Knowledge Discovery for Clinical Decision Support

Pedro Pereira Rodrigues

CINTESIS & LIAAD – INESC TEC

Faculty of Medicine – University of Porto, Portugal

@ECMLPKDD – Nancy, France

Who am I?

A privileged one, who being educated in machine learning, gets to teach medical students on research methodology and data science ;-)

- MSc (2005) and PhD (2010) on clustering data streams and stream sources.
- Last 6 years involved in medical informatics, clinical research and medical education.

Coordinator of the **BioData - Biostatistics and Intelligent Data Analysis** group of **CINTESIS - Centre for Health Technologies and Services Research** (100+ PhD research unit to start officially in 2015) and collaborator in **LIAAD – INESC TEC** (original research unit since 2003).

Agenda IV

- Uncertainty and evidence-based medicine
- Data science in the EBM loop
- Biostatistics and probabilistic decision support
- Bayesian networks as formalization of uncertainty for decision support
- Toy and real-world examples of Bayesian nets for clinical decision support
- Lessons learned

Uncertainty and Evidence Based Medicine

Uncertainty in clinical decision

Uncertainty in clinical decision analysis

- The consequences of a medical decision are uncertain by the time of decision.
- Clinical exam and diagnostic tests are inperfect.
- Therapeutic actions, as well as their risks and benefits, might be vaguely defined or even unknown.
- For a large group of clinical problems,

there is no information about clinical trials,

or it simply **isn't generalizable** for the patient.

D. Owens and H. Sox, “Biomedical decision making: probabilistic clinical reasoning,” in Biomedical Informatics, Chapter 3, Springer Verlag, 2006, pp. 80–132.

Evidence-based medicine

Conscient, explicit and criterious use of the **best available evidence** in clinical decision:

- personal clinical experience;
- best external clinical evidence from quality clinical research;
- values, needs, expectations and individual context of each patient.

Sackett D. et al. (1996)

Evidence based medicine: what it is and what it isn't

BMJ 312:71-2

Take away message

M1: During inference and decision support, uncertainty needs to be reduced.

S1: Better focus on the variables that reduce uncertainty the most (e.g. when sugesting a test).

Evidence-based medicine

Conscient, explicit and criterious use of the **best available evidence** in clinical decision:

- personal clinical experience;
- best external clinical evidence from quality clinical research;
- values, needs, expectations and individual context of each patient.

UNCERTAIN

Sackett D. et al. (1996)

Evidence based medicine: what it is and what it isn't

BMJ 312:71-2

Evidence-based medicine

Conscient, explicit and criterious use of the **best available evidence** in clinical decision:

UNCERTAIN

- personal clinical experience;
- best external clinical evidence from quality clinical research;
- values, needs, expectations and individual context of each patient.

UNCERTAIN

Sackett D. et al. (1996)

Evidence based medicine: what it is and what it isn't

BMJ 312:71-2

Evidence-based medicine

Conscient, explicit and criterious use of the **best available evidence** in clinical decision:

- personal clinical experience;
- best external clinical evidence from quality clinical research;
- values, needs, expectations and individual context of each patient.

Sackett D. et al. (1996)

