Predicting the Black Swan

Ludic Fallacy and Predictive Self-Healing in Future Cellular Networks

ITU Workshop on "Machine Learning for 5G and beyond"

June 17, 2019

Janne Ali-Tolppa, Henning Sanneck

Nokia Bell Labs Research, Munich, Germany

5G Network Management
Addressing the challenges

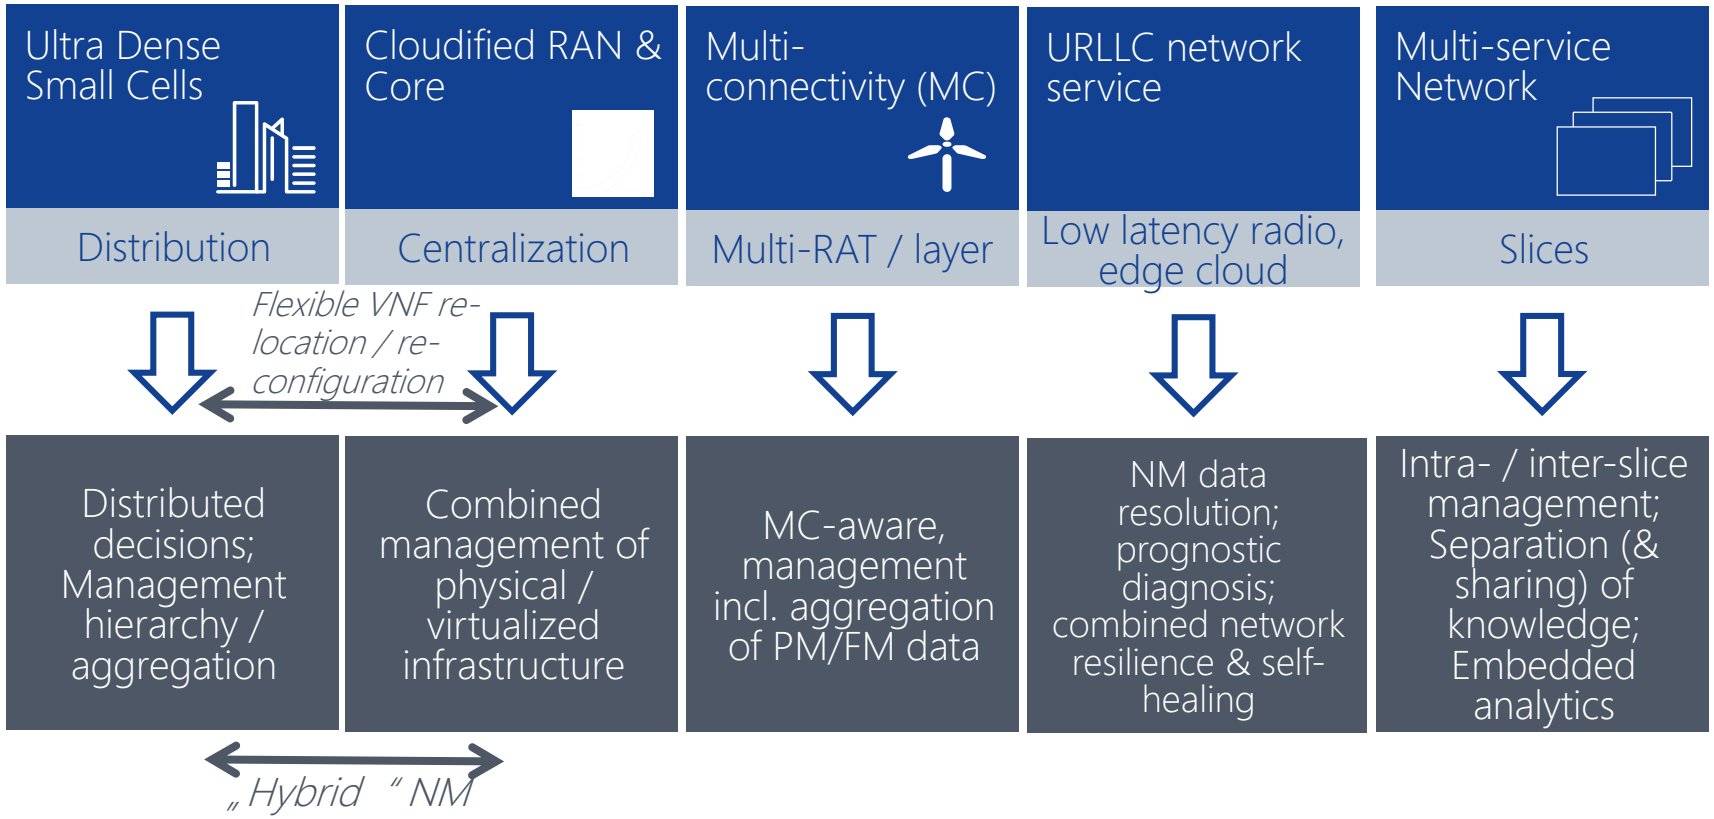

5G Network Management
Addressing the challenges

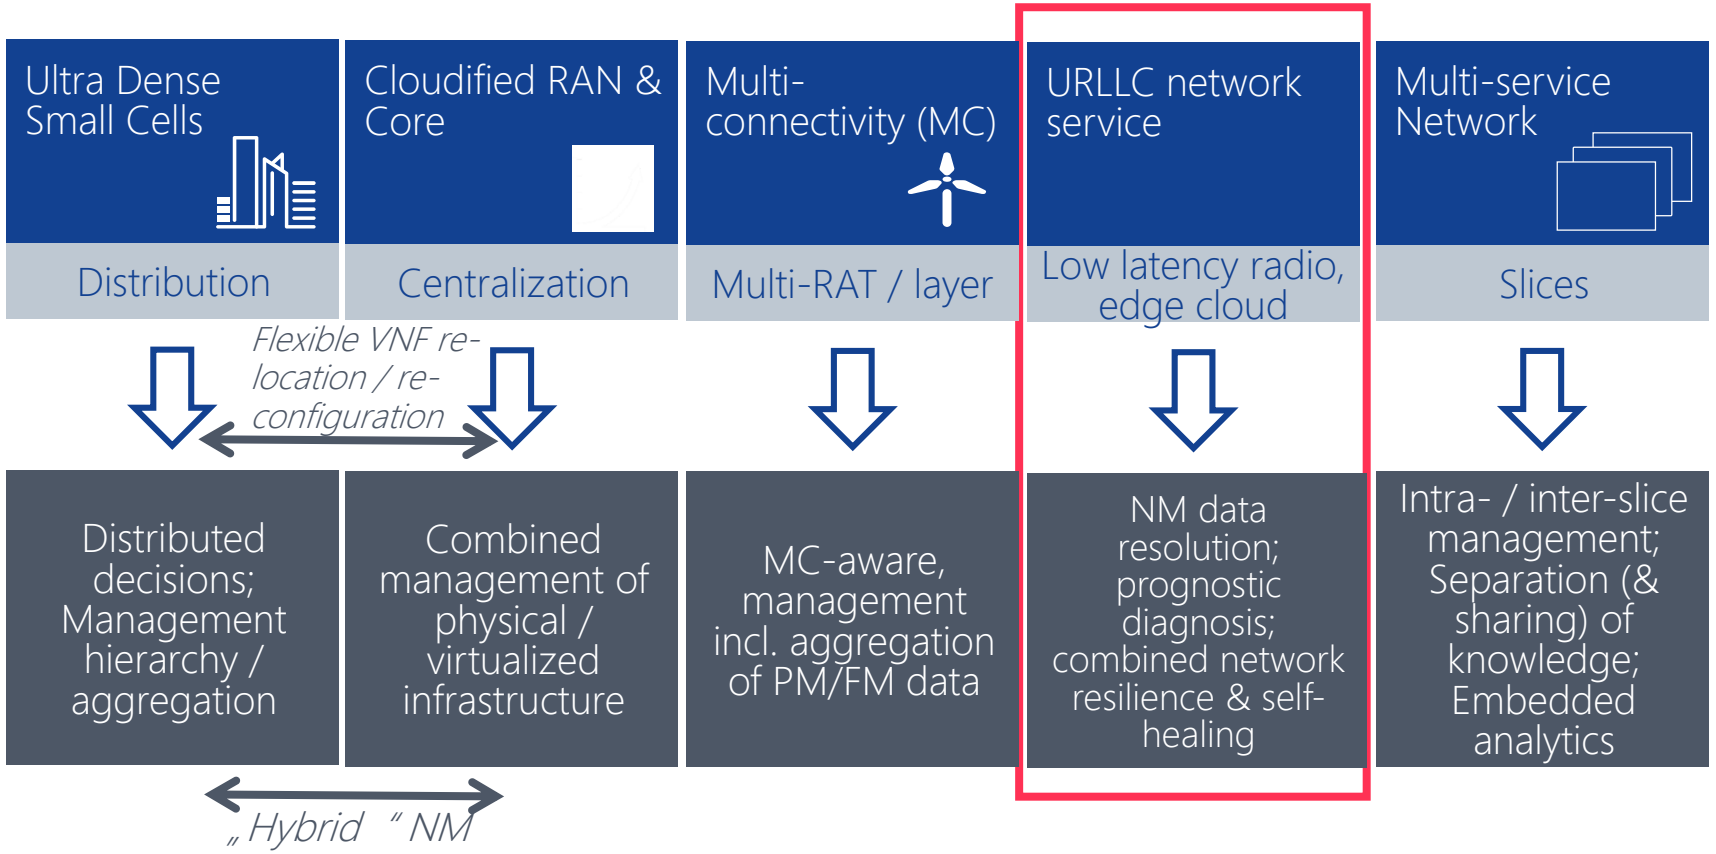

Self-Healing in Future Cellular Networks

Robustness and resilience

Self-Healing in Future Cellular Networks

Robustness and resilience

Robustness

"Capability of performing without failure under a wide range of conditions "

Resilience

"An ability to recover from or **adjust** easily **to** misfortune or **change**"

Merriam-Webster Dictionary

Self-Healing in Future Cellular Networks

Robustness and resilience

Networks are by nature dynamic and complex
→ Unforeseen circumstances are bound to happen

Use cases requiring ultra-high reliability (URLLC)
→ Robustness (redundancy etc.) is no longer alone enough, resilient networks are needed!

