MCPS: past, present, and future

Insup Lee

RPECISE Center

Department of Computer and Information Science

University of Pennsylvania

November 22, 2019

MCPS Mini-Workshop

NSF CPS PI Meeting

Penn Engineering

PRECISE

PENN RESEARCH IN EMBEDDED COMPUTING AND INTEGRATED SYSTEMS ENGINEERING

Early Years

High-Confidence Medical Device Software and Systems (HCMDSS)

- Organized workshop in Philadelphia, June 2005
- One hundred participants from
  - academia
  - medical sectors (care-givers, researchers, etc.)
  - industry
  - government agencies
- Sponsors: NSF, NCO, Penn Engineering
- Supporting government agencies: FDA, NIST, NSA, ARO
- Goals
  - Identify research challenges and emerging issues
  - Produce a comprehensive **report** on research needs and **roadmap** at the national level across multiple agencies
  - Create a new scientific community
- www.cis.upenn.edu/hcmdss/

Penn Engineering

PRECISE

Roadmap

- Working Groups
  - Foundations for Integration of Medical Device Systems/Models
  - Distributed Control & Sensing of Networked Medical Device Systems
  - Patient Modeling & Simulatio
  - Embedded, Real-Time, Netwo Infrastructures for MDSS
  - High-Confidence Medical Dev Development & Assurance an en Models
  - Certification of MDSS and Req

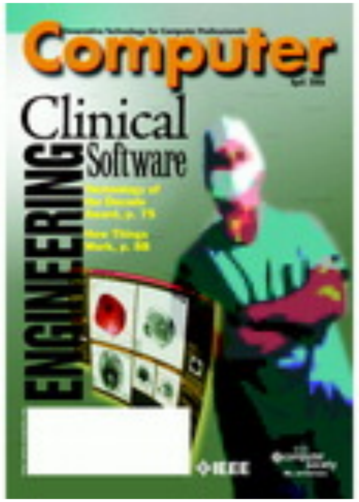

- Phase I (0-2 years)
  - Understand certification proce
  - Create a research community
  - Open experimental platforms

- Phase II (0-5 years)
  - Standards for secure data, communication, context.
  - middleware infrastructure
  - PnP device networks
  - ance and certification
  - clinical, system requirements
  - esign

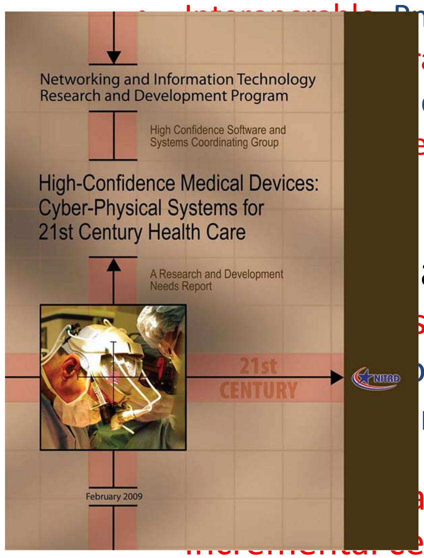

- ars)
  - s and simulators
  - r heterogeneous model-based design
  - nfigurable), fault-tolerant, distributed
  - ased verification/certification/testing
  - ertification

The Pacemaker Challenge

- Software Certification Consortium (SCC) in 2007
- In 2007. Brian Larson, who was at Boston Scientific, got permission to release a requirements document of a real, ten-year-old pacemaker
- Became the SCC challenge problem
- The SCC web site at http://www.cas.mcmaster.ca/wiki/index.php/Pacemaker
- Since then, much work on applying formal methods to modeling and verification of pacemakers

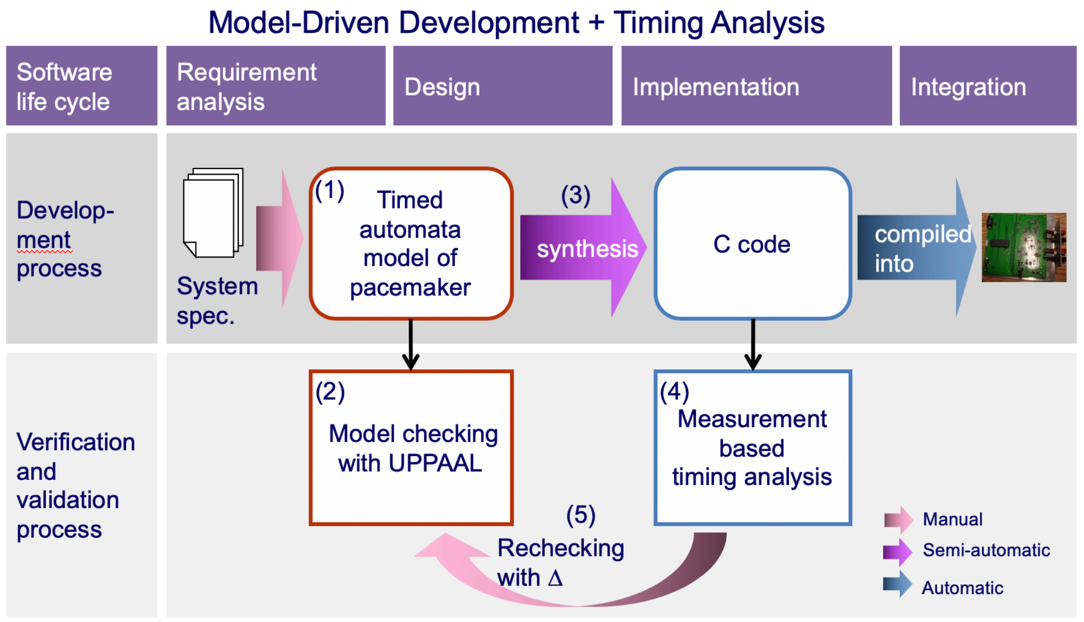

The Generic Infusion Pump

- Open infusion pump requirements and design specifications (incorporated lessons from earlier CARA work, 2002)
- Started as collaboration between FDA, Penn, Kansas, CIMIT, SEI, Minnesota, ASU, et al., 2007
- http://rtg.cis.upenn.edu/gip.php3
- FDA Infusion Pump Improvement Initiative, 2017
- https://www.fda.gov/medical-devices/infusion-pumps/infusion-pump-software-safety-research-fda