Evidence based medicine: what it is and what it isn't

BMJ 312:71-2

Clinical information/knowledge/decision

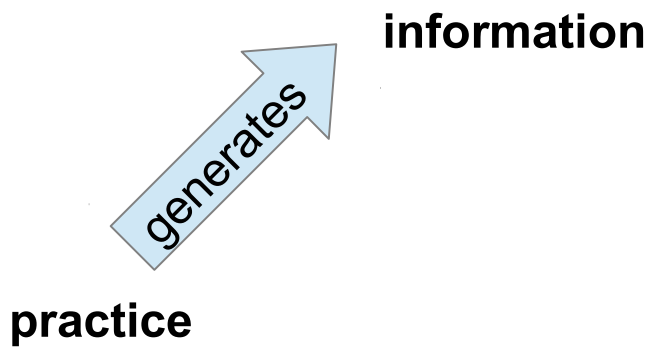

Clinical information/knowledge/decision

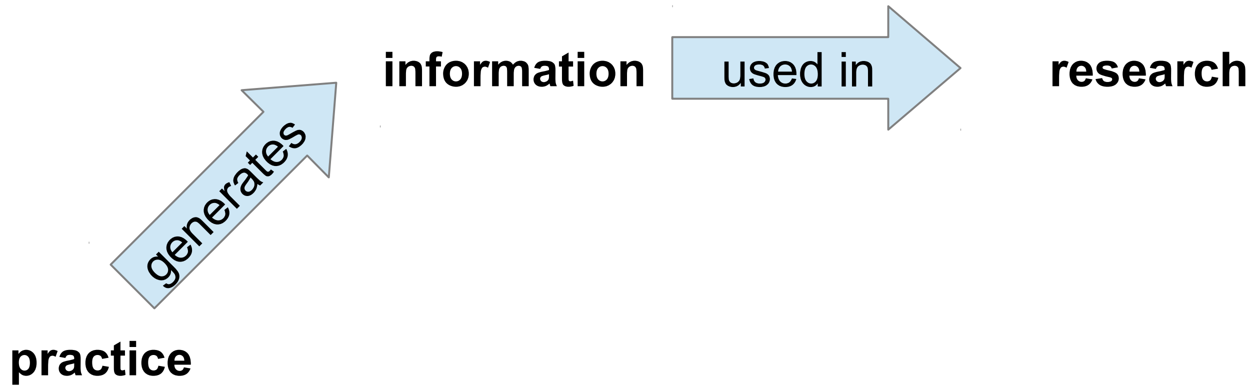

Clinical information/knowledge/decision

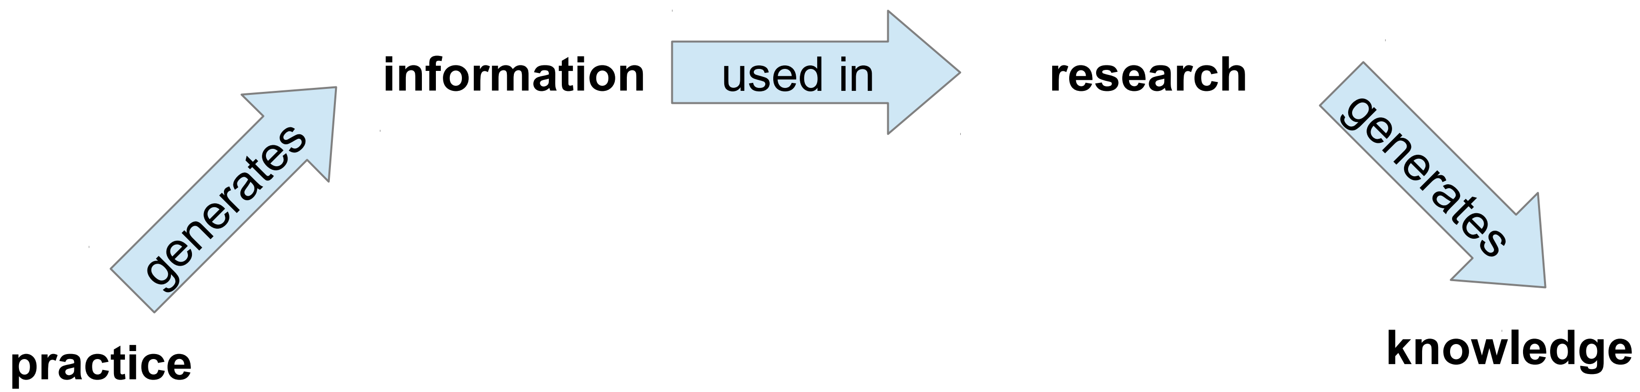

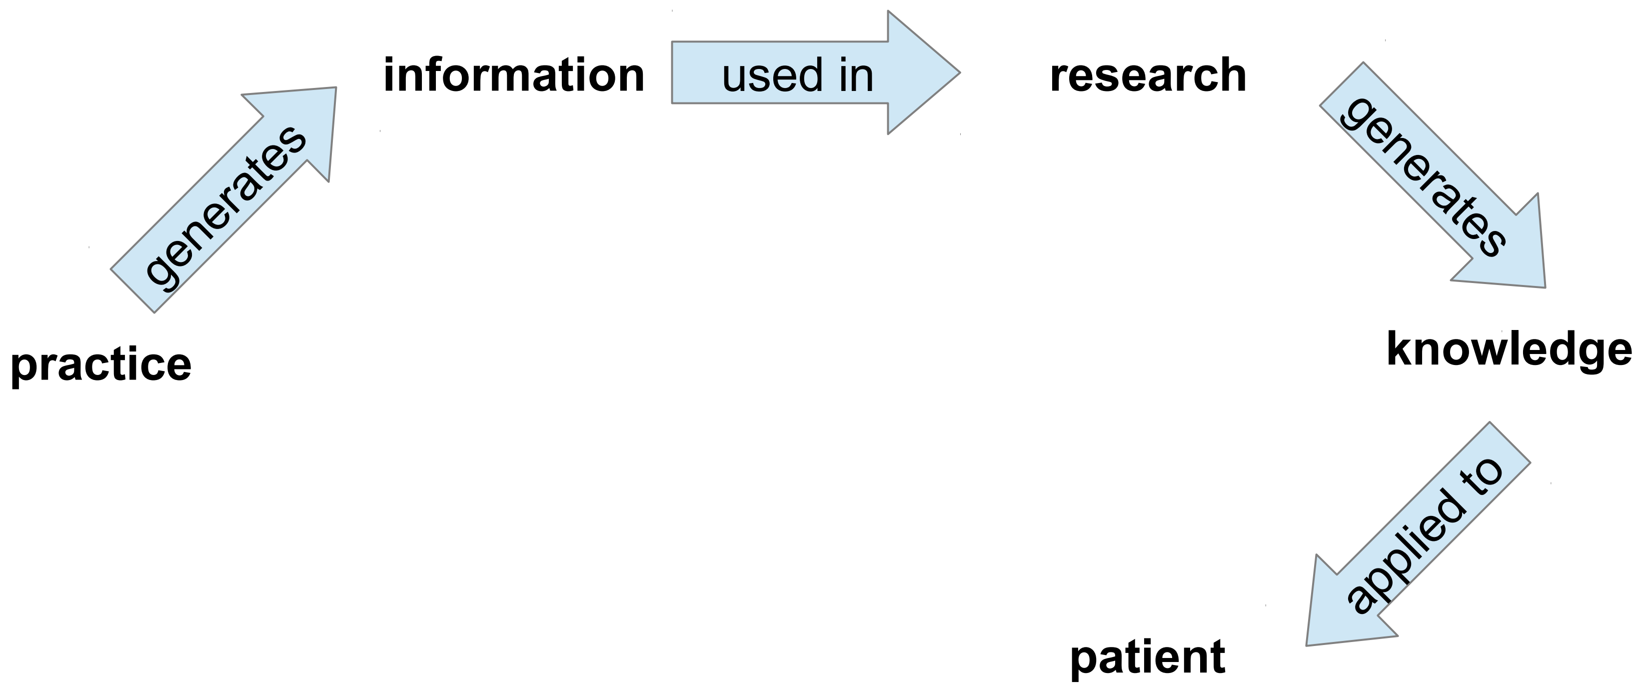
information
used in
research
generates
generates
practice
knowledge
applied to
patient