Self-Healing in Future Cellular Networks

Design principles for resilience

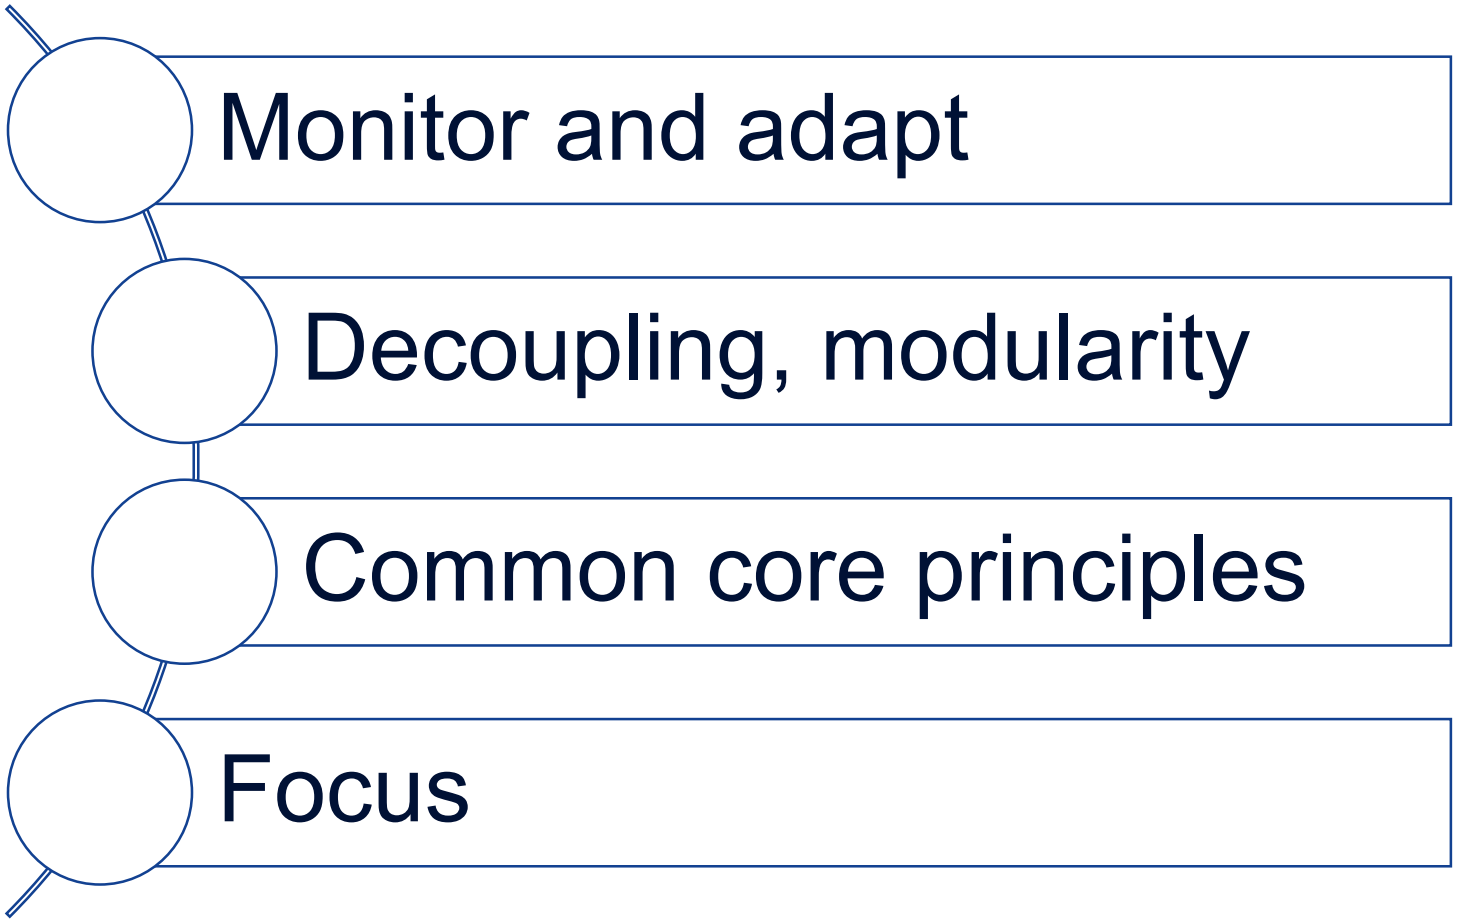

Self-Healing in Future Cellular Networks

Design principles for resilience

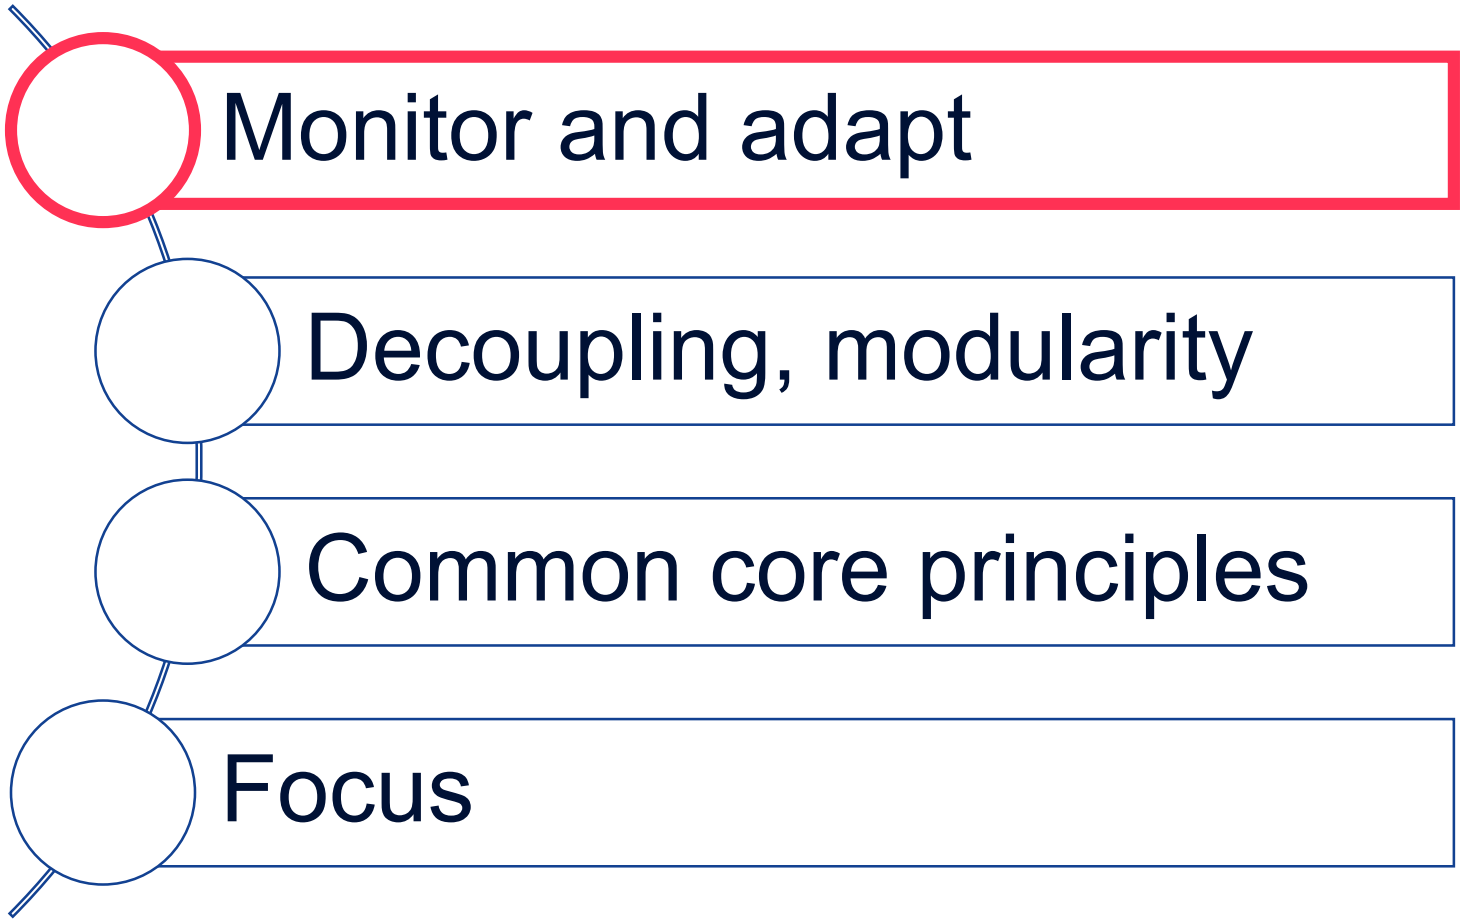

Self-Healing in Future Cellular Networks

Anomaly detection and diagnosis

General process:

1. Profiling – establishing a baseline definition of "expected"
2. Comparing measurements to the baseline, generating an anomaly value
3. Detecting anomalous events, both in time and space
   - An anomaly is something unusual, not necessarily a fault
4. Diagnosing the detected anomalies
   - Comparing the anomaly against a diagnosis knowledgebase of previously diagnosed anomalies (case-based reasoning)
   - Assess the impact and possible root cause
5. Plan and deploy potential corrective actions

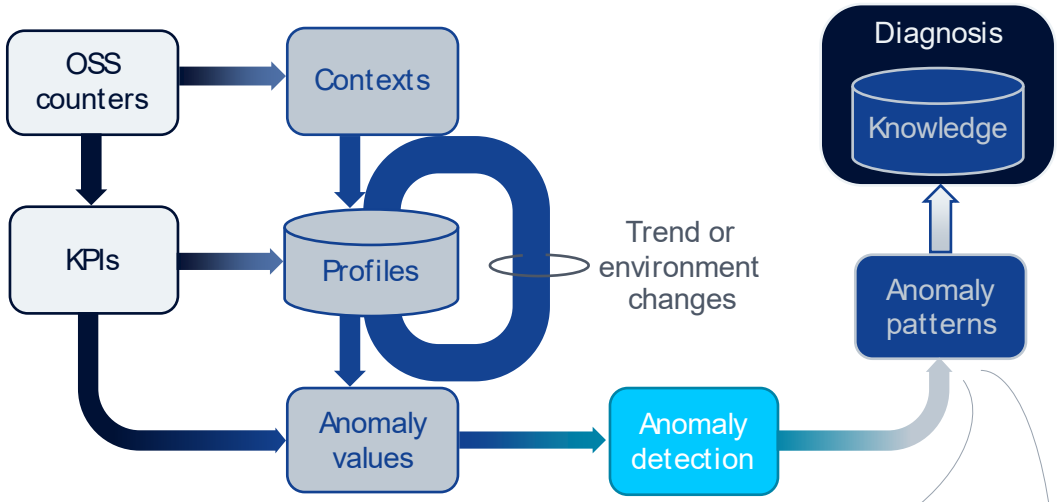