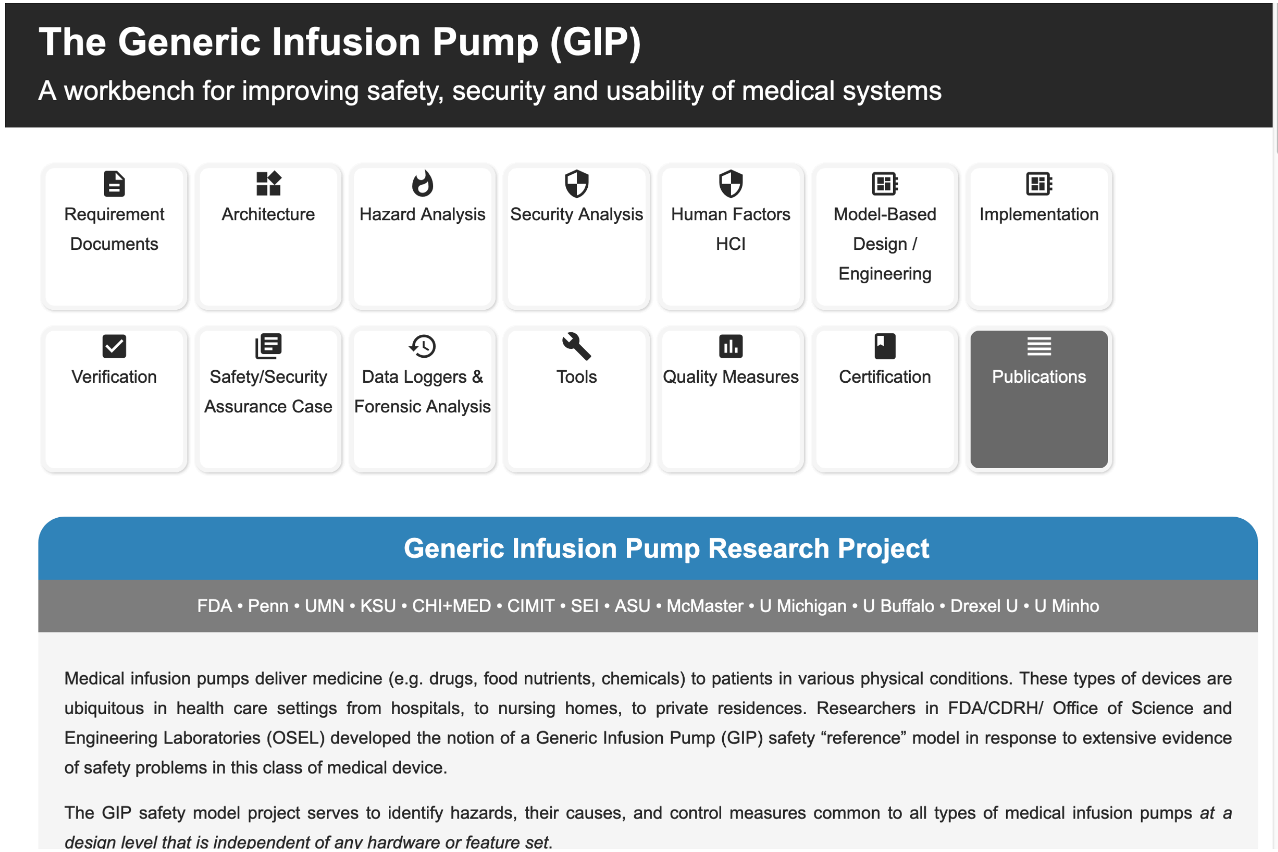

Fast Forward to Present

Internet of Medical Things

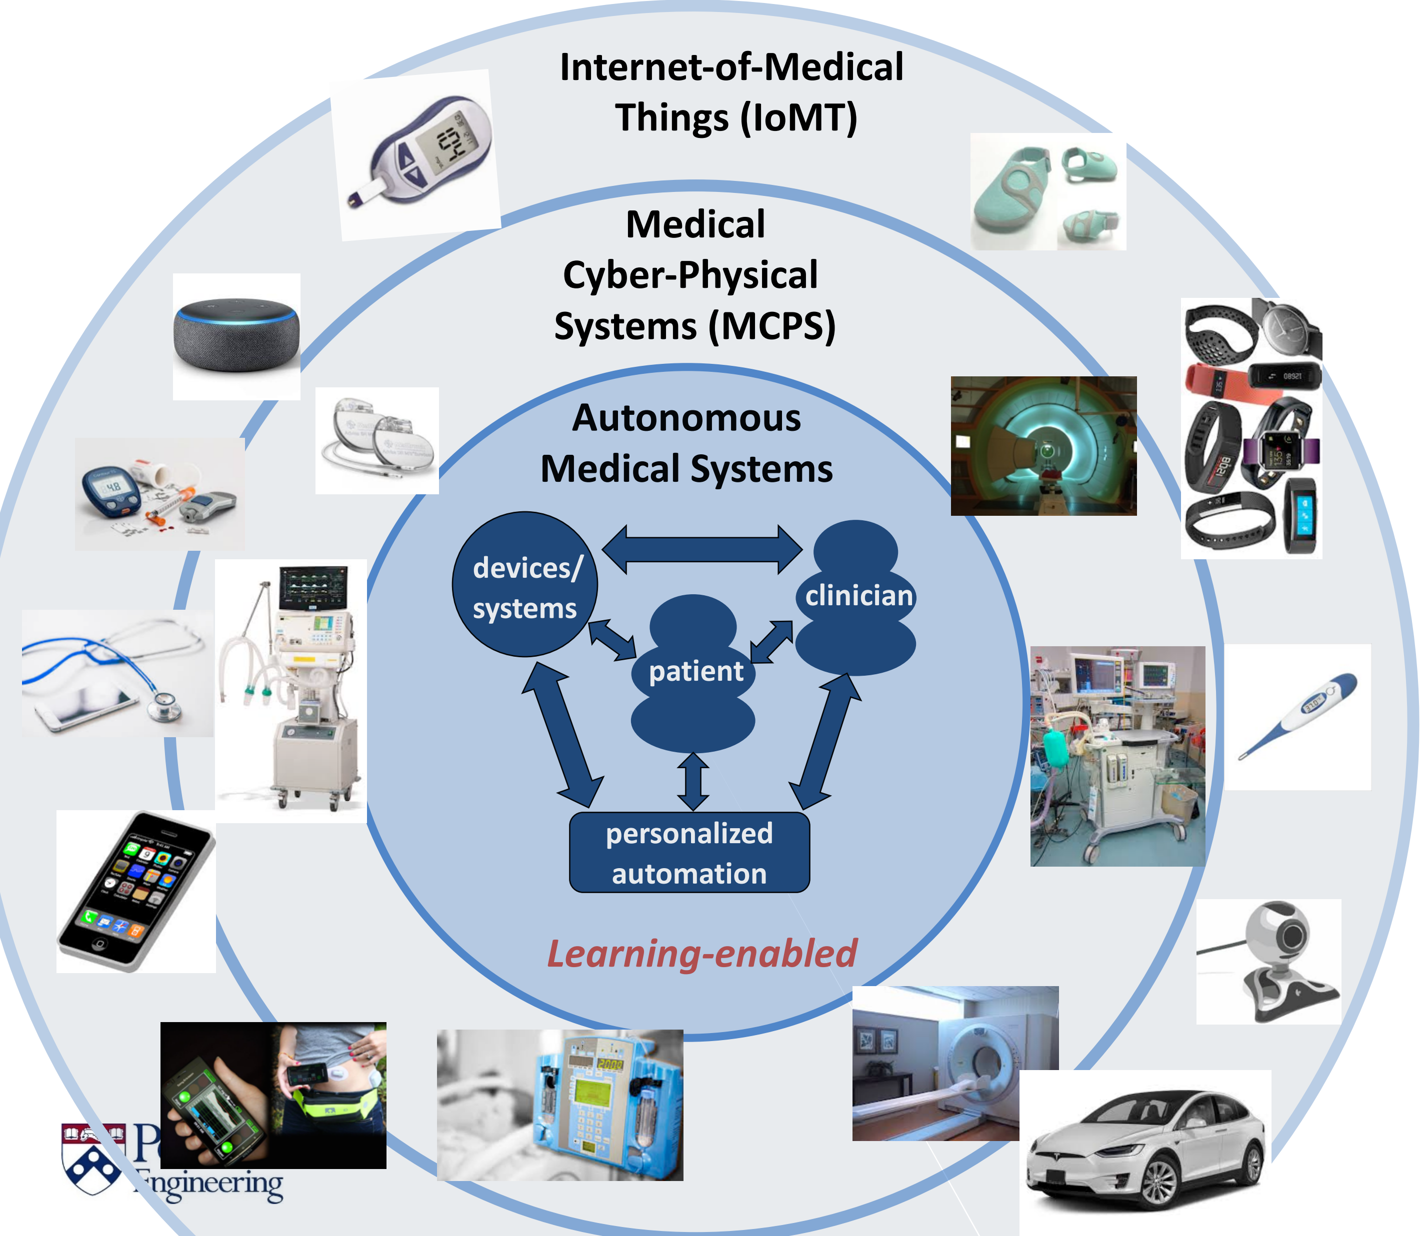

Software as Medical Devices (SaMD)

Medical Cyber-Physical Systems

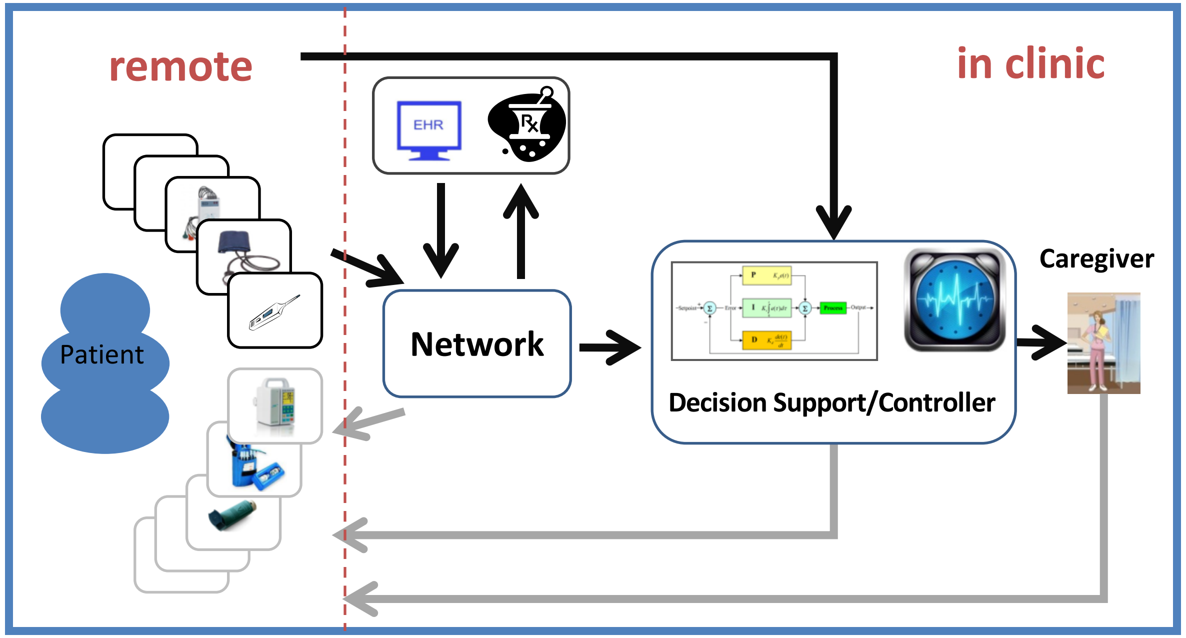

Penn Engineering

PRECISE

Smart Alarm and Clinical Decision Support Systems

CBS NEWS / *February 24, 2011, 12:37 PM*

"Alarm fatigue" blamed in hospital deaths

AP

Comment / 5 Shares /

A Boston Globe investigat… put patients at risk.

The newspaper says more …

THE WALL STREET JOURNAL.