Clinical information/knowledge/decision

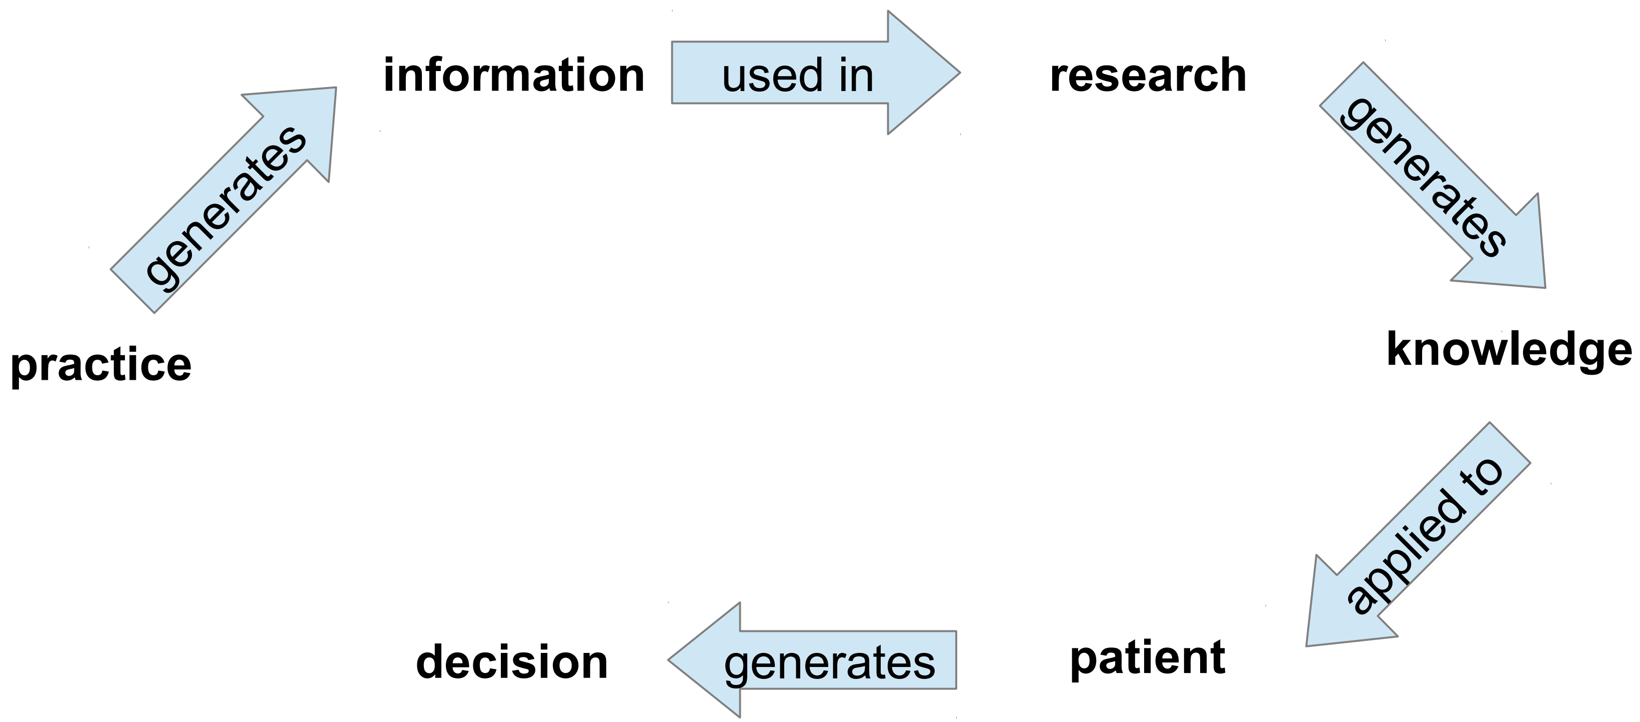

Clinical information/knowledge/decision

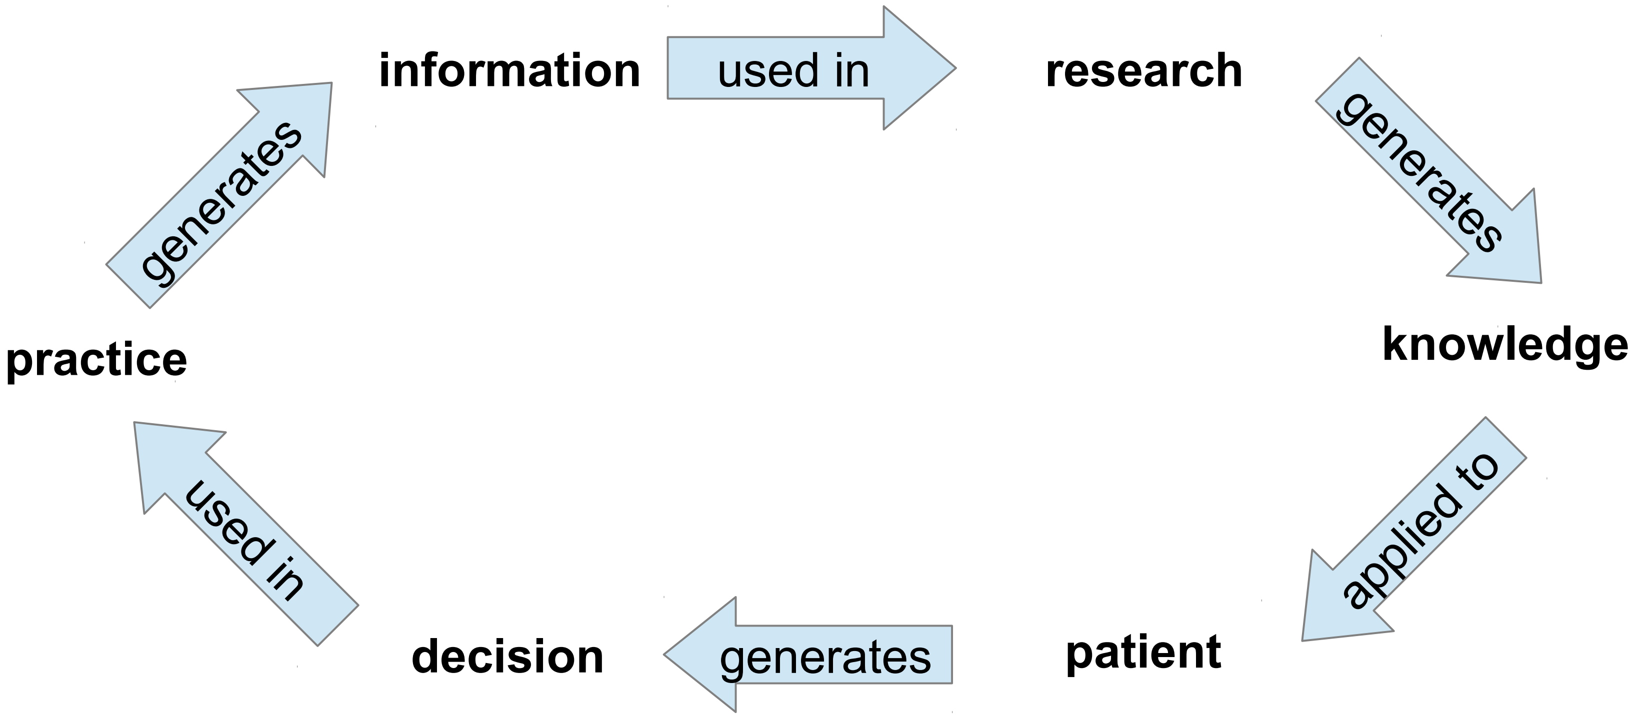

Where is data science involved?