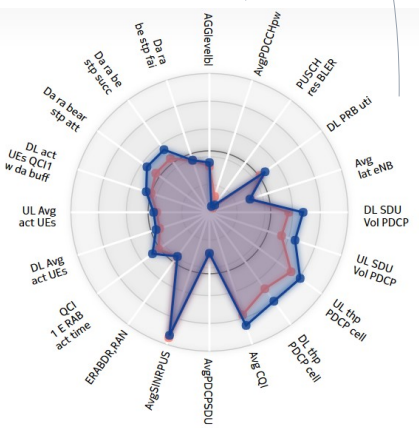

Self-Healing in Future Cellular Networks

Predicting network failures

Self-Healing in Future Cellular Networks

What about predicting faults?

- Of course prediction is still detection, i.e. detecting the early signs of a problem
- In addition to learning to detect and diagnose fault patterns, try to learn also to detect early signs of a degradation, if possible
- Including the temporal context, learning the sequence of events that led to a state labeled as degraded
- Different levels of prediction
  - Predict a previously diagnosed and labeled fault
  - Predict an anomalous state, which is not previously diagnosed (and may or may not be a fault)
  - Predict an anomaly, which is not well-represented in the current network element state model

Self-Healing in Future Cellular Networks

What/how to predict from the data?