Subscribe | Sign In

At the Hospital, Better Responses to Those Beeping Alarms

When patients' monitors beep falsely, nurses and doctors can get 'alarm fatigue' and miss real warnings

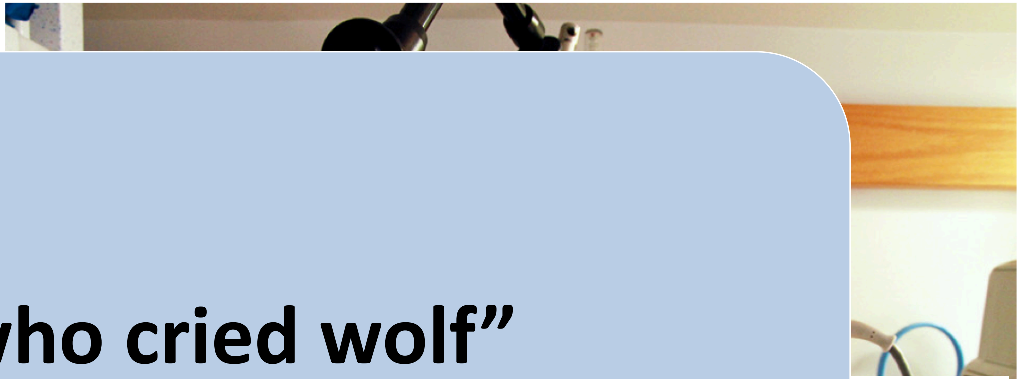

Boston Globe

…aff; the

"Boy who cried wolf"

SEPTEMBER 10, 2015 | THE PULSE

Beep, beep, beep - hospital alarms sound mostly without real cause

Beep, beep, beep - hospital alarms sound mostly without re…
Listen 0:00 / 9:31

Penn Engineering

PRECISE

Smart Alarms in Hospitals

- Medical device alarms are non-informative
  - between 80% and 99% of all alarms are false
- Clinicians have developed **alarm fatigue** and may not respond to alarms
  - A top 10 health technology hazard since 2007 by ECRI
- ECRI 2019 #7 Health Tech Hazard: Improper Customization of Physiologic Monitor Alarm Settings May Result in Missed Alarms
- Smart alarm suppression
  - Maximally suppress non-informative alarms without suppressing actionable alarms
- Case study: low SpO2 alarms

Low SpO2 Smart Alarm Results

[CHASE2018]

- Dataset: 100 children at CHOP
  - Alarms are annotated by clinicians via video feedback
- AdaBoost with reject option algorithm
  - Prioritize achieving a low false negative rate for each weak learner
  - Only make a decision when almost all weak learners agree
  - identifies "easy to silence" alarms
- Results:
  - 100% sensitivity to actionable alarms
  - Silenced 23% of alarms (413 of 1786)
  - Significantly outperforms other learning approaches

- Future work: Smart Alarm 2.0
  - Context-aware smart alarm
  - How to model caregivers and tune system?
  - In-clinic vs. at-home monitoring?

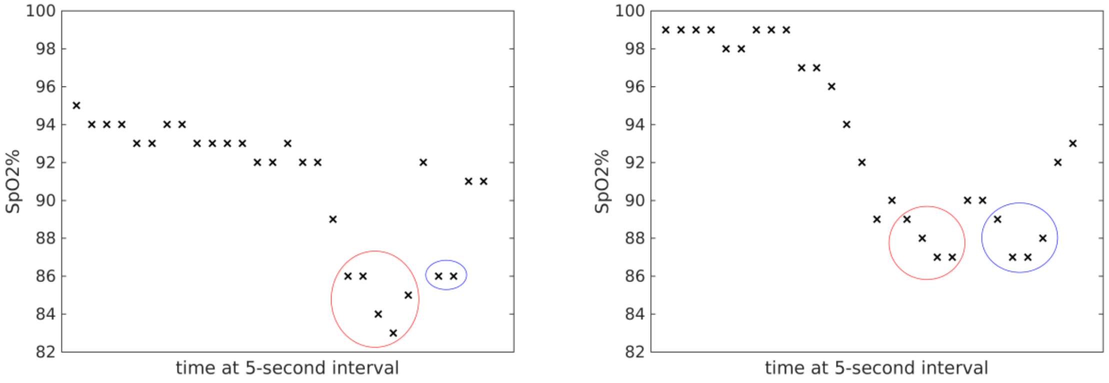

(a) Case #1 (b) Case #2

Penn Engineering

PRECISE

Challenges of Clinical Decision Support

Actionable **Clinical Decision Support (learning-enabled)**

Inter/Intra Patient Variability

…

Presents a Fundamental Technical Challenge

We want Big Deep data

Consider a population's data

… but we have Big Thin data.

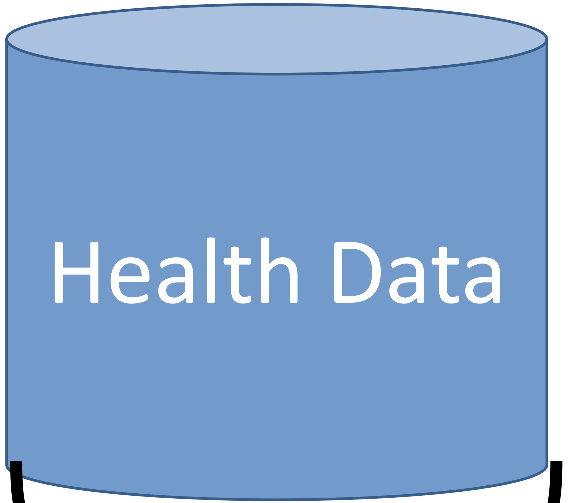

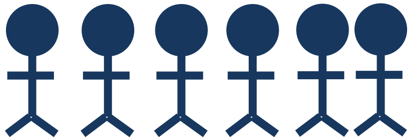

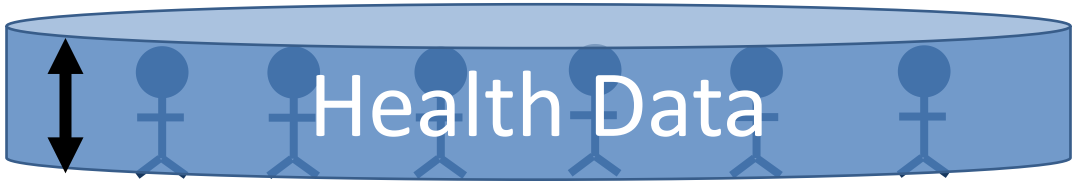

"Deep Data" e.g., good sampling of distribution

"Thin Data" – poor sampling of distribution

- Inter/intra-patient variability
- Anomalies (e.g., bad data)
- Limited sensing/actuation

Lots of data (i.e., "Big Data")

