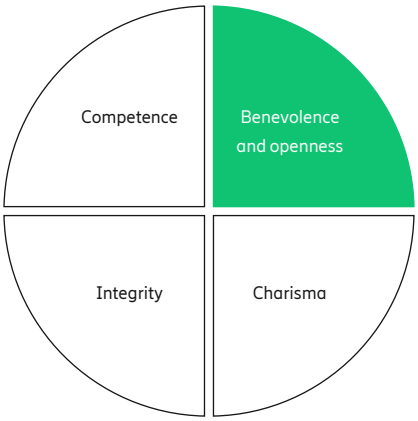

Showing a system is open to influence from the user is a big building block of trust.

Supervising network performance

Figure 7: Controllability, enabling users to take a participatory role in the decision-making process

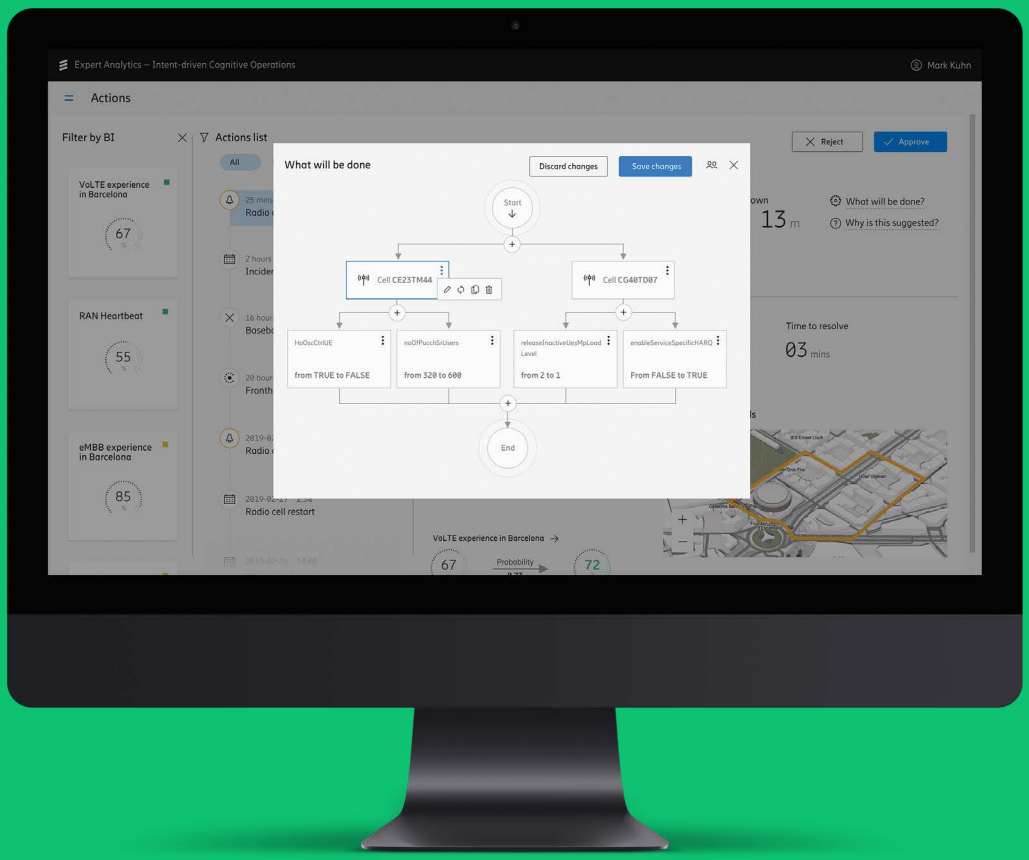

Figure 8: Adaptability, showing the user that their preferences have influence

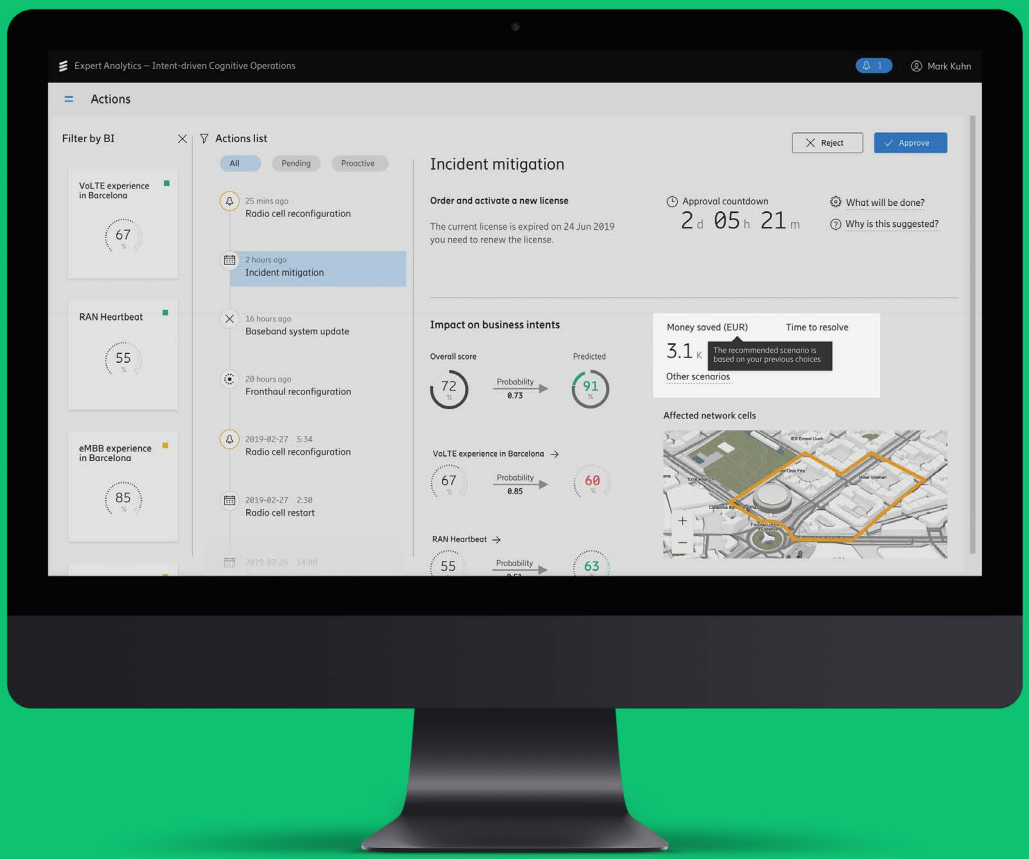

Integrity

Do you share my values?

The concept of integrity in an AI system comes down to whether the user feels that the system is honest, and whether it adheres to the same high ethical standards as the user.

There are two ways UX design in AI contributes to the impression of integrity in a system:

Veracity of promises
Setting the right expectations for the user by clearly communicating the capabilities and limitations of the AI system – knowing what it can promise to do and follow through on and what it cannot or is not designed to do.

Transparency on safety, security and permissions
Making sure the user understands what kind of data is collected, how it is collected, for what reason and how it will be used.

Figure 9: Components of trust – integrity

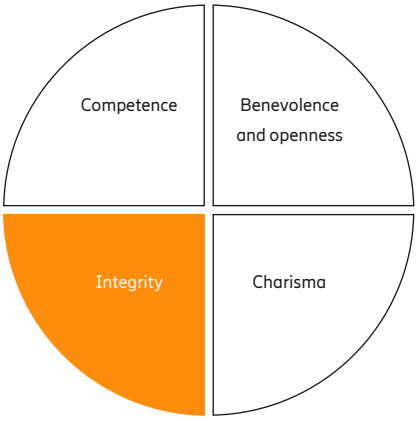

Figure 10: Setting the right expectations and showing the user the different possible outcomes of the AI's recommendations