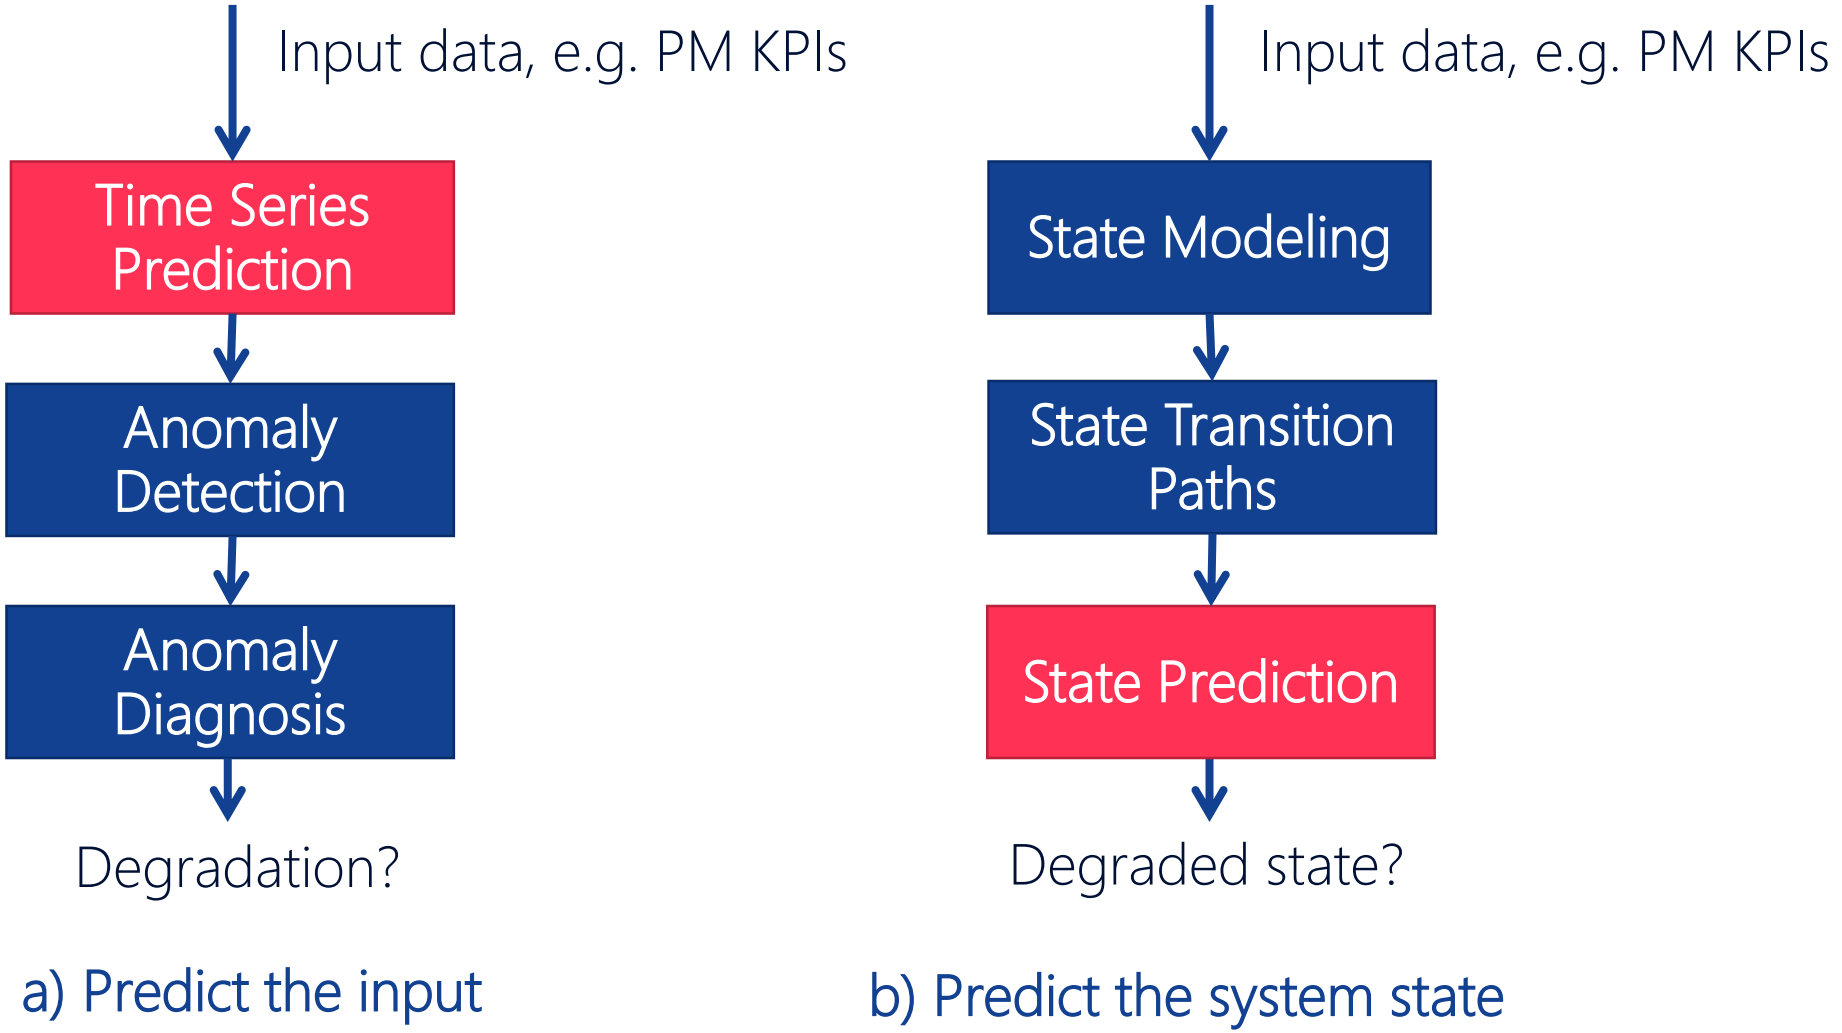

a) Predict the input

b) Predict the system state

Self-Healing in Future Cellular Networks
Algorithms (example)

- ARIMA
- Triple Exponential Smoothing / Holt-Winters
- Recurrent Neural Networks
  - Long-Short Term Memory (LSTM)
- Convolutional Neural Networks
- Sparse auto-encoders (for state modeling)

Can we predict also the more complicated and even previously unencountered faults?

NOKIA Bell Labs

Self-Healing in Future Cellular Networks

Three Aspects of the Black Swan

1. It is an outlier, as it lies outside the realm of regular expectations, because nothing in the past can convincingly point to its possibility.
2. It carries an extreme 'impact’.
3. In spite of its outlier status, human nature makes us concoct explanations for its occurrence after the fact, making it explainable and predictable

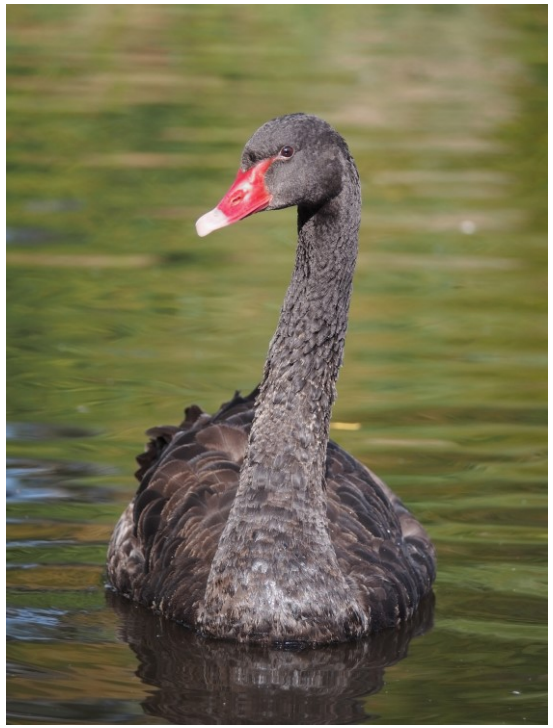

Photograph by Francis C. Franklin / CC-BY-SA-3.0

[1] Nassim Nicholas Taleb, *The Black Swan*, Random House, 2007

Self-Healing in Future Cellular Networks
A Ludic Fallacy

An argument against applying simplified, game-like statistical models in complex domains [1]:

- It is impossible to have the entirety of information available
- Even the smallest variations can have very significant effects
- Because of this, what we tend to see, post hoc, as a failure of the predictive capability of a model, may very well be a Black Swan, something that was real, but unpredictable (or *predictable only after the fact*)