Big Thin Data

Penn Engineering

PRECISE

Challenges of Clinical Decision Support

Actionable* Clinical Decision Support** (learning-enabled) | ***Inter/Intra Patient Variability | ***Bias/Fairness*** (learning-enabled) | ***Security/ Privacy*** | ***Regulatory/ Certifiability*** …

- Inter/intra-patient variability
- Anomalies (e.g., bad data)
- Dataset shifts (e.g., co-morbidites)
- Limited data, bias, fairness
- Privacy, differential privacy
- Regulatory Challenge*

*[FDA, Proposed Regulatory Framework for Modifications to AI/ML Based Software as a Medical Device, 2019]

Data-driven Behavior Modeling & Learning-enabled Closed-the-Loop Systems

Human-in/on-the-loop CPS

IoMT applications interact with **human operators** who act as supervisors or collaborators

Global Hawk UAV, 98-2003, 19991206, FSPM 1201A

The Global Hawk incident caused by the lack of proper coordination between operator and autonomy

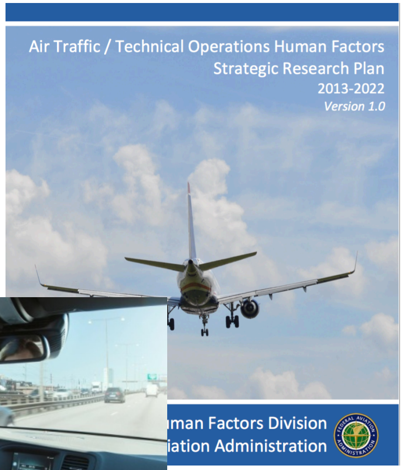
Air Traffic / Technical Operations Human Factors
Strategic Research Plan
2013-2022
Version 1.0
man Factors Division
iation Administration

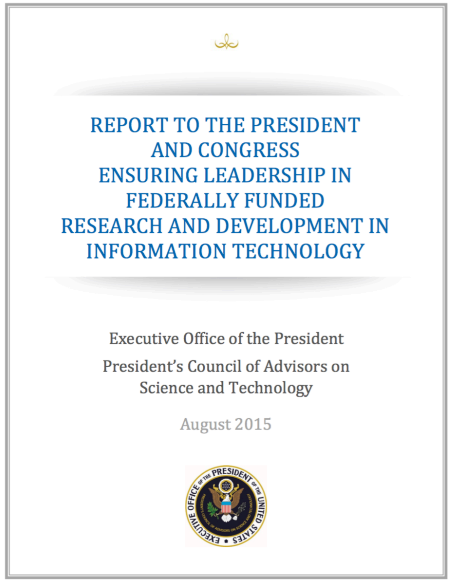
REPORT TO THE PRESIDENT
AND CONGRESS
ENSURING LEADERSHIP IN
FEDERALLY FUNDED
RESEARCH AND DEVELOPMENT IN
INFORMATION TECHNOLOGY

Executive Office of the President
President's Council of Advisors on
Science and Technology

August 2015

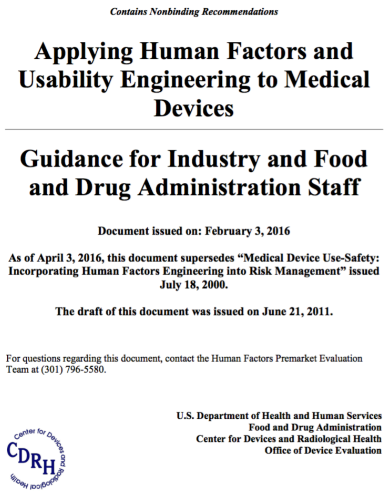
Contains Nonbinding Recommendations

Applying Human Factors and Usability Engineering to Medical Devices

Guidance for Industry and Food and Drug Administration Staff

Document issued on: February 3, 2016

As of April 3, 2016, this document supersedes "Medical Device Use-Safety: Incorporating Human Factors Engineering into Risk Management" issued July 18, 2000.

The draft of this document was issued on June 21, 2011.

For questions regarding this document, contact the Human Factors Premarket Evaluation Team at (301) 796-5580.

U.S. Department of Health and Human Services
Food and Drug Administration
Center for Devices and Radiological Health
Office of Device Evaluation

"There is a need for research focused on ***human interaction with systems that operate in the physical world***, particularly around issues of **safety**, **trust**, and **predictability of response**."

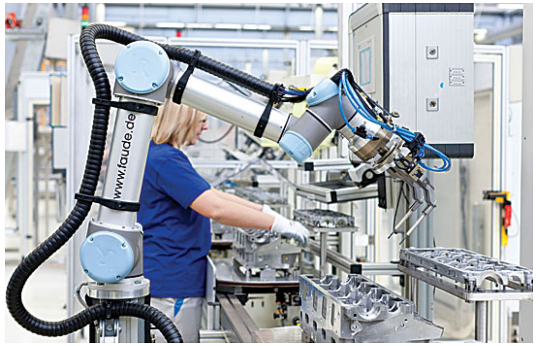

Penn Engineering

PRECISE

Data-driven behavior modeling for diabetic patients

- 30% - 40% T1D patients in the US use insulin pumps
- Requires user supervision
  - Input meal information, approve pump-suggested boluses, give non-mealtime boluses, calibrate CGM sensor
- American Association of Clinical Endocrinologists report highlights critical needs for better understanding the ***physiological and psychological*** impacts of insulin pumps on diabetic users

