Nonparametric Distributed Detection Using One-Sample Anderson-Darling Test and *p*-value Fusion

Aalto University
School of Electrical Engineering

Topi Halme and Visa Koivunen

Introduction

- We address the problem of statistical inference and decision making in a network of independent, low-bandwidth sensors
  - IoT, Sensor Networks, radar, radio spectrum sensing, environmental surveillance, cyber-physical systems
- It is not feasible to specify accurate probability models for each sensor in large scale applications
- Data driven approach: empirical models and a fully nonparametric inference
- Distributions are learned from data: each sensor can adjust to its operational environment

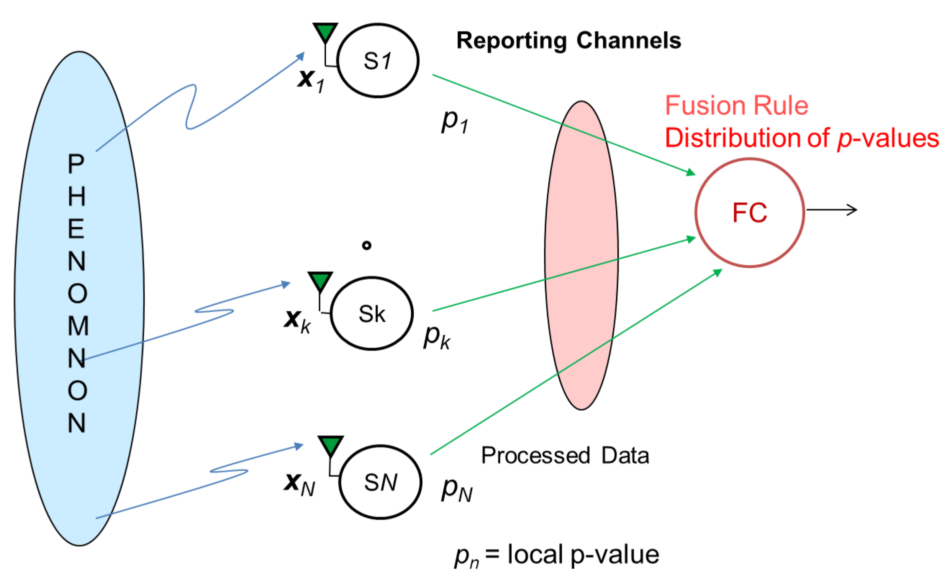

A parallel topology, where the observers only have a one-way communication channel with a global Fusion Center is assumed

Contributions

- Fully nonparametric distributed detection approach
- Underlying probability models approximated by empirical distributions
- Distributions of the test statistics are learned from the data by bootstrapping at each sensor
- The local test compares the observed data to the learned distribution under null hypothesis using the one sample Anderson-Darling (AD) test
- Each sensor sends its *p*-value from a local AD test to the FC that that makes final decision
- FC performs a test on the distribution of *p*-values
- Concentration towards small *p*-values: the evidence for rejecting the null hypothesis is strong
- Fisher's Chi-square, Stouffer's Z-score and Tippett's minimum value tests considered at FC
- Strict control of the error levels in decision making

System Model

- Hypothesis testing problem where a training data set X from unspecified distribution F is collected in each sensor under H_0 conditions
- F and test statistics distributions are learned from data
- Distribution F is compared to observation data Y generated from distribution G, and a local test is performed with following hypotheses:
  - H_0: Both X and Y obey the same distribution $F = G$
  - H_1: X and Y obey different distributions $F \neq G$

One-Sample Anderson-Darling Test

- In most applications, one affords to collect lot of training data, and empirical approximation $\hat{F}$ for F using bootstrapping is thus accurate
- One-Sample version of AD test comparing Y with $\hat{F}$ is more accurate than two-sample approach of contrasting X and Y
- Tradeoff is higher computational demand, mainly in local offline computation in training phase
- The one-sample test consistently outperforms the two-sample test independent of the underlying distributions