[1] N. N. Taleb, *The Black Swan*, Random House, 2007

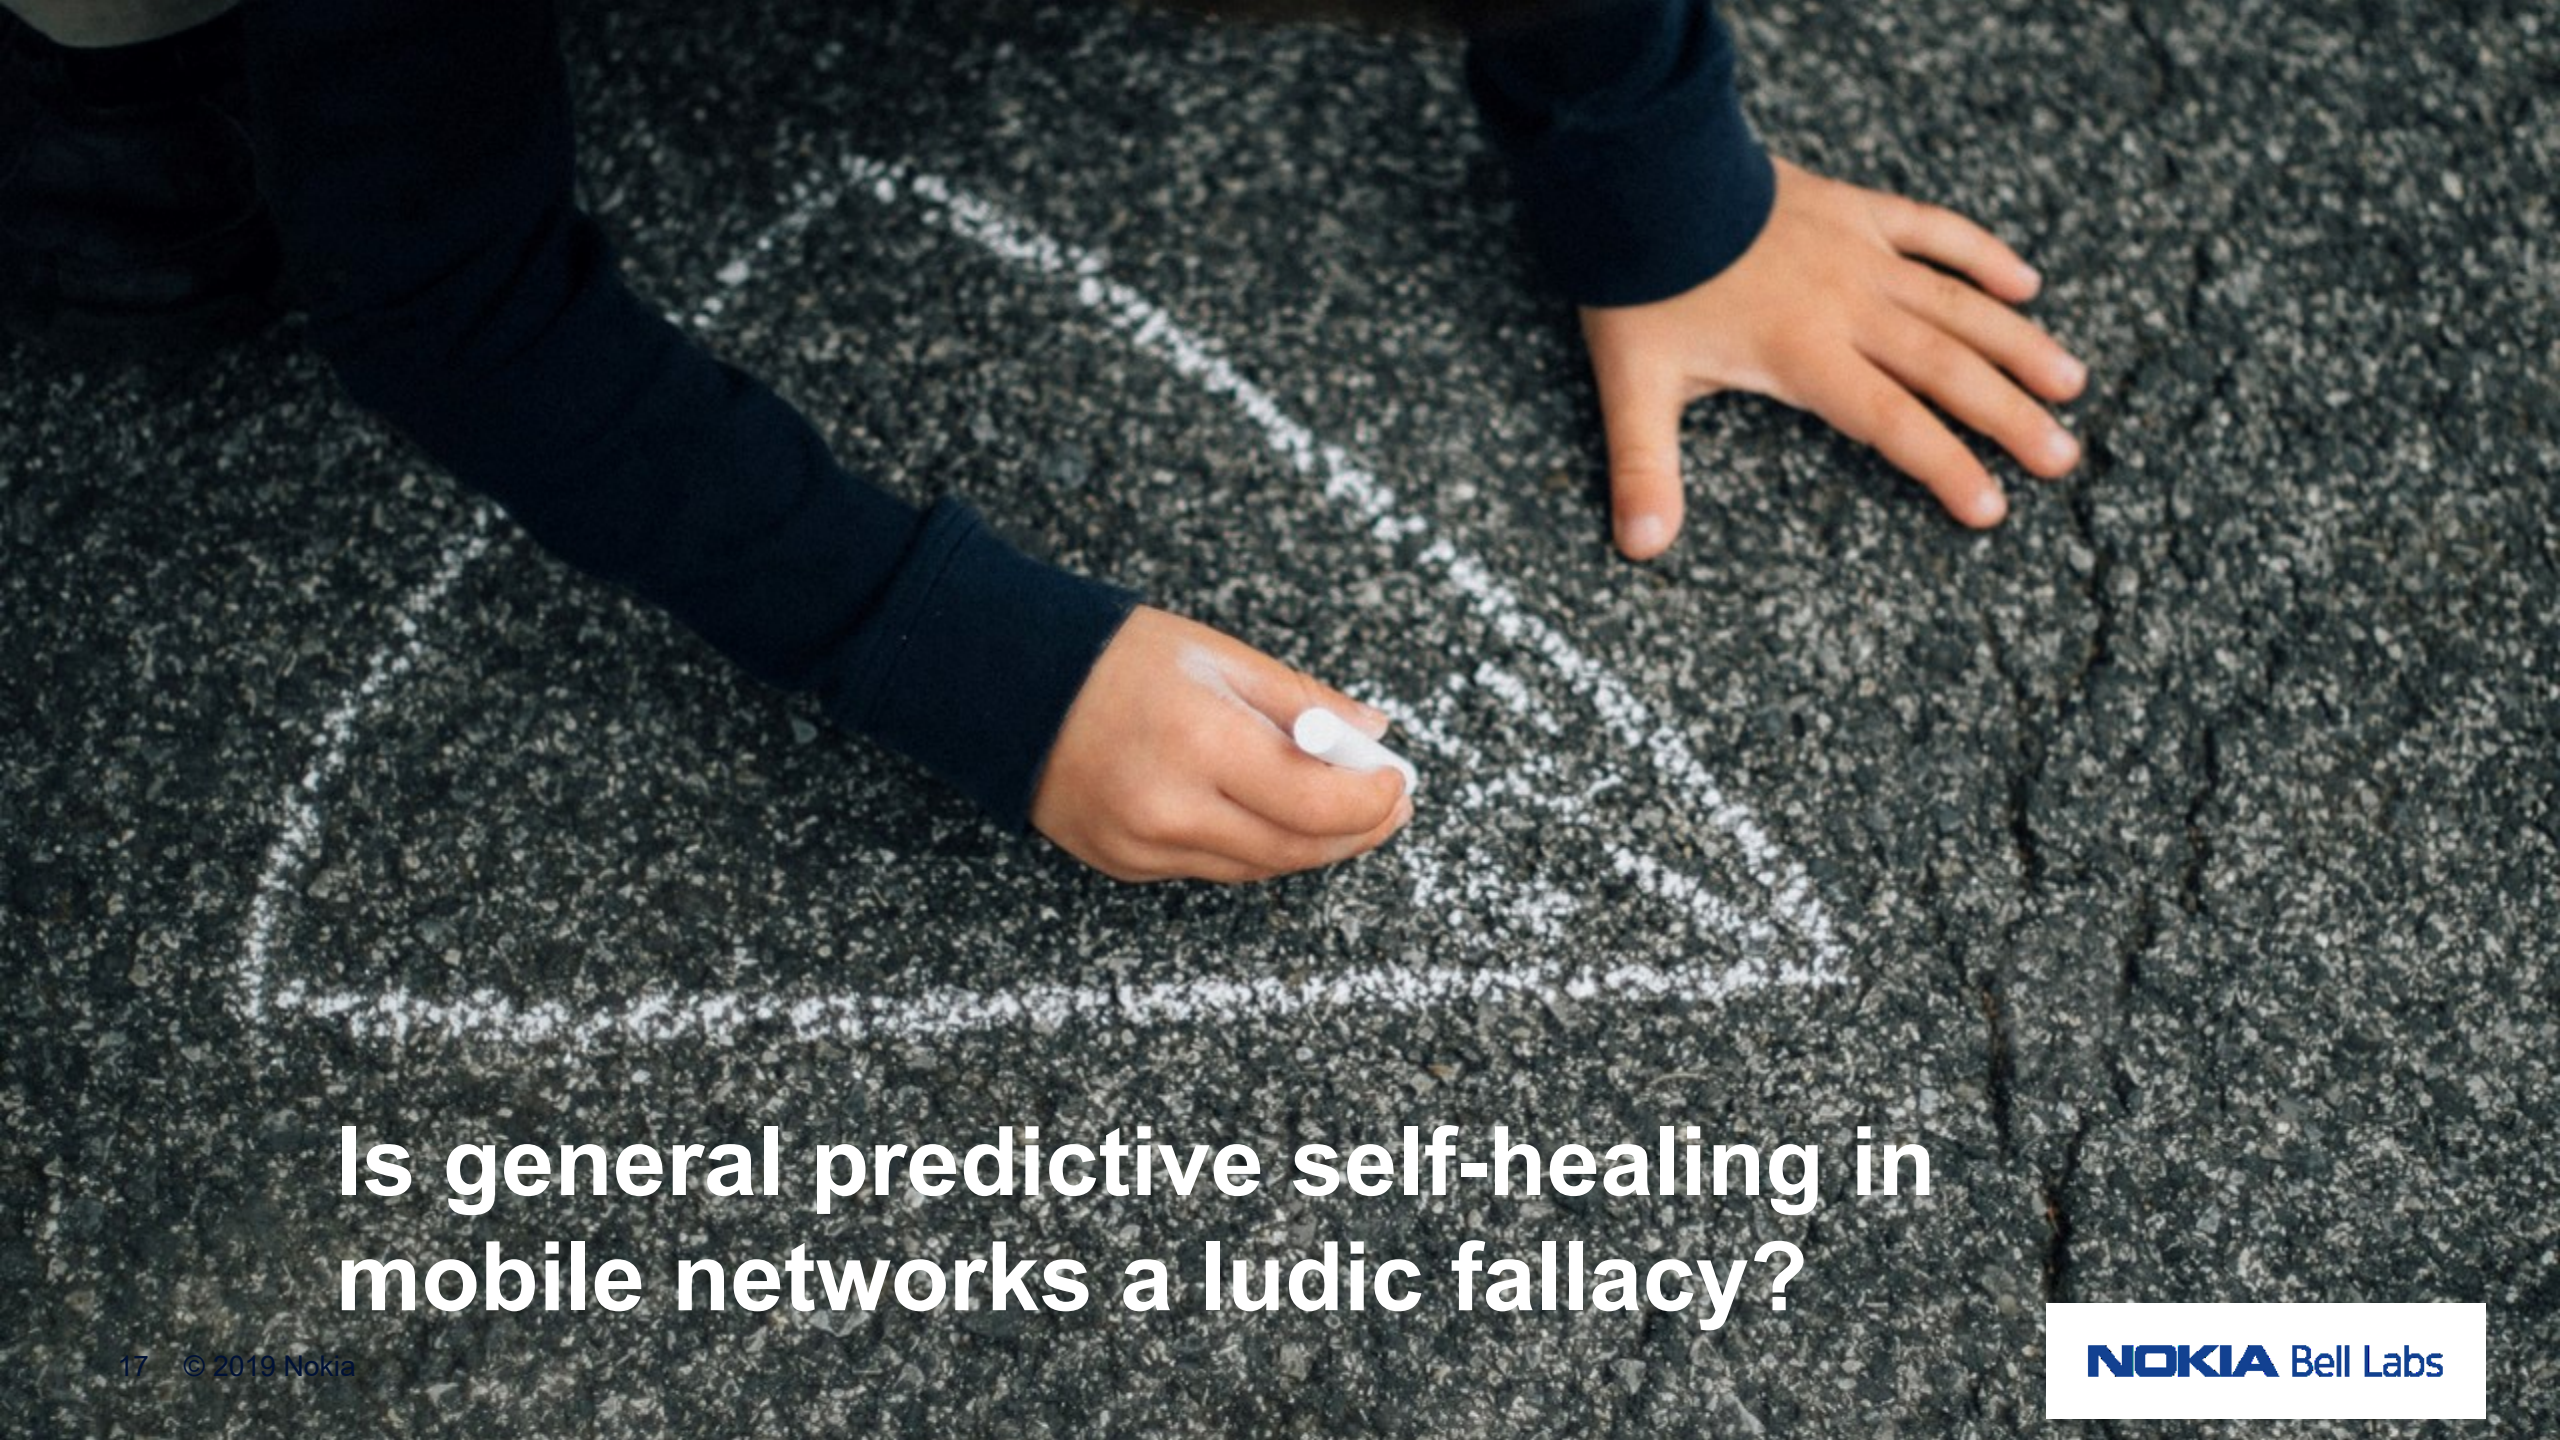

Is general predictive self-healing in mobile networks a ludic fallacy?

NOKIA Bell Labs

Self-Healing in Future Cellular Networks

What can you predict? Some examples?

- Software problem (e.g. a memory leak)?
- Hardware problem (e.g. a power amplifier failure)?
- (Human-induced) misconfiguration?
- Environmental impacts (building, weather)?
- Unexpected network traffic patterns?

- Often failures are a combination of several root causes, each of which alone would not have led to a failure
- What can be detected, depends on the available data

Self-Healing in Future Cellular Networks

What information do we have available?

Self-Healing in Future Cellular Networks
What data do we have available?

- Network management data: CM, FM, PM, UE measurements
- Cloud resource KPIs
- Logs
- Traces
- Context data
  - Weather
  - (Road) traffic
  - Calendar with special events
  - Production system data in an Industry 4.0 private networks
  - Etc.
- Service-based architecture, ML-pipelines etc. enable better utilization of all available data

Self-Healing in Future Cellular Networks

Collecting knowledge – transfer learning

Self-Healing in Future Cellular Networks
Collecting knowledge