Clinical information/knowledge/decision

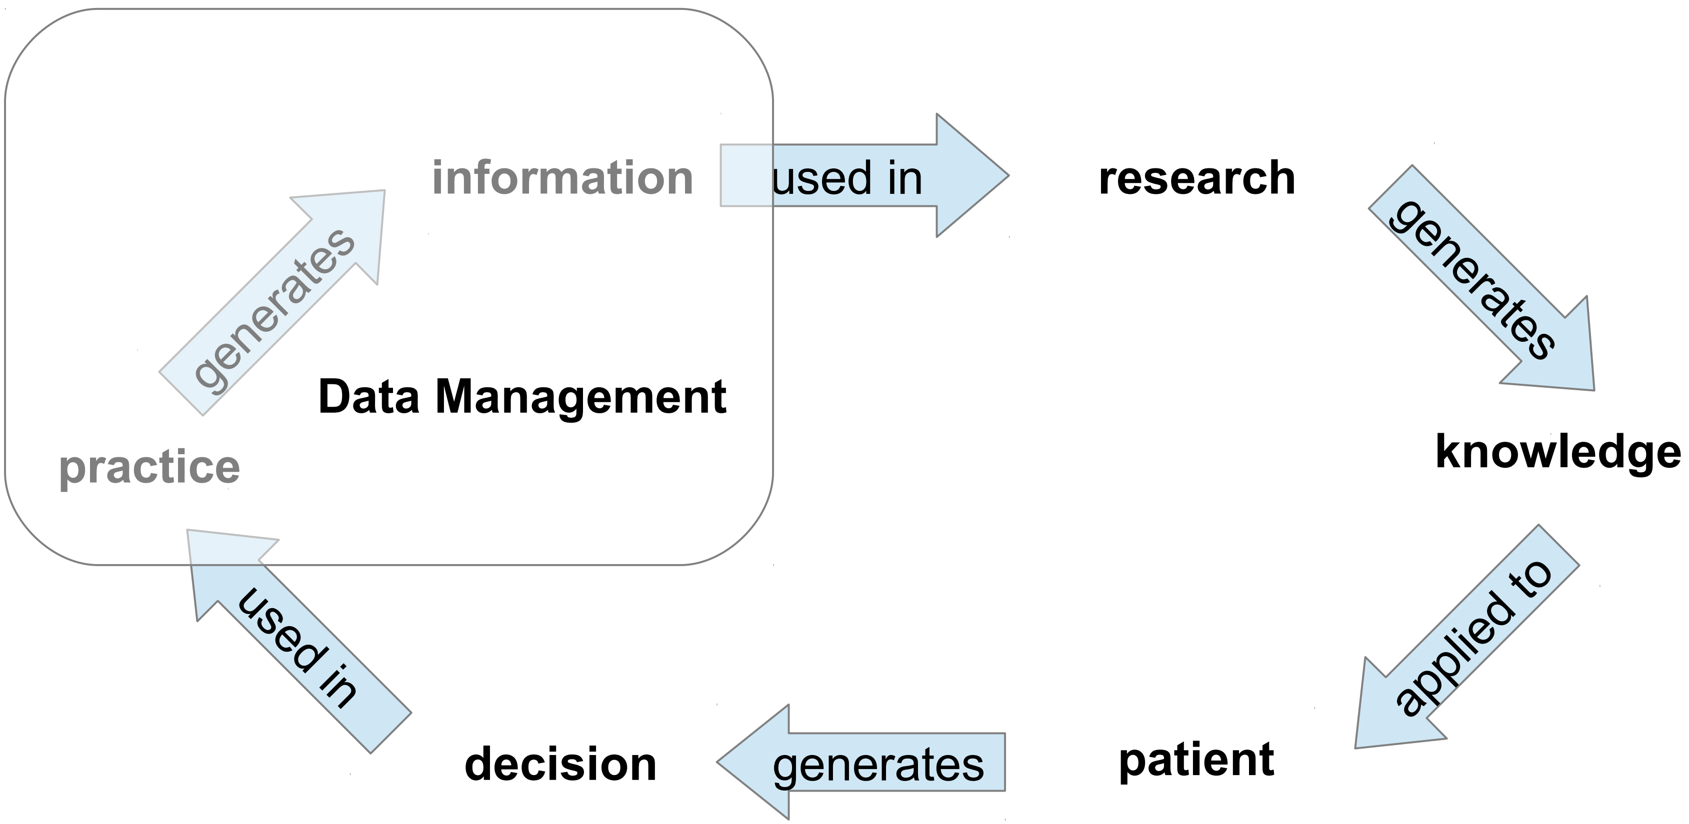

Clinical information/knowledge/decision

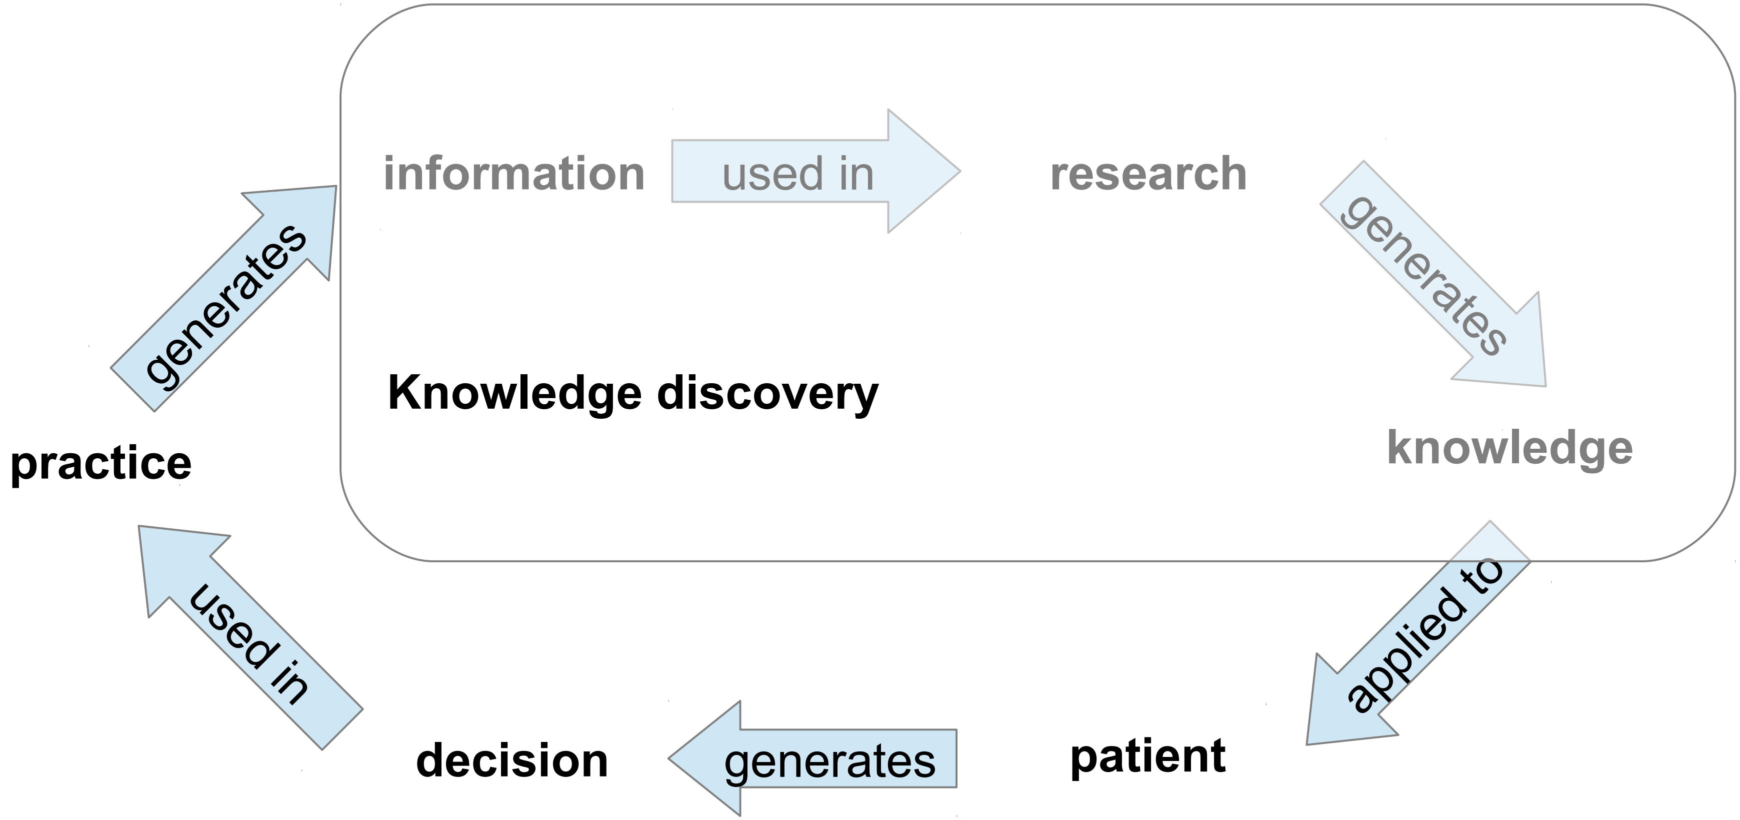

Clinical information/knowledge/decision