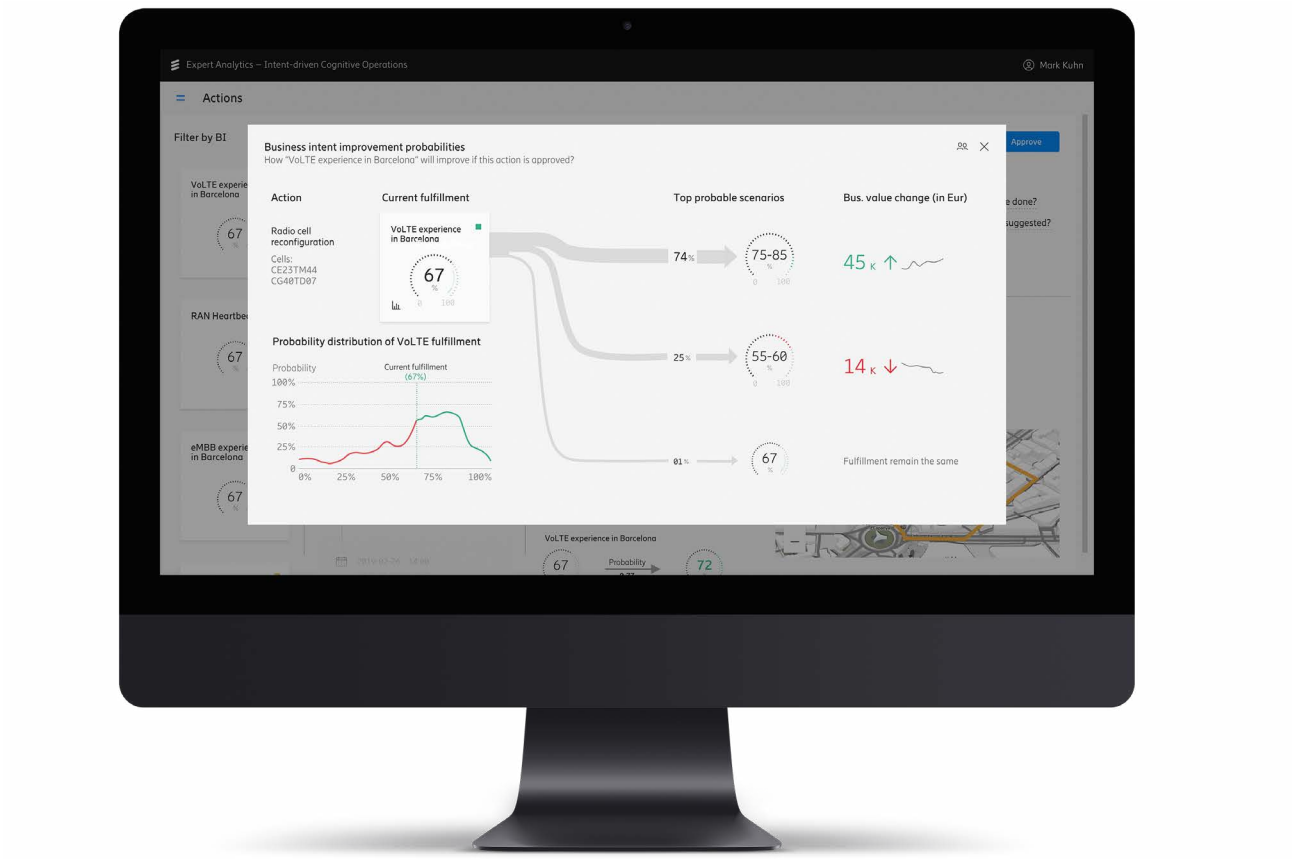

Charisma

Do I like you?

And finally – charisma. Charisma in a AI system comes down to crafting it in a way that gives it general charm and appeal, and that the system looks and sounds appropriate to the task it is handling.

UX designers and practitioners can contribute to the attractiveness of an AI system by implementing:

Visual appeal
Crafting the system's look and feel in an aesthetically pleasing and visually organized way, so that the human user perceives it to be more efficient and understandable.

Tone-of-voice suitability
Making sure that the style and tone of the copywriting and voice interactions are aligned with the message that you want to convey, the desired personality of the system, and the traits of the targeted user group.