- Diagnoses for case-based reasoning are collected in a diagnosis knowledgebase
- The knowledgebase or the predictive model can also contain knowledge of the state transition paths that led to a fault
- The main challenge is that the faults are rare occasions
- Collecting sufficient amount of analyzed samples is difficult for diagnosis and for prediction the dimensionality is even higher

Self-Healing in Future Cellular Networks
Transfer Learning

Transfer of *knowledge*

between *different domains*

with *different learning tasks*

- A ***domain*** $\mathcal{D}$ consists of a feature space $\mathcal{X}$ and a marginal probability distribution $P(X)$, with $X \in \mathcal{X}$.
- A ***task*** $\mathcal{T}$ consists of a label space $\mathcal{Y}$ and an objective predictive function $f(\cdot)$, which is not observed but can be learned from the training data
- The ***goal****:* Given a source domain $\mathcal{D}_S$ and learning task $\mathcal{T}_S$, a target domain $\mathcal{D}_T$ and learning task $\mathcal{T}_T$, transfer learning aims to help improve the learning of the target predictive function $f_T(\cdot)$ in $\mathcal{D}_T$ using the knowledge in $\mathcal{D}_S$ and $\mathcal{T}_S$, where – in general, but not necessarily – $\mathcal{D}_S \neq \mathcal{D}_T$ or $\mathcal{T}_S \neq \mathcal{T}_T$.