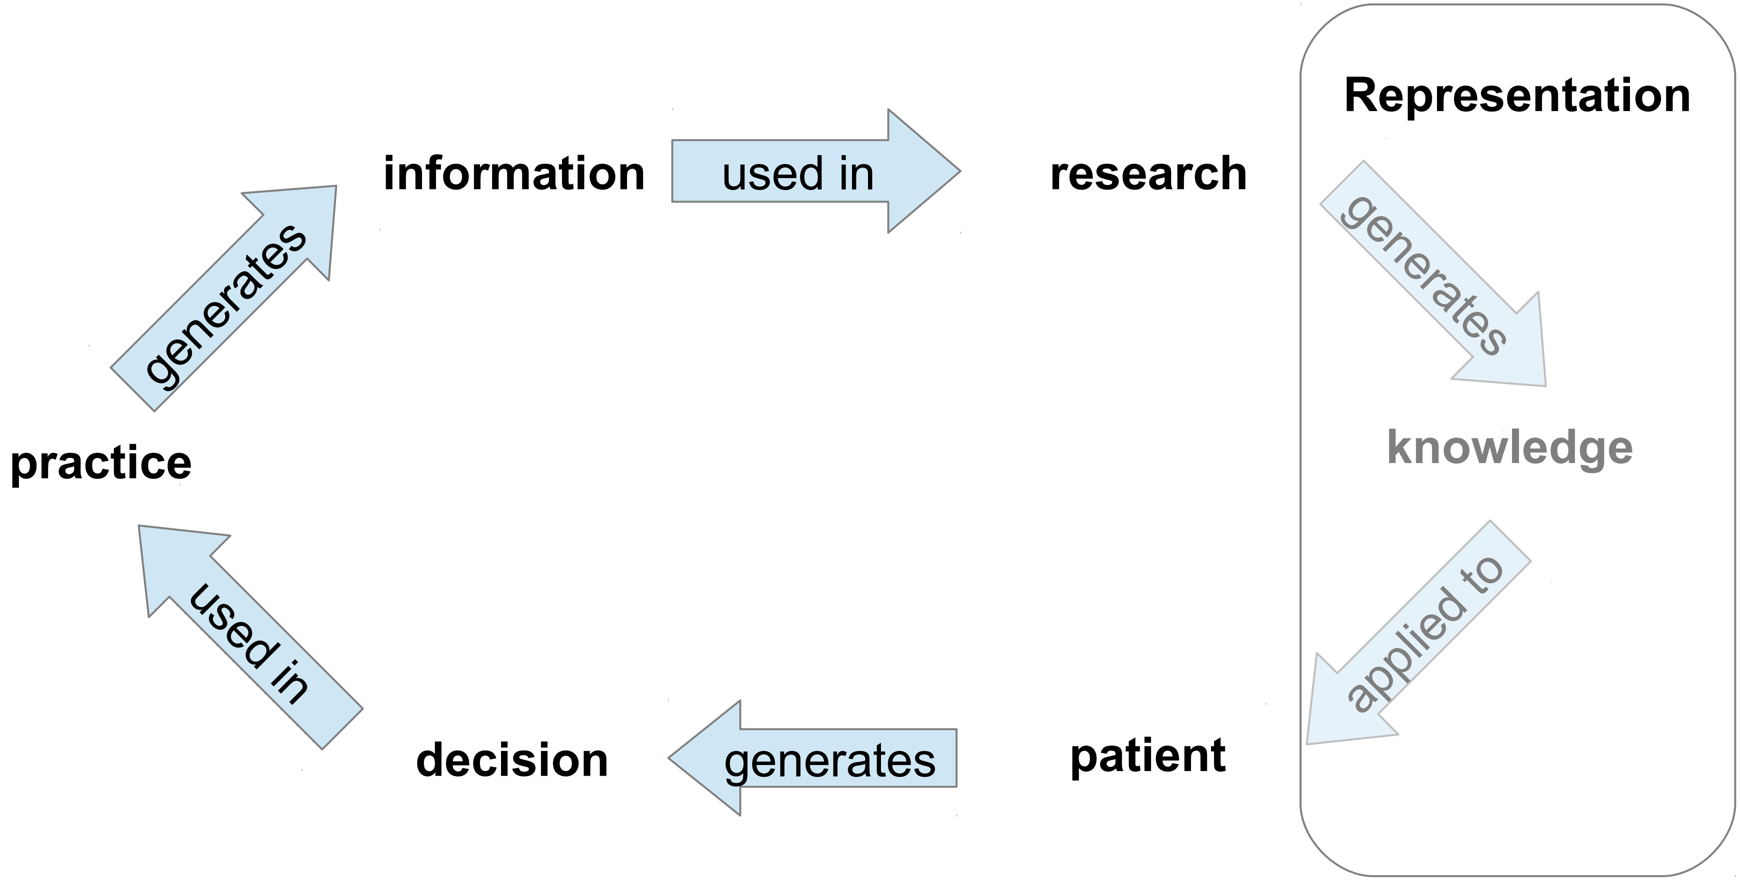

Clinical information/knowledge/decision

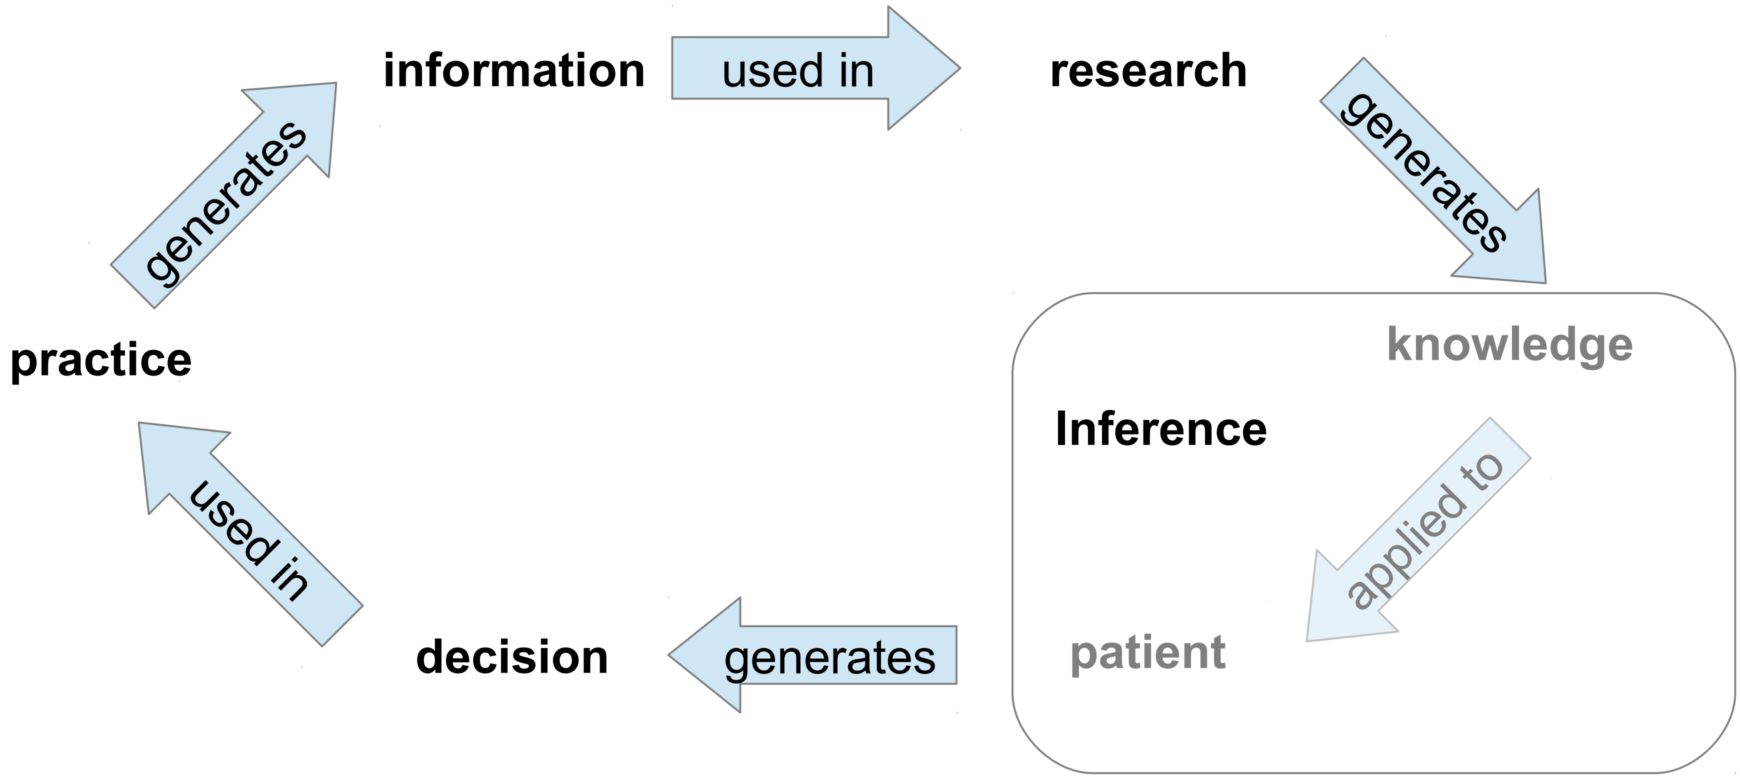

Clinical information/knowledge/decision

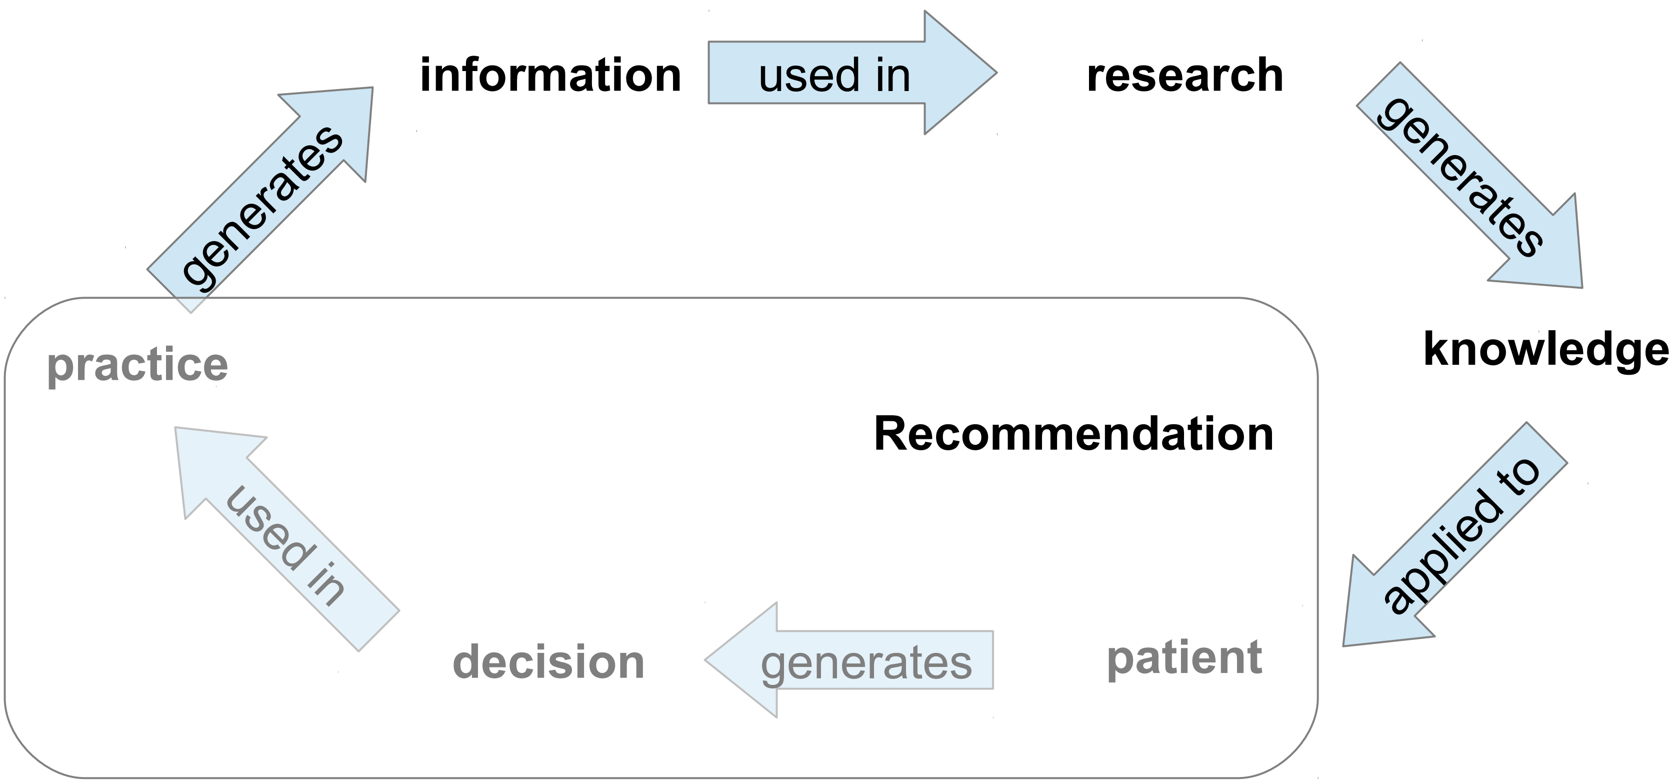

Clinical information/knowledge/decision