Figure 11: Components of trust – charisma

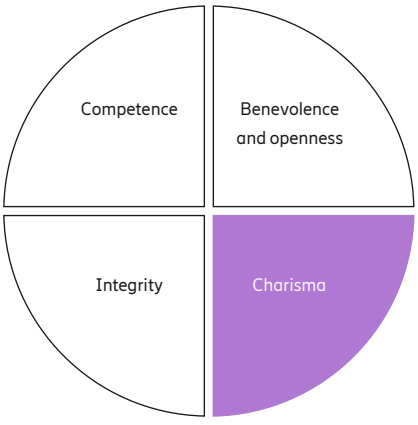

Figure 12: Appropriate tone of voice for the scenario

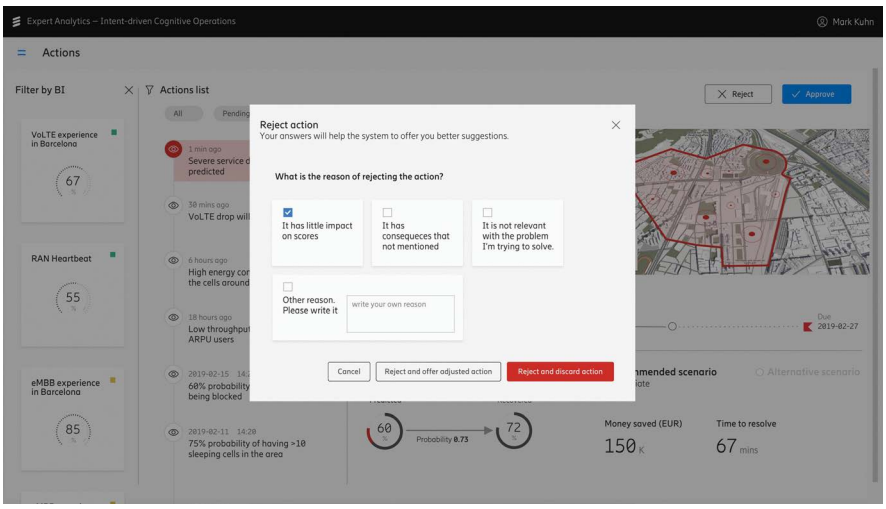

When rejecting an action, the system behaves like a junior engineer and asks about the reason for the rejection, so it can learn from the user's experience

Figure 13: Clear and organized visual layout contribute to increased perceived trustworthiness

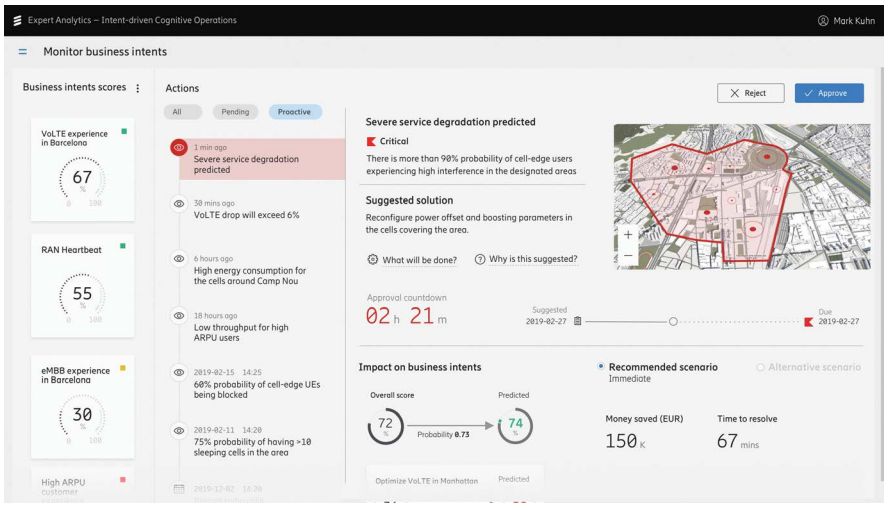

Beyond the building blocks

Other factors that can affect trust in AI

All the elements that have been mentioned so far in this framework are those that are represented in one way or another in the interaction and the interface of the AI system, and are therefore relevant to the UX design.

There are numerous other factors that can affect trust in an AI system, and although these additional factors cannot be translated into the interface as UX, they are nevertheless essential elements in the building of a trust-based relationship. For example:

Accuracy
If the output results, predictions, recommendations and decisions of the system are not accurate to begin with, then the system will not meet the aforementioned criteria needed to satisfy the "competence" component.

Bias in AI
One widely discussed topic in the AI community is the concept of bias – specifically that ML selects incomplete, uninclusive or biased data sets to train the model, whether deliberately or otherwise. This will result in output that is not fair and biased towards a group of users – meaning that the core pillar of "integrity" will not be met.

Laws and ethics
Without sufficient and clear laws and a code of ethics that regulates the relationship between the user and the AI system, for example defining who is responsible if the output of the system affects the user in a negative way, then the trust in the human-AI relationship will not survive any potential mistakes the system makes.