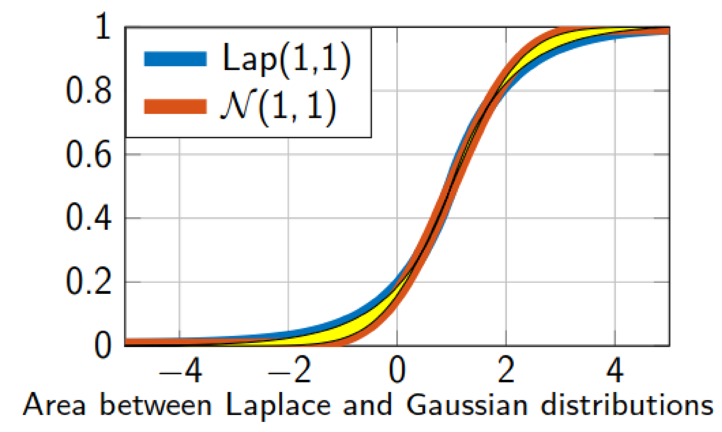

Area between null hypothesis distribution and sample distribution

Weighting function $F(x)(1-F(x))$ that emphasizes deviations at the tails

$$\tau_{AD} = n\int_{-\infty}^{\infty}[F(x) - G_Y(x)]^2\varphi(x)\,dF(x)$$

AD test statistic τ is the weighted area between distributions

Learning the Test Statistic Distribution

- When working only with empirical distributions, the distribution of τ under null cannot be derived analytically
- Bootstrapping provides an accurate approximation

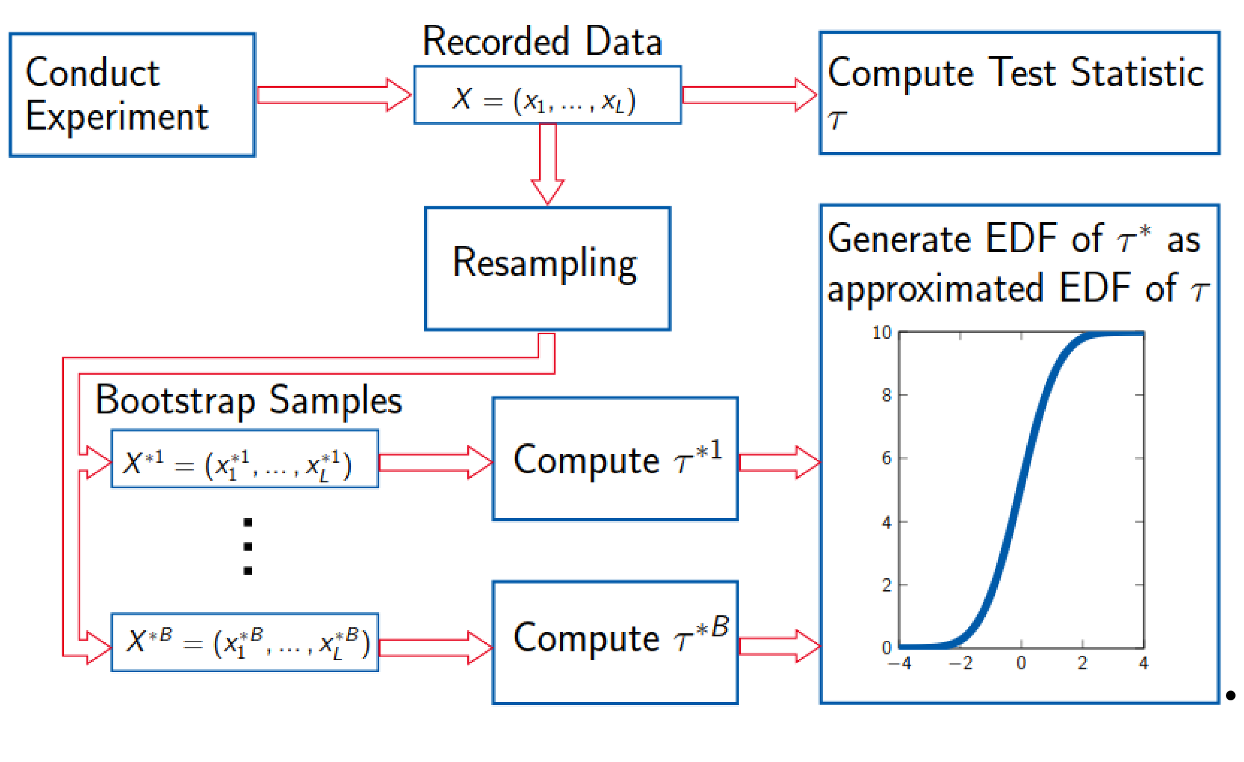

- A local p-value is conveniently obtained as the proportion of bootstrapped samples more extreme than obtained τ