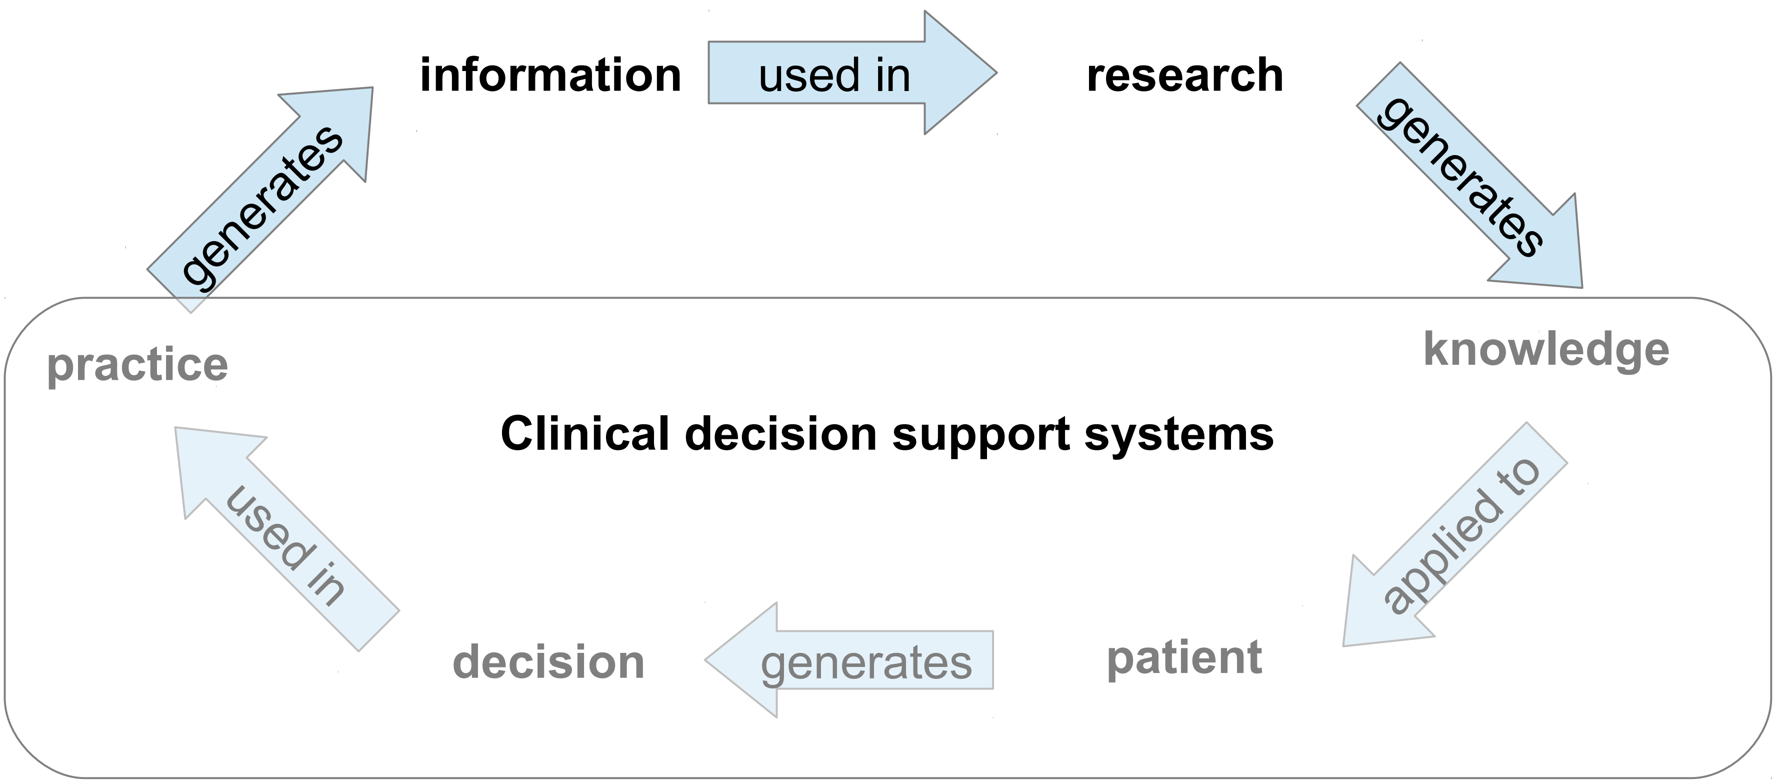

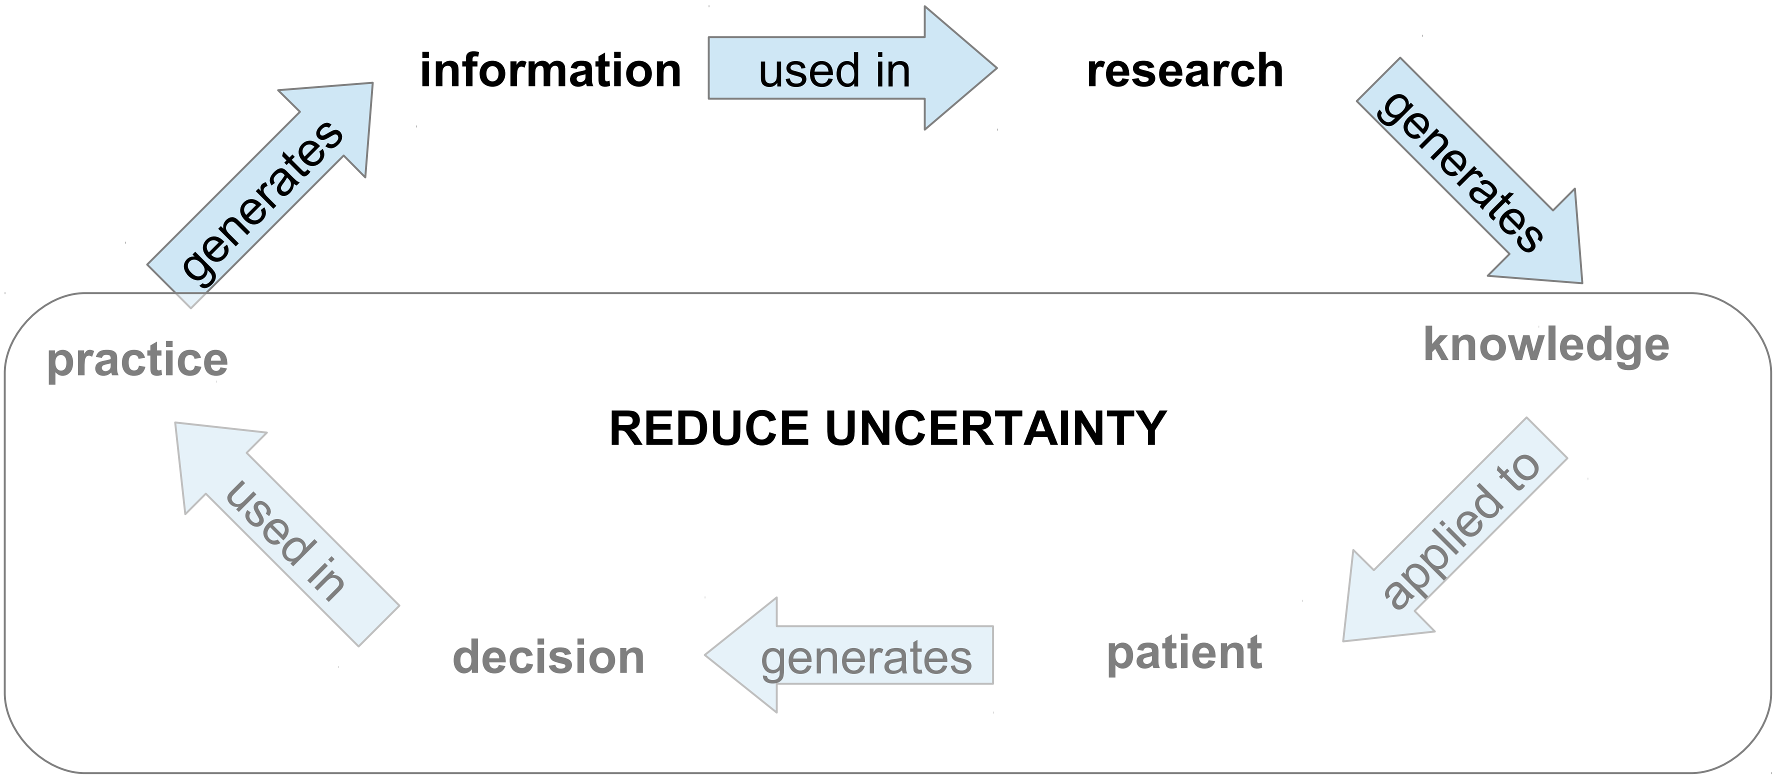
information
used in
research
generates
generates
practice
knowledge
REDUCE UNCERTAINTY
used in
applied to
decision
generates
patient