<u>Real patients' data</u>
- 55 T1D patients
- ≈ 30 days

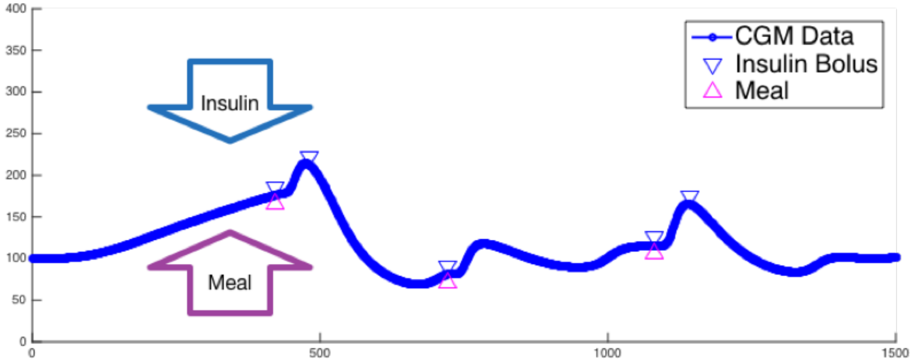

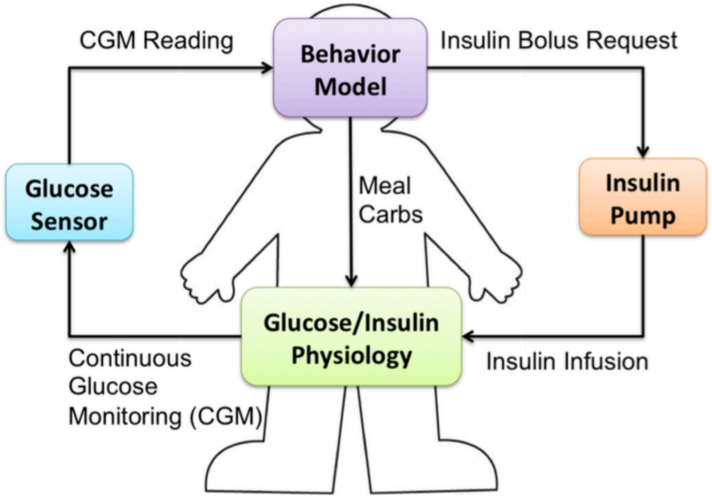

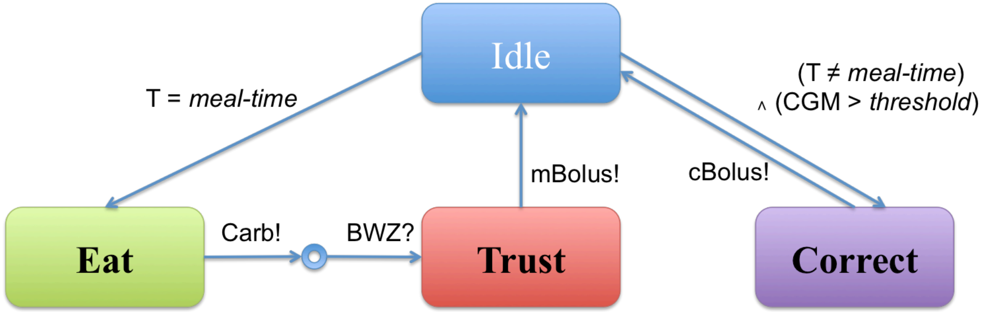

Clinically-relevant insights from *in silico* analysis via quantitative verification

Trust Clusters

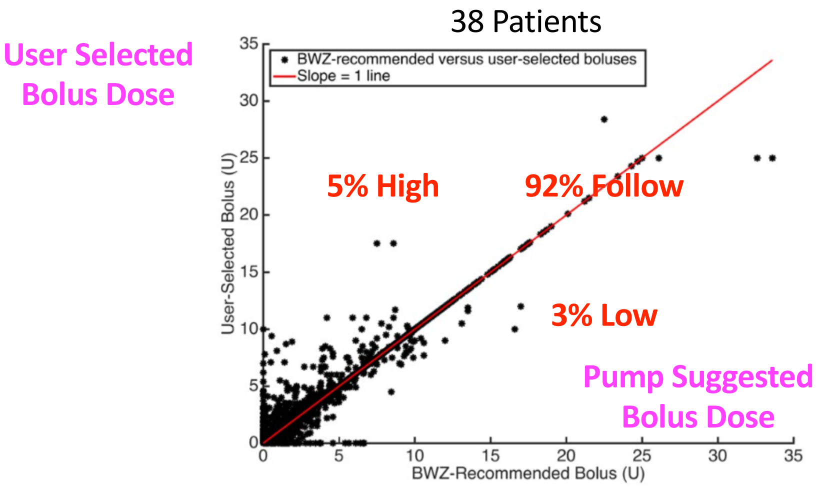

(a) Trust T1: high probability of following BWZ-recommended doses

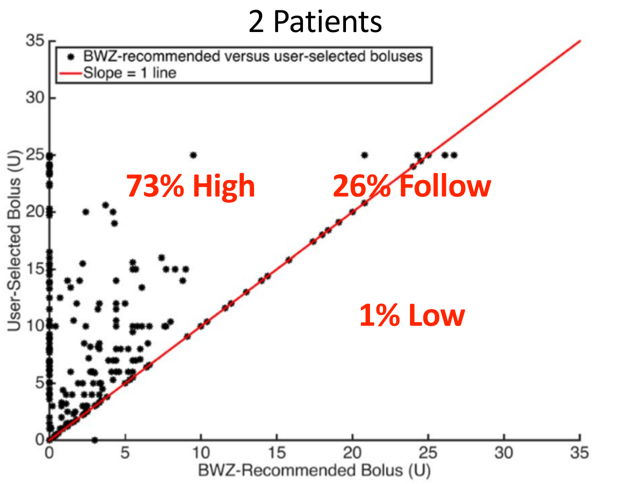

(b) Trust T2: high probability of increasing BWZ-recommended doses

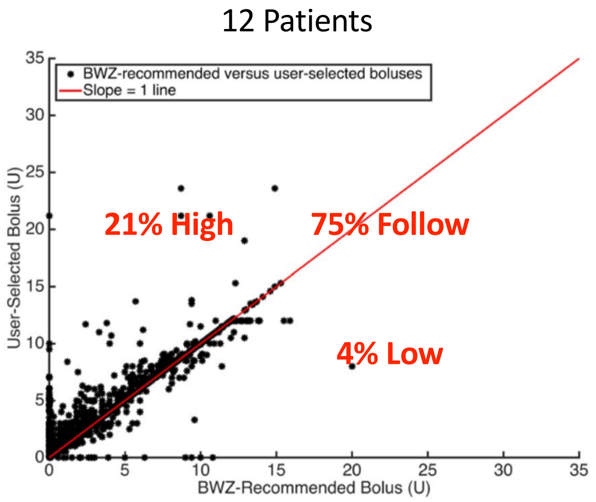

(c) Trust T3: moderate probability of increasing BWZ-recommended doses

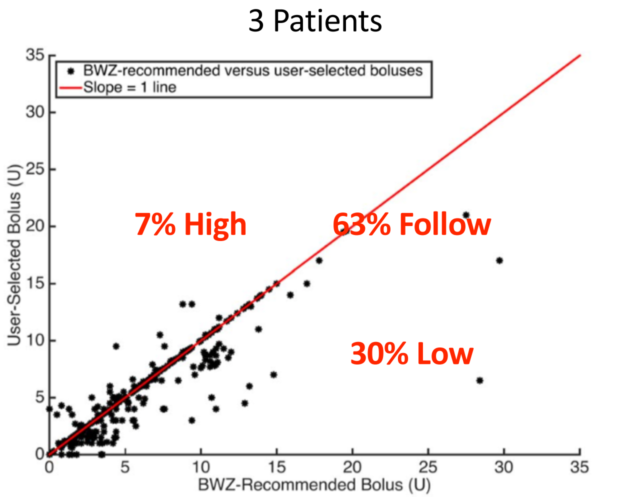

(d) Trust T4: high probability of decreasing BWZ-recommended doses

[ICHI'15]