Self-Healing in Future Cellular Networks
Diagnosis Cloud

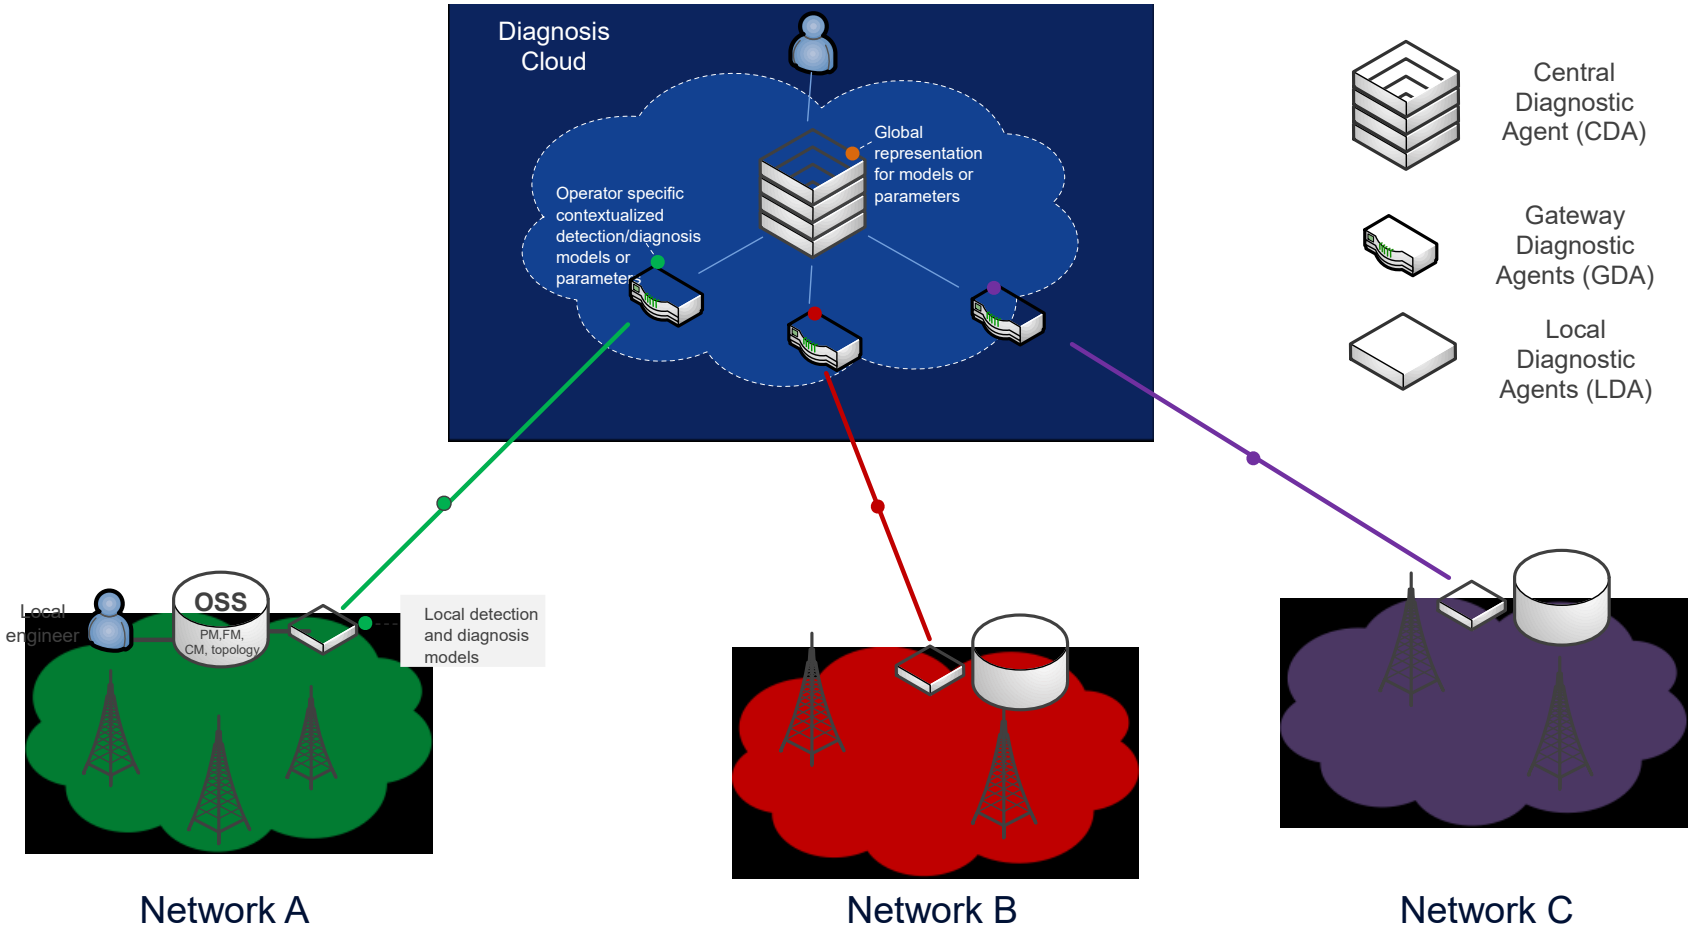

Ciaocarlie et al: Diagnosis cloud: Sharing knowledge across cellular networks

Self-Healing in Future Cellular Networks

Inherent Resilience

Self-Healing in Future Cellular Networks
Inherent resilience

- Self-healing functions are most effective, when used in conjunction was other resilient design principles
- Mobile networks typically have distributed and decoupled by nature and have built-in redundancy
- This makes them less susceptible for catastrophic failures due to minor changes

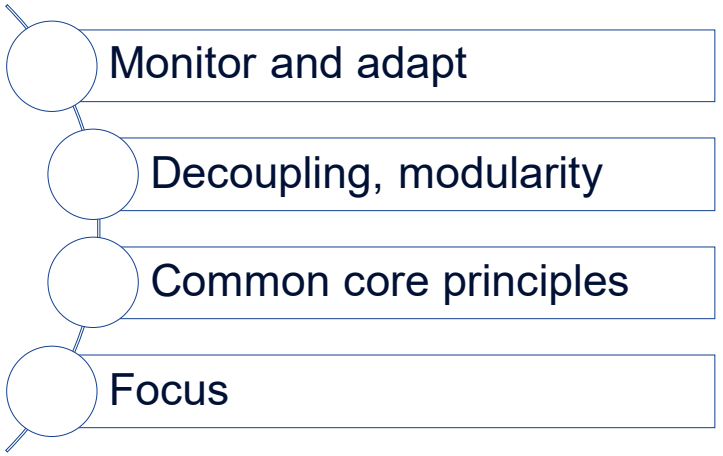

Self-Healing in Future Cellular Networks

A practical example

Self-Healing in Future Cellular Networks

5G Slice Analytics & Diagnostics

Hamburg Seaport Testbed

- Hamburg seaport testbed: evaluate concepts in a controlled environment (ground truth available; closed-loop automation possible)
- Results
  - Slice-aware Network Element (NE) state model: quantization of NE KPIs into a selected number of states
  - States -> Long-Short Term Memory (LSTM) Recurrent Neural Network (RNN) -> State Prediction
  - Additionally Mobility Pattern Prediction (MPP) on the barges moving in the harbor area
  - Of 6220 sequences in a validation set, we were able to predict 97.6% of the low-SINR events with 1.4% false positive rate of the total number of sequences

Predicted states (T+3):

Original states:

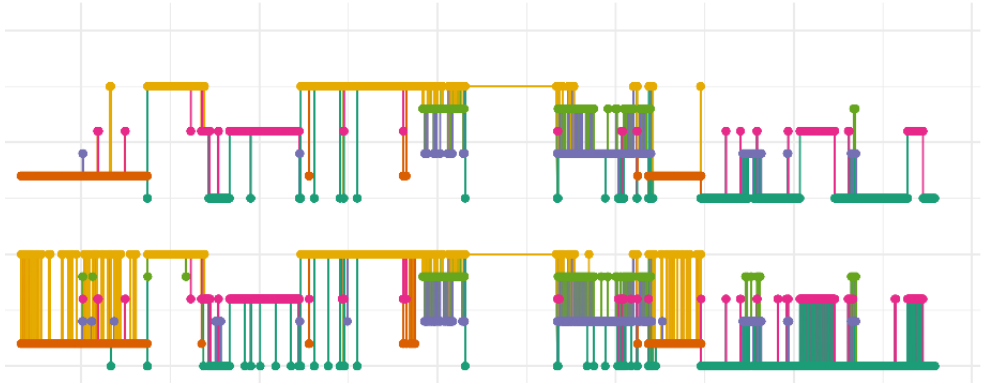

Public

NOKIA Bell Labs

Self-Healing in Future Cellular Networks

Summary

Self-Healing in Future Cellular Networks
Black Swan

1. An outlier, an anomaly, unpredictable
   - We can at least detect them with anomaly detection
2. Extreme impact, especially in the sense of a surprise
   - The more automated diagnosis knowledge we collect, the more prepared we are for a wide variety of issues

   -> Less high-impact surprises and we are prepared to react quickly to correct or mitigate any problems
3. With points 1. and 2. covered or at least mitigated, the human bias for hindsight may also be manageable ☺

Self-Healing in Future Cellular Networks

Is predictive mobile network self-healing a ludic fallacy?