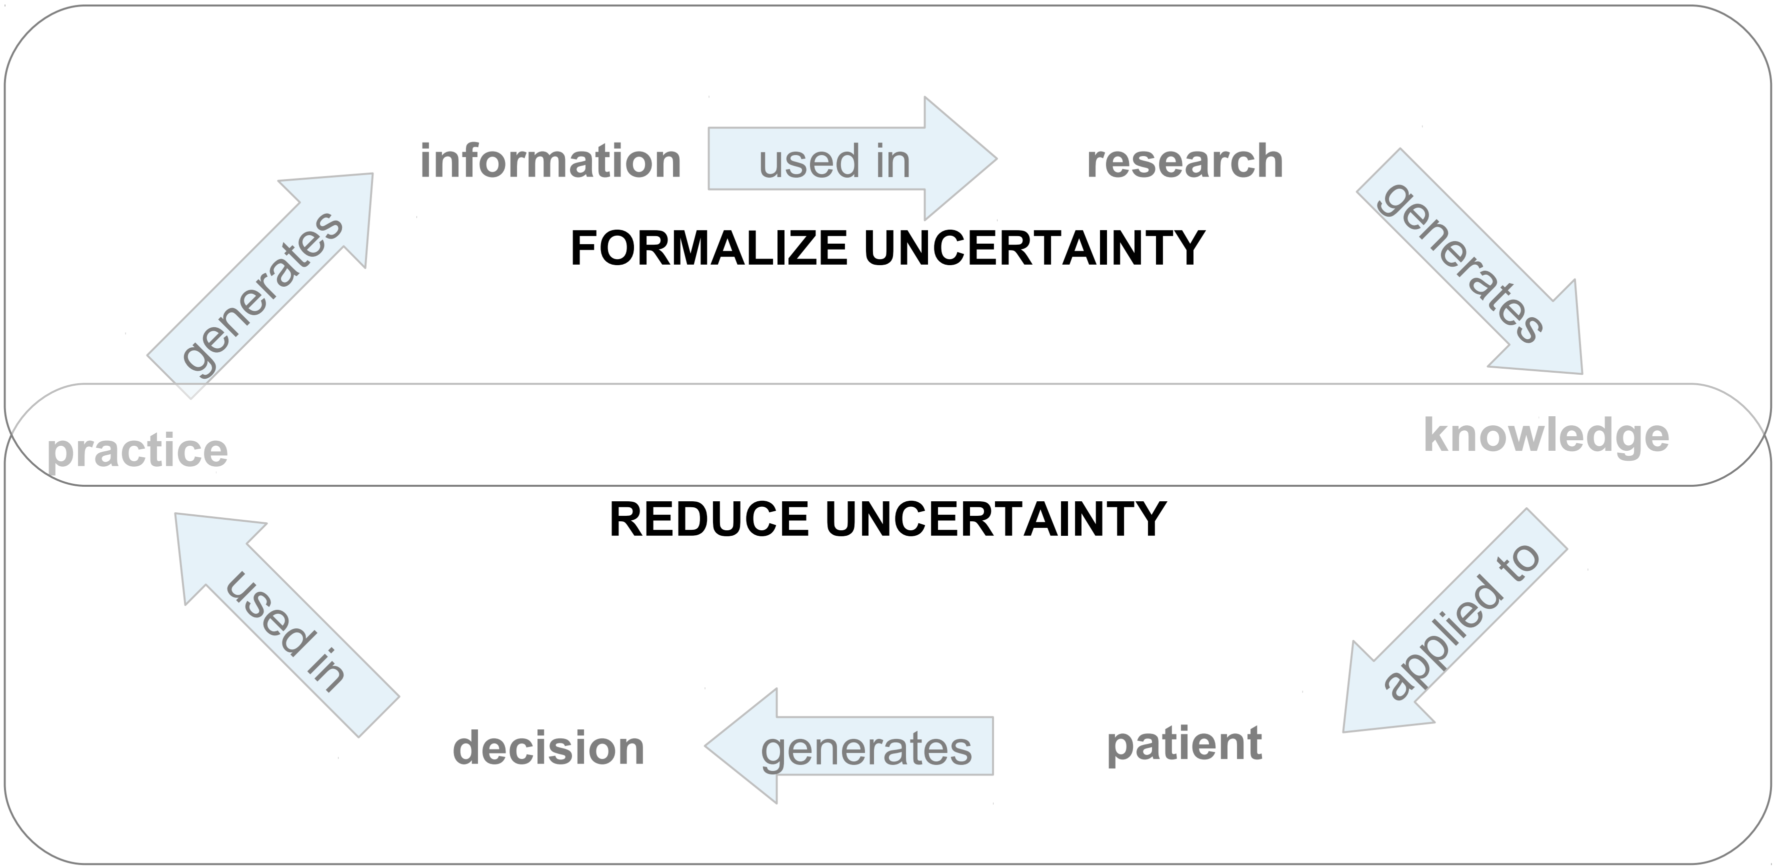
information
used in
research
generates
FORMALIZE UNCERTAINTY
generates
practice
knowledge
REDUCE UNCERTAINTY
used in
applied to
decision
generates
patient

Uncertainty and probability

We use terms such as **frequent, possible** or **rare** to express uncertainty.

Probability is a numeric expression of the likelihood that an event will occur.

We can then use probability to **express uncertainty without ambiguity**...

… and compute the efect of new information in the probability of disease, using the Bayes theorem.

Knowledge modeling for decision support

Risk and predictive factors

To support clinical decisions, we need to define:

Outcome - result variable (diagnosis, prognosis, treatment, etc.)

Factors - associated with the outcome (clinical history, demographic, etc.)

- Risk (of developing the disease or worse prognosis)
- Prediction (useful to predict but not necessarily of risk)

Association between factors and outcome

D. Bowers, A. House, and D. Owens, Understanding clinical papers. 2006.

Prevalence/Incidence

P = (a+c) / n

Risk ratio

RR = a/(a+b) / c/(c+d) = a(c+d)/c(a+b)

Odds ratio

OR = exposition odds (cases) / exposition odds (controls) = (a/c) / (b/d) = (ad) / (bc)

Sensitivity and specificity of factor as predictor of outcome

Sens = a / (a+c), Spec = d / (b+d)

		Outcome		
		Yes	No	Total
Factor	Yes	a	b	a+b
	No	c	d	c+d
	Total	a+c	b+d	n

A. Petrie and C. Sabin, Medical statistics at a glance. Blackwell, 2009, p. 180 pages.

Prevalence/Incidence

P = (a+c) / n

Risk ratio

RR = a/(a+b) / c/(c+d) = a(c+d)/c(a+b)

Odds ratio

OR = exposition odds (cases) / exposition odds (controls) = (a/c) / (b/d) = (ad) / (bc)

Sensitivity and specificity of factor as predictor of outcome

Sens = a / (a+c), Spec = d / (b+d)

		Outcome		
		Yes	No	Total
Factor	Yes	a	b	a+b
	No	c	d	c+d
	Total	a+c	b+d	n

These can all be interpreted as

(ratios of) conditional probabilities…

A. Petrie and C. Sabin, Medical statistics at a glance. Blackwell, 2009, p. 180 pages.

Clinical decision support

Evidence-based medicine relies on these simple, yet powerful, statistical measures as means for **evidence assessment**, yielding:

- Easy computation
- Formal representation of uncertainty (probability-based)
- Human-interpretable evidence

(e.g. RR > 1 means increased risk for exposed individuals compared to non-exposed ones)

Knowledge modeling

"The complicated nature of real-world biomedical data has made it necessary to look beyond traditional biostatistics."

"Bayesian statistical methods allow taking into account prior knowledge when analyzing data, turning the data analysis a process of updating that prior knowledge with biomedical and health-care evidence."

Peter Lucas (2004) Current Opinion in Critical Care

P. Lucas, "Bayesian analysis, pattern analysis, and data mining in health care.," Curr. Opin. Crit. Care, vol. 10, no. 5, pp. 399–403, Oct. 2004.

Knowledge modeling

"Bayesian networks offer a general and versatile approach to capturing and reasoning with uncertainty in medicine and health care."

Peter Lucas et al. (2004) *Artificial Intelligence In Medicine*

P. J. F. Lucas, L. C. van der Gaag, and A. Abu-Hanna, "Bayesian networks in biomedicine and health-care," Artif. Intell. Med., vol. 30, no. 3, pp. 201–14, 2004.

Graph representation where:

the attributes are represented by the graph **nodes**, and

the **arcs** represent dependencies among attributes,

using **conditional probabilities**.

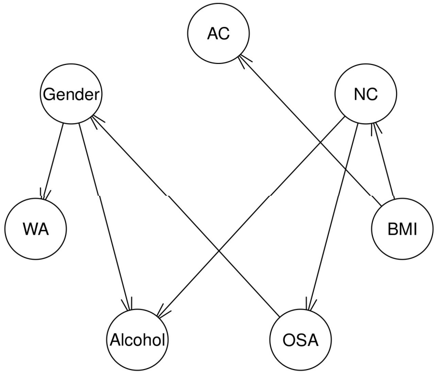

Easily human-interpretable representation, since it uses a

probabilistic reasoning similar to the usual uncertainty in human reasoning.