Our Approach & Analysis Results

- Behavioral modeling: unsupervised learning
- Physiological modeling: fitting a standard physiological model
- Closed-loop analysis: probabilistic model checking
- Patient education/peer-support: how behaviors affect outcomes

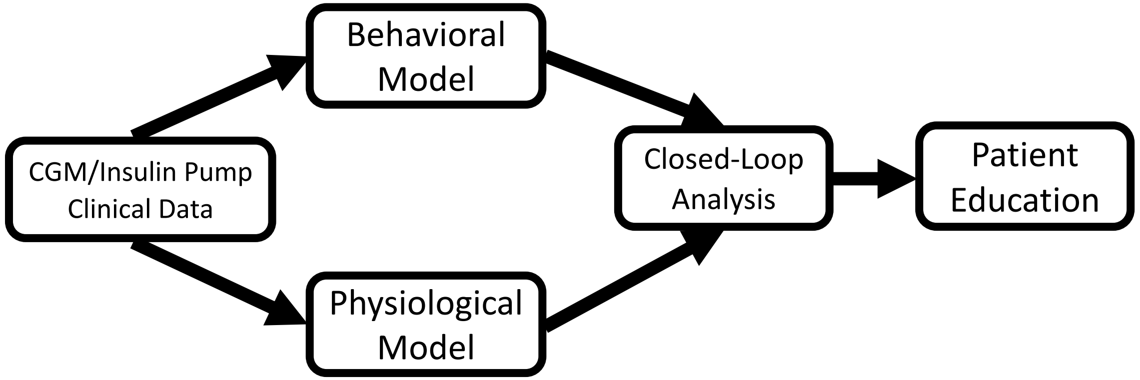

- PRISM model checker
  - Support probabilistic transitions
  - Enables exhaustive check all execution paths of a model
- Integrate individualized physiological model and behavioral models
  - Explore how changing behavior types may impact outcomes
  - Hypoglycemia: % of CGM readings < 70 mg/dL
  - Hyperglycemia: % of CGM readings > 180 mg/dL

	ETC Type	Hypoglycemia Rate (%)	Hyperglycemia Rate (%)
Actual type	E3T2C1	**6.93**	8.43
Change E subtype	E1T2C1	6.20	12.78
	E2T2C1	5.99	13.72
Change T subtype	E3T1C1	0.02	10.33
	E3T3C1	0.04	10.09
	E3T4C1	0.02	11.05
Change C subtype	E3T2C2	7.04	6.30
	E3T2C3	6.95	7.93
Change multi-subtypes	E2T1C1	0.04	16.46
	E2T2C1	5.99	13.72
	E3T1C3	0.10	9.76
	E2T1C3	0.08	15.42

	ETC Type	Hypoglycemia Rate (%)	Hyperglycemia Rate (%)
Actual type	E1T1C1	0	**43.92**
Change E subtype	E2T1C1	0	44.38
	E3T1C1	0	41.62
Change T subtype	E1T2C1	0	39.13
	E1T3C1	0	43.46
	E1T4C1	0	45.31
Change C subtype	E1T1C2	0	41.59
	E1T1C3	0	43.47
Change multi-subtypes	E1T2C2	0	37.22
	E3T2C1	0	35.45
	E3T1C2	0	38.01
	E3T2C2	0	32.56

Safe Learning-Enabled Physiological Control

Challenge: *assuring safety w/ learning-enabled components in-the-loop*

- **Key observation**: derivative of a sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}, \qquad \frac{d\sigma}{dx}(x) = \sigma(x)\cdot(1-\sigma(x))$$

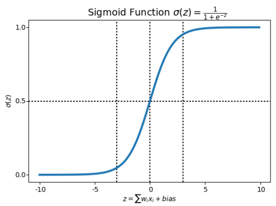

- Introduce a proxy function

$$g(t,x) = \frac{1}{1+e^{-tx}}$$

  - where $g(1,x) = \sigma(x)$ and $\dot{g} = xg(t,x)(1-g(t,x))$, $g(0,x) = 0.5$
- *Each neuron in DNN can be encoded by a corresponding g*

deep neural network (sigmoid, tanh, swish, etc.)

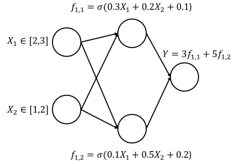

Keras model

physiological model (hybrid system)

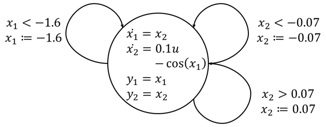

SpaceEx model

Safety property **[HSCC'19]**

SpaceEx config file

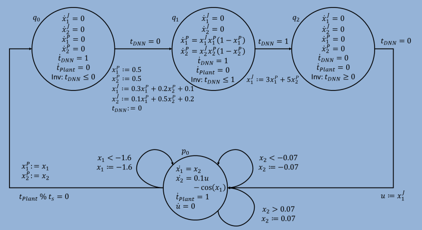

Verisig (www.verisig.org)

current applications

Mechanical Ventilation Weaning

- **Adaptive Support Ventilation (Hamilton)**
- **SmartCare/PS (Drager)**

Type 1 Diabetes

- **Medtronic 670G / 690G**
- **Dexcom + t:slim X2**
- **OpenAPS – NOT APPROVED**

Autowean: Foundations for Autonomous Medical CPS for Mechanical Ventilation Weaning (NIH 1R01EB029767-01)