- It is impossible to have the entirety of information available
  - Correct, but by instrumenting the system much more than before and by including also data sources for context data, a lot of the required information is available
  - Methods like transfer learning can help to collect knowledge for diagnosis
  - Already ongoing in 5G and beyond
- Even the smallest variations can have very significant effects
  - By following the design principles of redundancy, decoupling and modularity, the networks can be made inherently resilient -> small variations are unlikely to have catastrophic impacts
- In any system with such complexity as in mobile networks, there are and will always unavoidably be “black swans”, i.e. completely unpredictable major events
- Still, we argue that predictive self-healing is **not** a ludic fallacy. By proper data and knowledge collection, advanced self-healing methods can significantly improve the resiliency of future mobile networks.
- To enable this, the work done on defining ML-friendly architecture and enablers, e.g. for collecting and sharing data and information, are critically important.

NOKIA Bell Labs

Self-Healing in Future Cellular Networks
Automatic network element state definition

Clustering

Clustering methods group observations into clusters, and try to represent all observations in a group with a single example, a **prototype**. The prototype is usually selected through a minimization of an error measure.

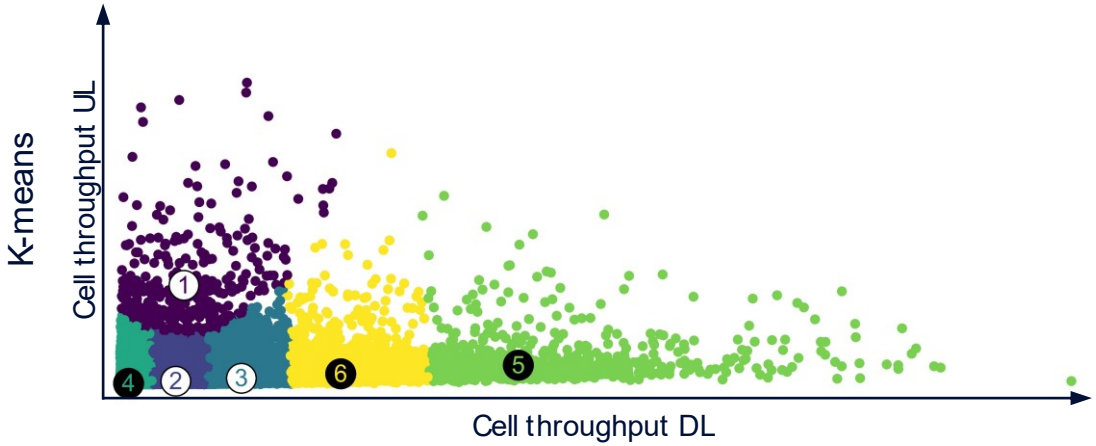

Bounding-Sphere Quantization

Previous work [2] involved automated state definition for anomaly detection and visualization uses. For these cases, a quantization algorithm was developed called Bounding-Sphere Quantization. BSQ strives for a quantization where the maximum error is kept small.

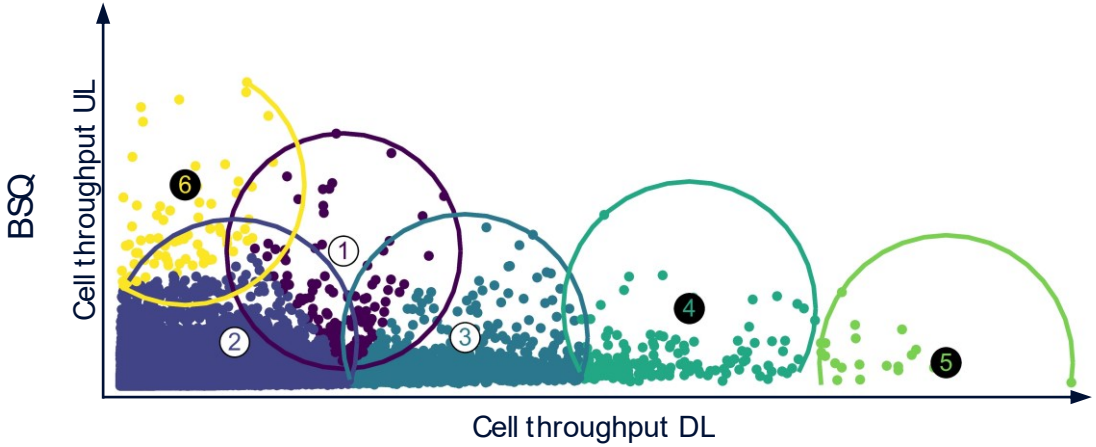

[2] https://ieeexplore.ieee.org/document/8406263

Self-Healing in Future Cellular Networks

Network element state prediction

- Advantage is that the high dimensional network management data is reduced to a series of discrete states
- Well-suited for Recurrent Neural Networks (RNNs)
- The performance depends highly on the state model
- Tradeoff between short-term and longer-term accuracy

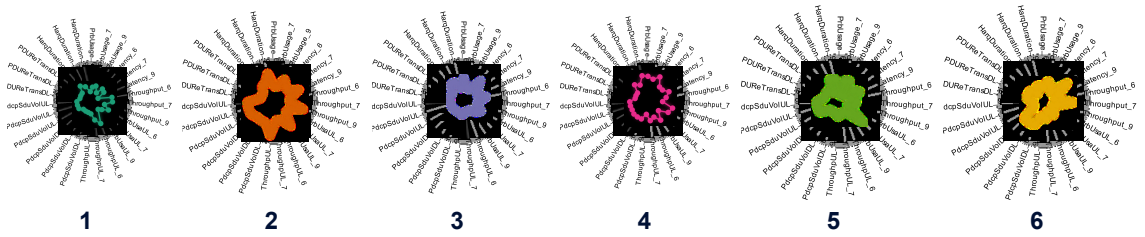

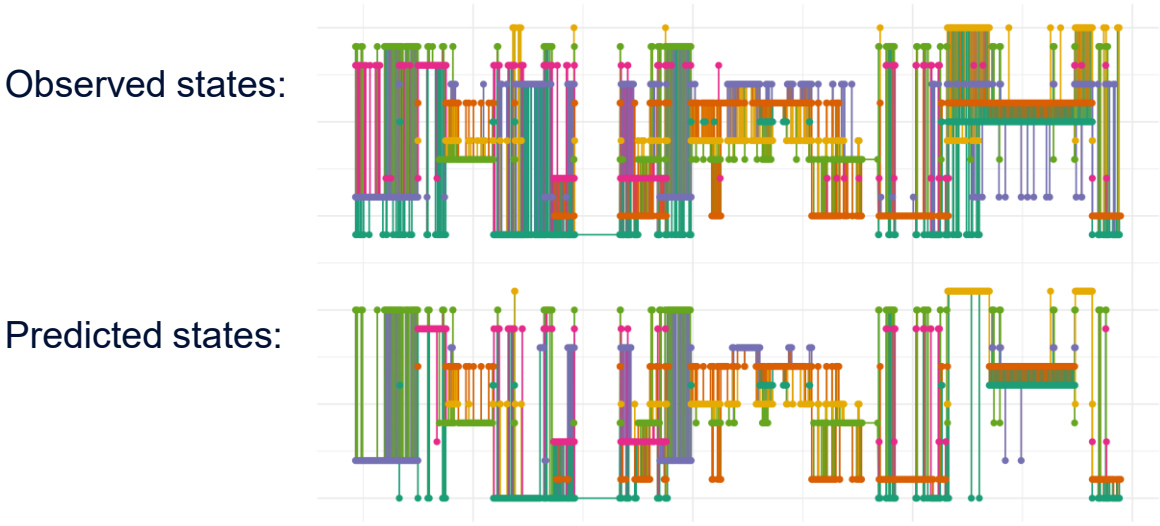

Self-Healing in Future Cellular Networks
Holistic optimization and self-healing

- In a complex system, improving the resilience of only one part or level of the system can sometimes (unintentionally) introduce fragility in another
- To improve the resilience, it is often necessary to work in more than one domain, scale and time granularity at a time
- If there are several optimization and self-healing functions, possibly operating at different scopes, their actions may need to be coordinated
    - E.g. on UE-level vs. on NE-level