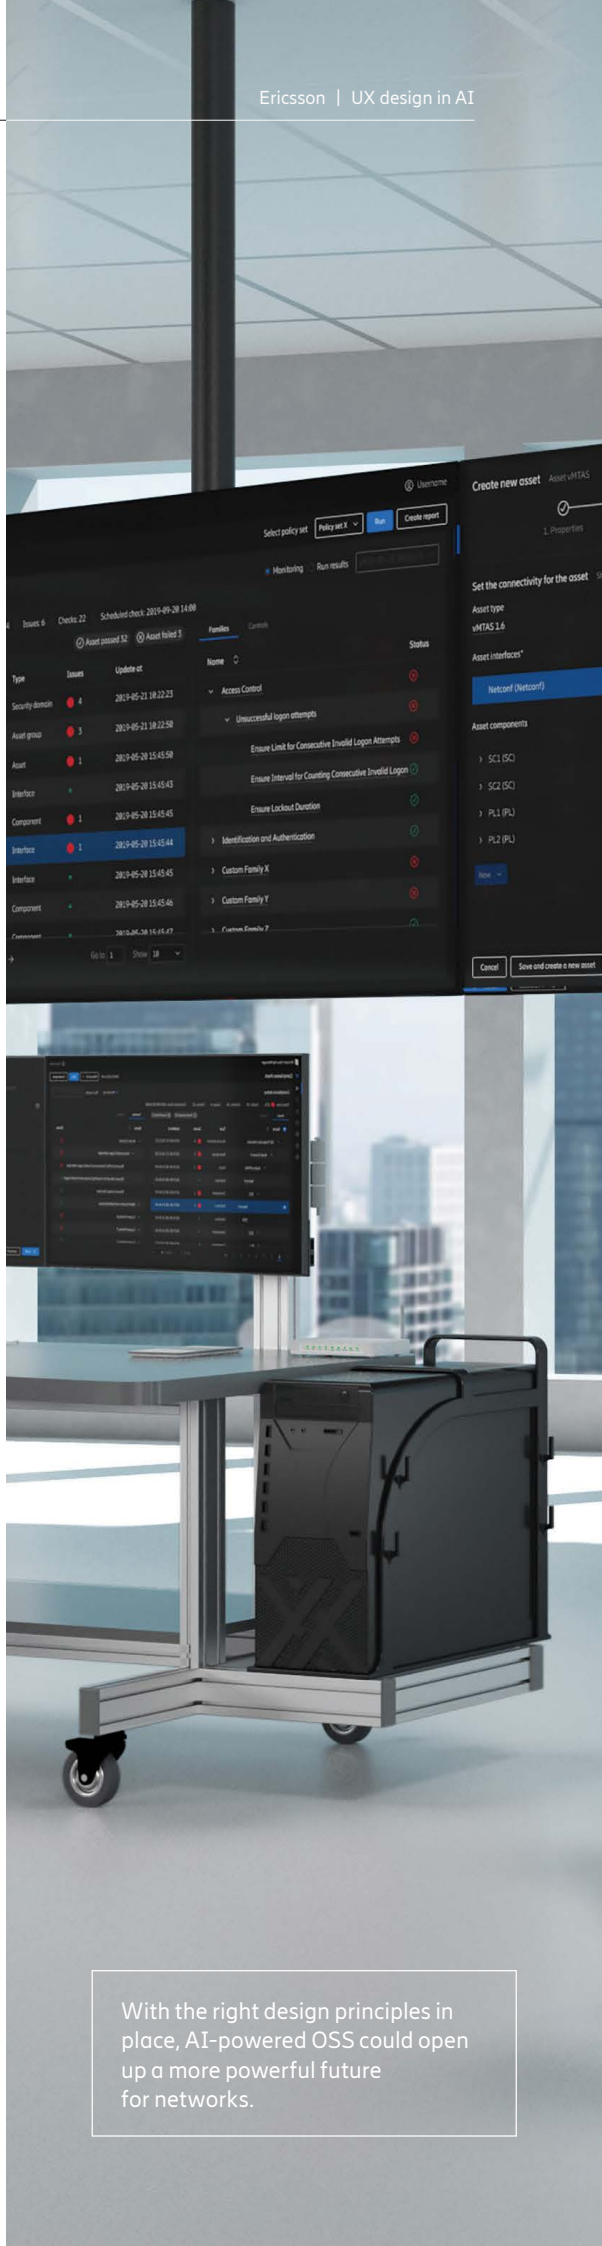

With the right design principles in place, AI-powered OSS could open up a more powerful future for networks.