PRECISE
PENN RESEARCH IN EMBEDDED COMPUTING AND INTEGRATED SYSTEMS ENGINEERING

Human-in-the-Loop

- Interactions between caregivers, patients, and autonomous medical systems
- How can we assure safety?
- Analysis of safety and effectiveness needs to take their behaviors into consideration
  - Better modeling techniques
  - How much the user trusts the system
  - When and how the user interferes with automation

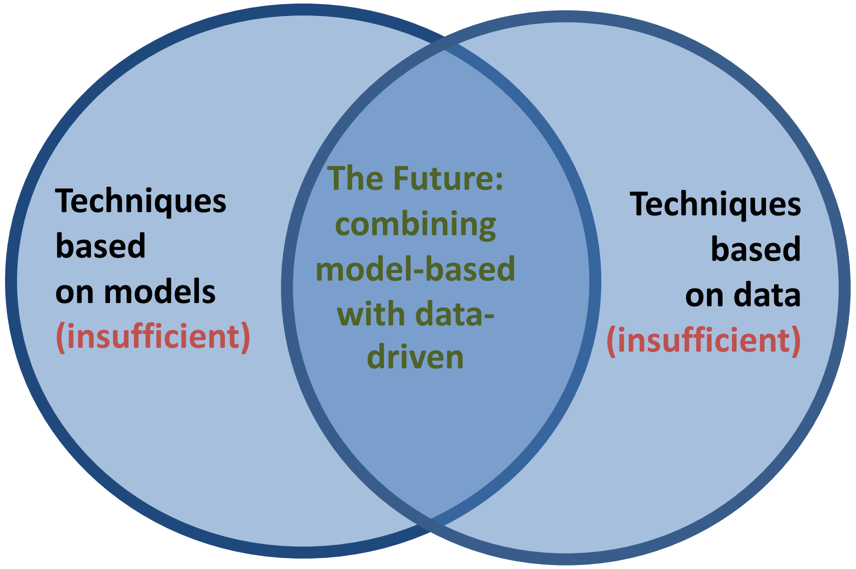

Next-Generation Medical Systems

Levels of autonomy

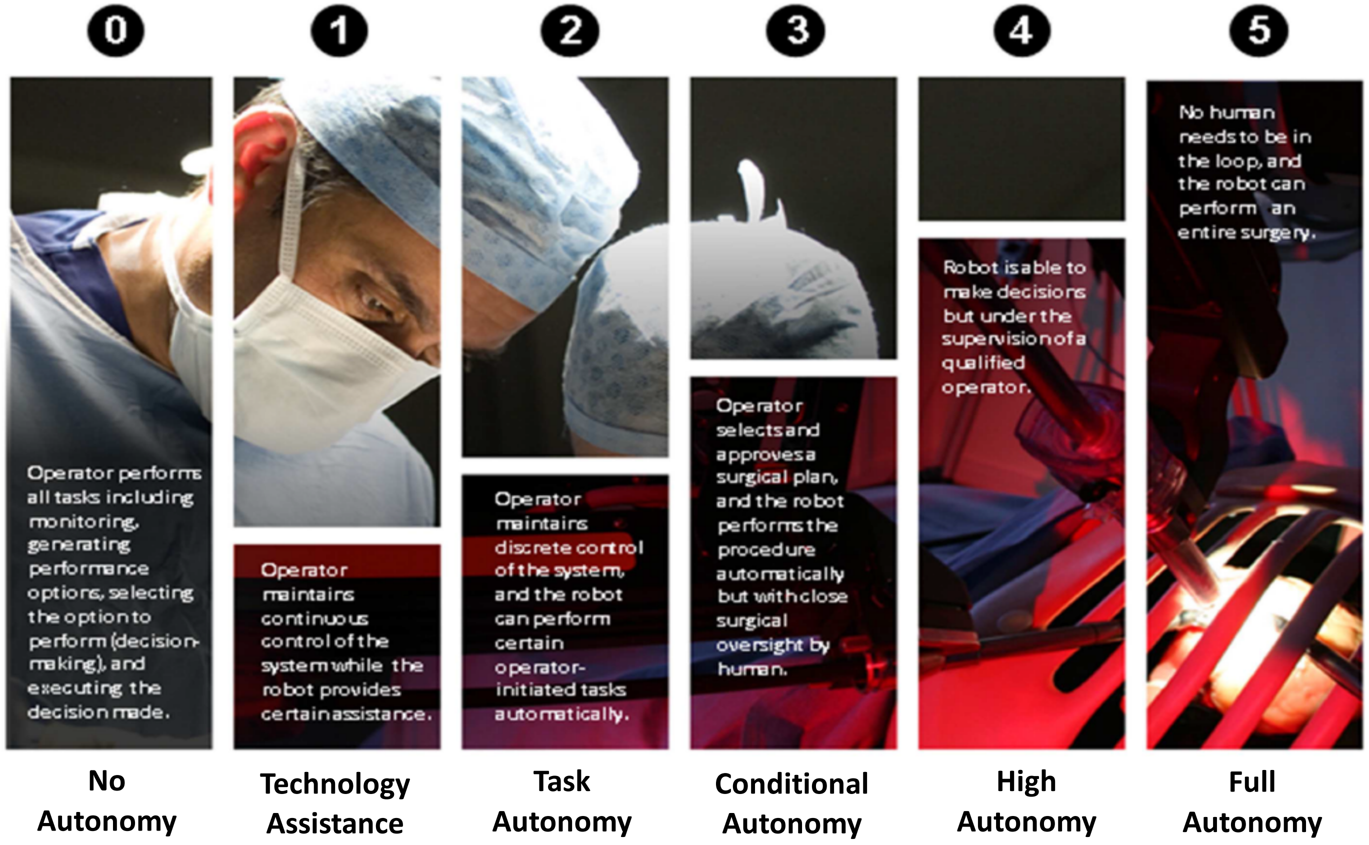

Challenges of learning-enabled autonomous medical systems

- Modeling patient physiology
- Understanding human-automation interactions
- Formal design of closed-loop systems with human-in-the-loop
- Human factors in MCPS
- Safety and fairness assurance (and certification)
- Security and privacy
- Interoperable medical device platforms that are open and trustworthy
- Benchmark problems/testbeds

Thank You!

Penn Engineering

PRECISE
PENN RESEARCH IN EMBEDDED COMPUTING AND INTEGRATED SYSTEMS ENGINEERING