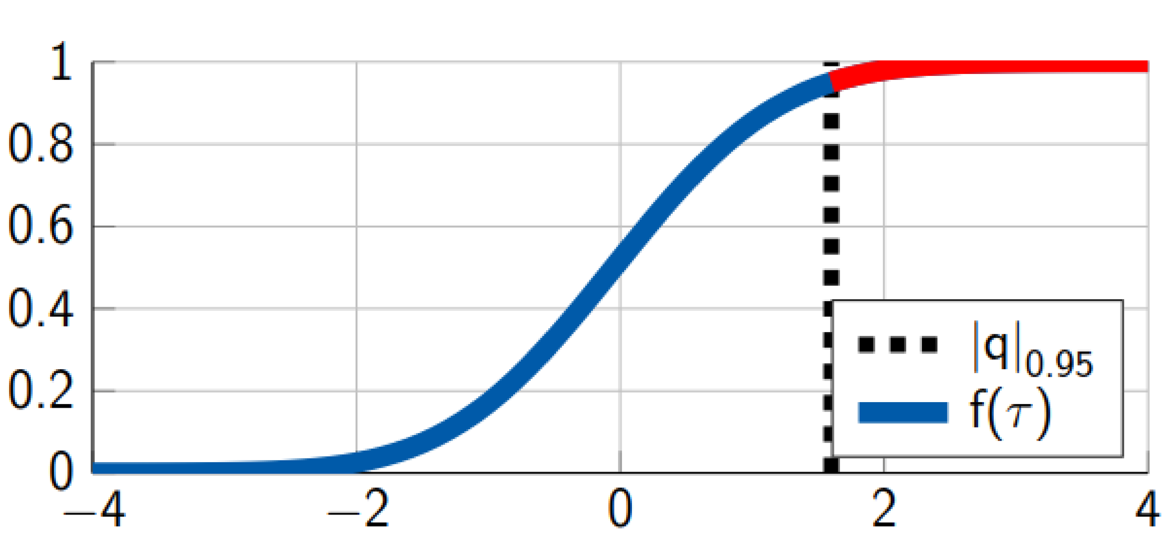

Fusion of Local Tests

- In a continuous sample space, **the *p*-value is uniformly distributed between 0 and 1 if the null hypothesis is true**
- Comparison of *p*-value fusion methods shows that for our purposes methods evaluating the distribution of local *p*-values are most efficient
- Fisher's Chi-square, Stouffer's Z-score and Tippett's minimum value tests employed at FC

1000 local p-values produced from simulating AD-test with X = Normal, Y = Laplace and X = Normal, Y = Normal

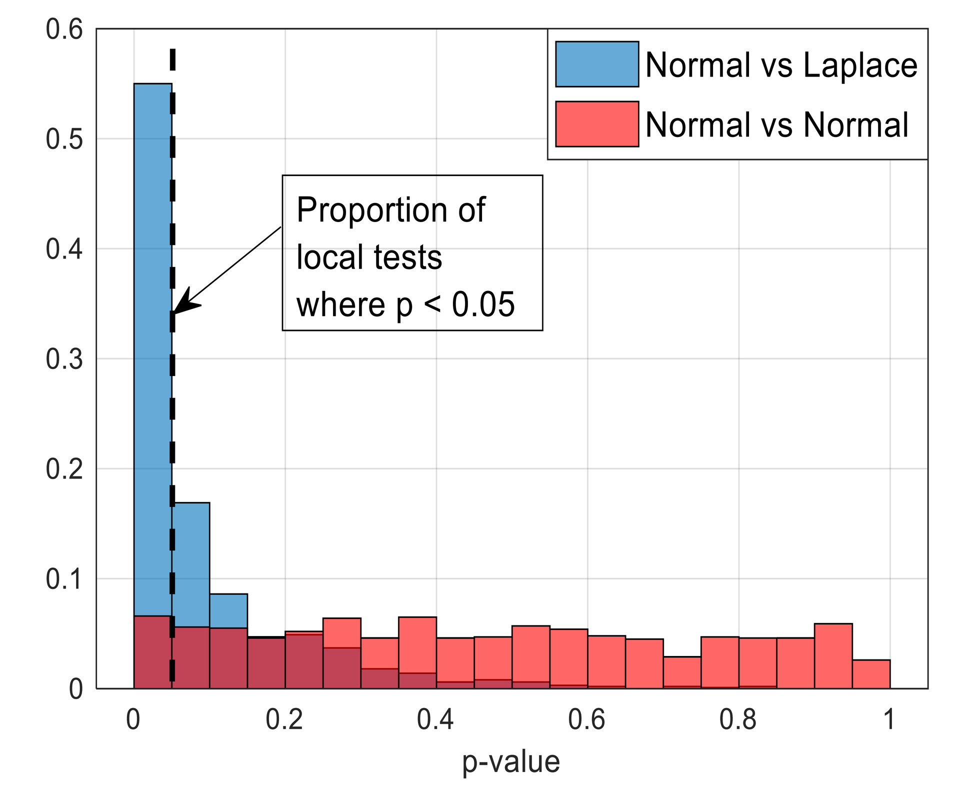

- Fisher's and Stouffer's methods exploit this property to construct transformations that follow a well known distribution under H_0