At the controls of the entire network

Conclusion

The future of human-tool relationships

When introducing AI-powered software like our Cognitive OSS services, the notions of what a network is, what the owner or network operator's role is and what a system provider contributes are changing.

The interaction will be on a much higher and more abstract level. Instead of changing gears in the car, the focus will be on the passenger's journey. Instead of having field technicians manually climbing towers to fine-tune the radio, business operators will collaborate with the AI machine to reach the organization's business intent, impacting the roles of the system providers and the network operator.

A user-centric design process will be even more important when designing AI-powered services than in traditional services. If users and organizations are going to trust the AI-powered system, for example an airplane without a pilot, the trust must be designed into the system and the relationship from the very beginning. Without it, these services will fail.

Designing an OSS AI solution that takes critical decisions that can affect the performance of an entire network is about more than focusing on building better AI models and algorithms. Trust will be the most vital factor in human-AI interaction. If we want the users of our AI solutions to accept handing over more critical tasks and decisions to AI, we need to design them to be trustworthy.

> Designing for trust is a cornerstone of building successful AI systems.

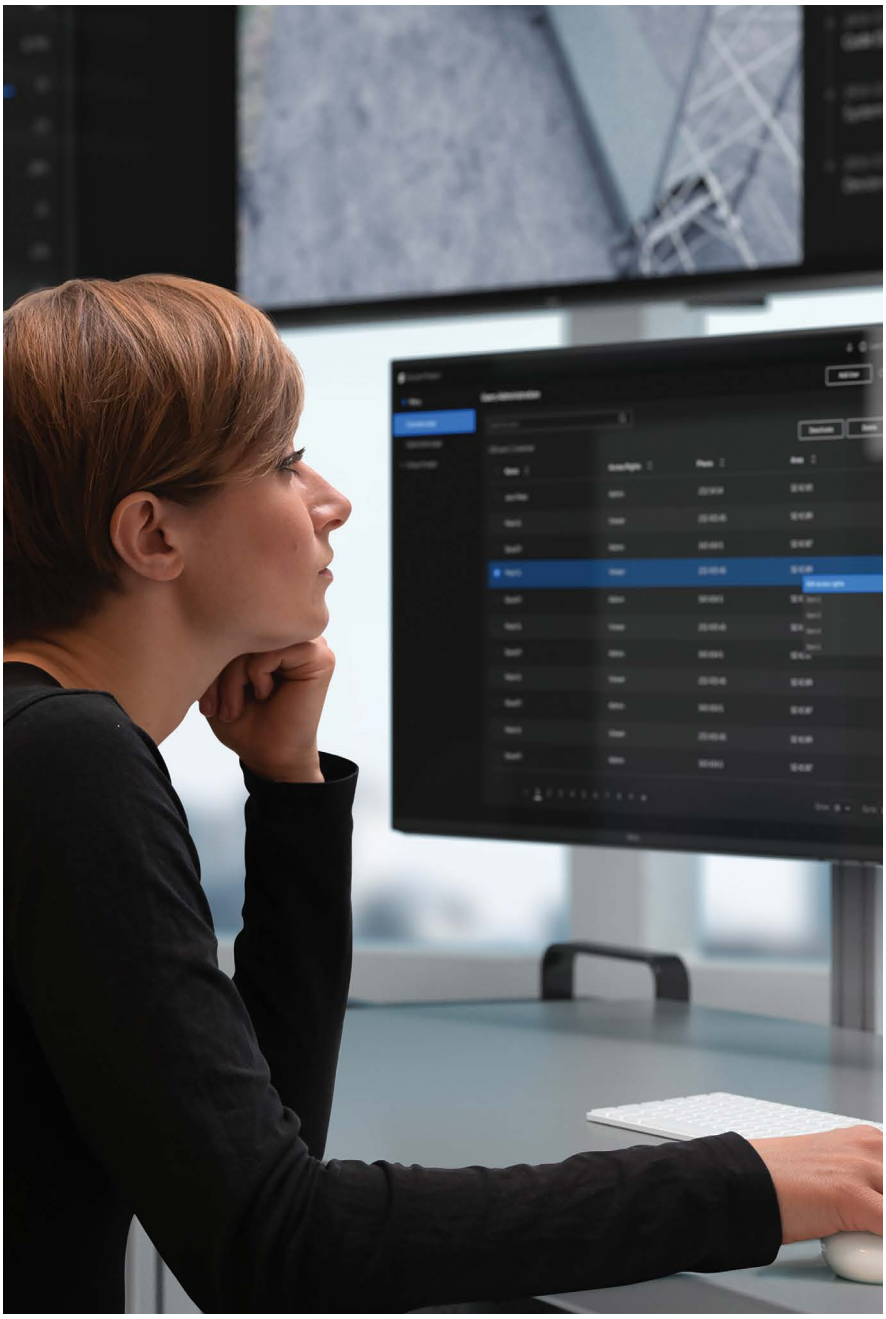

The essential human in the loop

Authors and references

Mikael Eriksson Björling
Mikael is Experience Design Line Manager at Ericsson Experience Design Lab and an Ericsson Evangelist. He was previously Director at the Networked Society Lab. His specialty is in understanding how new behavior, emerging technologies and new industry logics are shaping the future society and in the intersection of these areas design great user and customer experiences. Mikael believes that with the ongoing digital transformation we have a great opportunity to shape a better world. Mikael joined Ericsson in 1998.

Ahmed H. Ali
Ahmed is visual and user experience designer at Ericsson Experience Design Lab. With over 15 years of experience, his career inside and outside Ericsson was focused on designing digital systems that satisfy users' needs and help them achieve their goals by bringing design thinking to the product development process, applying human-computer interaction best practices, and delivering UI/UX concepts and insights. Ahmed joined Ericsson in 2018 and he holds an M.A. in visual design from the University of Hertfordshire in the UK.

- Andras, P., Esterle, L., Guckert, M., Han, T. A., Lewis, P. R., Milanovic, K., ... & Urquhart, N. (2018). Trusting intelligent machines: Deepening trust within socio-technical systems. IEEE Technology and Society Magazine, 37(4), 76-83.
- Bathaee, Y. (2017). The artificial intelligence black box and the failure of intent and causation. Harv. JL & Tech., 31, 889.
- Commerce is a conversation: Survey on Amazon Echo and Voice Assistants – Experian Insights, 2016.
- Gobble, M. M. (2019). The Road to Artificial General Intelligence.
- Goertzel, B. (2007). Artificial general intelligence (Vol. 2). C. Pennachin (Ed.). New York: Springer.
- Intelligent Economies: AI's transformation of industries and society – The Economist Intelligence Unit, and Microsoft, 2018.
- Minsky, M. (1982). Semantic information processing.
- Muir, B. M. (1987). Trust between humans and machines, and the design of decision aids. International journal of man-machine studies, 27(5-6), 527-539.
- Noessel, C. (2017). Designing agentive technology: AI that works for people. Rosenfeld Media.
- Ransbotham, S. et all. Artificial Intelligence In Business Gets Real – Pioneering Companies Aim for AI at Scale. MIT Solan, 2018.
- Sanders, K., Schyns, B., Dietz, G., & Den Hartog, D. N. (2006). Measuring trust inside organisations. Personnel review.
- Schoorman, F. D., Mayer, R. C., & Davis, J. H. (2007). An integrative model of organizational trust: Past, present, and future.
- Siau, K., & Wang, W. (2018). Building trust in artificial intelligence, machine learning, and robotics. Cutter Business Technology Journal, 31(2), 47-53.
- Stern, M. J., & Coleman, K. J. (2015). The multidimensionality of trust: Applications in collaborative natural resource management. Society & Natural Resources, 28(2), 117-132.

Ericsson enables communications service providers to capture the full value of connectivity. The company's portfolio spans Networks, Digital Services, Managed Services, and Emerging Business and is designed to help our customers go digital, increase efficiency and find new revenue streams. Ericsson's investments in innovation have delivered the benefits of telephony and mobile broadband to billions of people around the world. The Ericsson stock is listed on Nasdaq Stockholm and on Nasdaq New York.

www.ericsson.com

Ericsson
SE-164 80 Stockholm, Sweden
Telephone +46 10 719 0000
www.ericsson.com

GFMC-20:000630 Uen