Fisher	Stouffer
$-2\sum_{i=1}^{N}\log p_i \sim \mathrm{X}^2_{2N}$	$\frac{\sum_{i=1}^{N}\Phi^{-1}(1-p_i)}{\sqrt{N}} \sim N(0,1)$

- The methods have near linear relationship, but performance differences still exist, with Fisher's method being more accurate in our simulations
- Both methods outperform most optimal binary fusion rules (Chair-Varshney)
- The methods look at the distribution of *p*-values, for large scale sensor networks: a few outliers will not determine the outcome. Robustness to malfunctioning sensors.
- Reliability of local sensor test is still important
- Secure communication of *p*-values to FC is under development

ROC graph of row 2 in Table 3: normal distributions with differing means

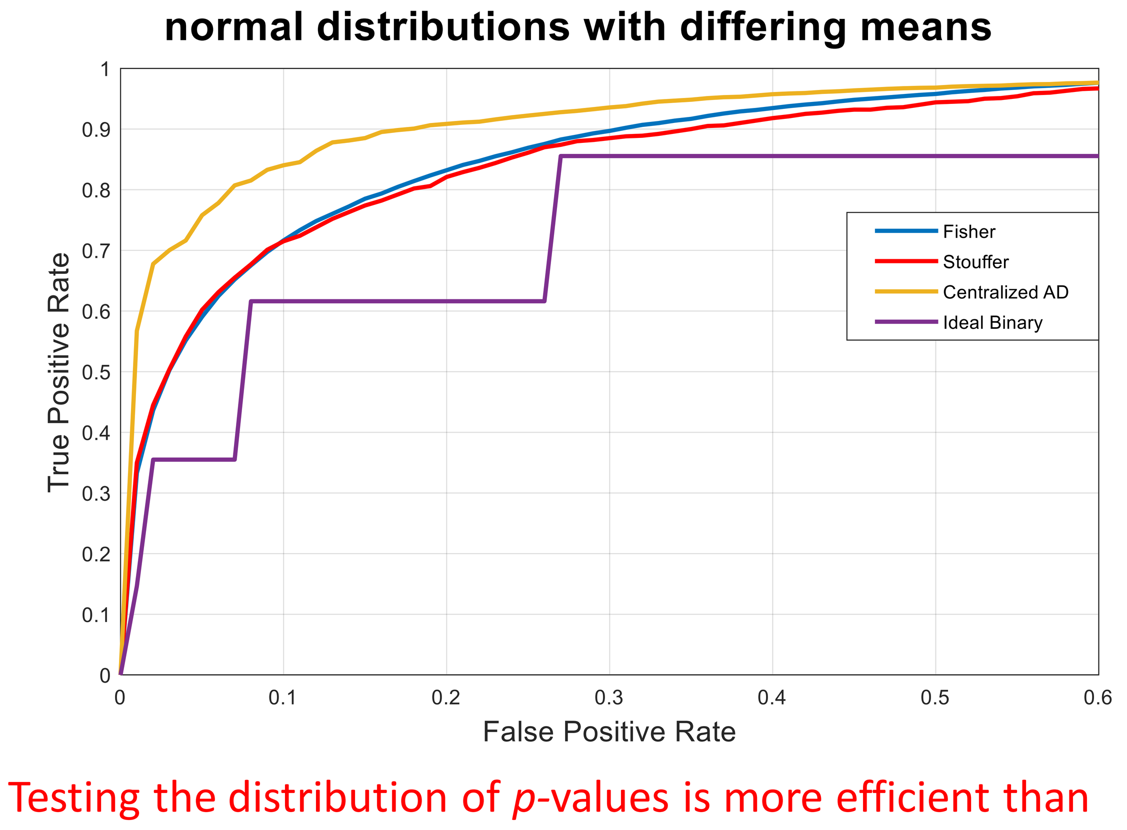

Testing the distribution of *p*-values is more efficient than making decision based on binary inputs. Information loss is experienced in comparison to a theoretical centralized scheme where sensors send all their data to FC, but the difference is not very significant.

Contact information:
visa.koivunen@aalto.fi