T. M. Mitchell, Machine Learning. McGraw-Hill, 1997.

D. Poole, A. Mackworth, and R. Goebel, Computational Intelligence: A Logical Approach. Oxford University Press, 1998.

Bayesian networks for clinical decision support

Bayesian networks **intrinsic uncertainty modeling** yields:

- Qualitative interpretation of **associations**
- Formal representation of **uncertainty** (probability-based)
- Human-interpretable **evidence** (a priori risk, a posteriori risk, relative risk, ...)
- Similar to traditional **biostatistics** (remember how measures are based on probabilities?)
- Decision support even with **unobserved variables**.

Complex **research questions** can be addressed by the same model:

Etiology and risk

Can a visit to China be the cause of patient's SARS?

Can a visit to China (and corresponding acquired SARS) be the cause of patient's dyspnea?

Diagnosis

The patient visited China; does he have SARS?

The patient has a high temperature reading; is it SARS?

Prognosis

The patient has fever and has visited China; without treatment, is he going to develop dyspnea?

Bayesian networks for clinical decision support

Sample of real examples

Content suppressed due to copyright constraints

2000

24h-prognosis of head-injured ICU patients

G. . Sakellaropoulos and G. . Nikiforidis, “Prognostic performance of two expert systems based on Bayesian belief networks,” Decis. Support Syst., vol. 27, no. 4, pp. 431–442, Jan. 2000.

Content suppressed due to copyright constraints

2005

Diagnosis of

ventilator-associated pneumonia

C. A. M. Schurink, P. J. F. Lucas, I. M. Hoepelman, and M. J. M. Bonten, "Computer-assisted decision support for the diagnosis and treatment of infectious diseases in intensive care units.," Lancet Infect. Dis., vol. 5, no. 5, pp. 305–12, May 2005.

Content suppressed due to copyright constraints

2008

Predicting maintenance

fluid requirement in ICU

L. A. Celi, L. C. Hinske, G. Alterovitz, and P. Szolovits, “An artificial intelligence tool to predict fluid requirement in the intensive care unit: a proof-of-concept study,” Crit. Care, vol. 12, no. 6, p. R151, Jan. 2008.

Content suppressed due to copyright constraints

2013

Breast cancer diagnosis

C.-R. Nicandro, M.-M. Efrén, A.-A. María Yaneli, M.-D.-C.-M. Enrique, A.-M. Héctor Gabriel, P.-C. Nancy, G.-H. Alejandro, H.-R. Guillermo de Jesús, and B.-M. Rocío Erandi, "Evaluation of the diagnostic power of thermography in breast cancer using Bayesian network classifiers.," Comput. Math. Methods Med., vol. 2013, p. 264246, Jan. 2013.

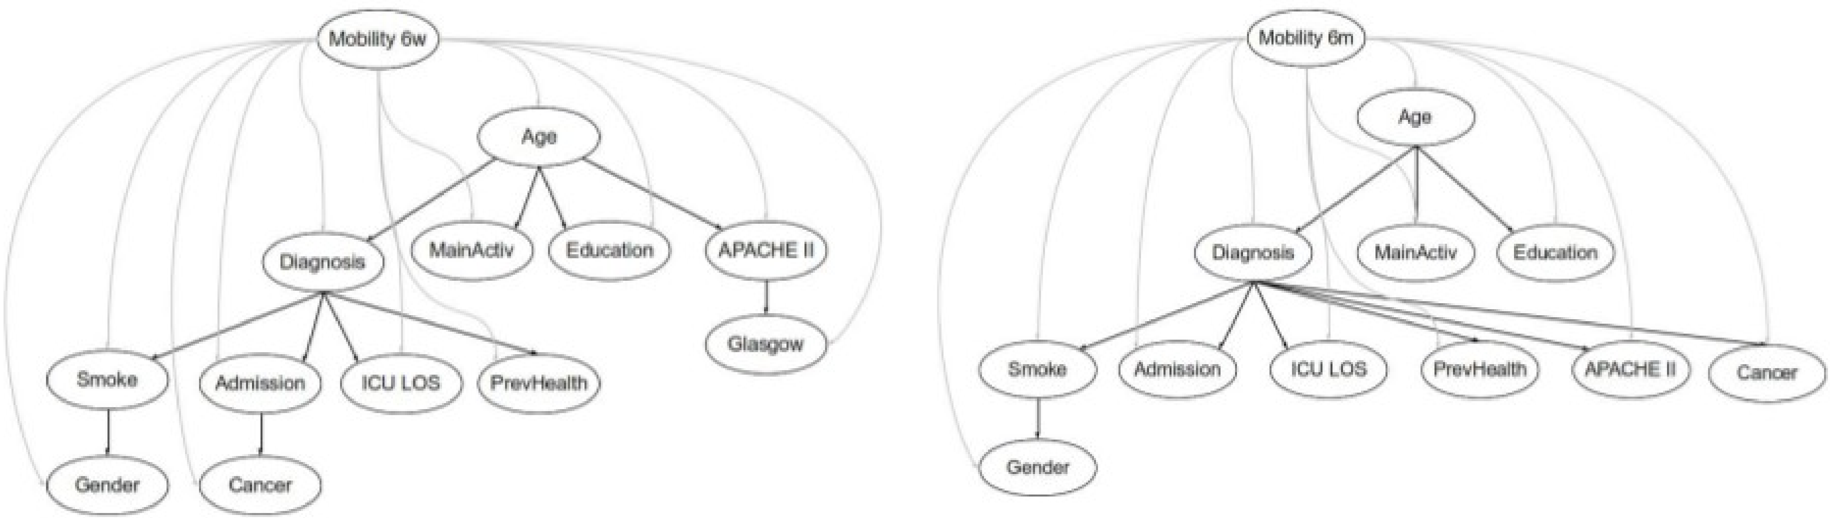

2014

Prognosis of quality of life after ICU stay

C. C. Dias, C. Granja, A. Costa-Pereira, J. Gama, and P. P. Rodrigues, "Using probabilistic graphical models to enhance the prognosis of health-related quality of life in adult survivors of critical illness," in 2014 IEEE 27th International Symposium on Computer-Based Medical Systems, 2014, pp. 56–61.

Bayesian networks for clinical decision support

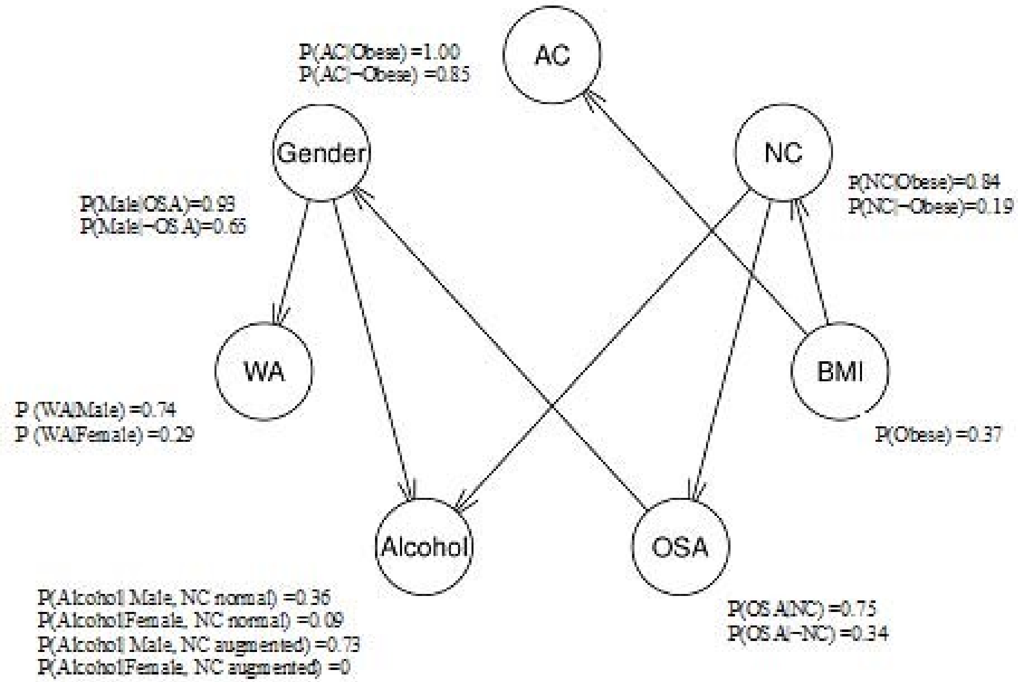

2014

Obstructive sleep apnea diagnosis

L. Leite, C. Costa-Santos, and P. P. Rodrigues, “Can we avoid unnecessary polysomnographies in the diagnosis of Obstructive Sleep Apnea? A Bayesian network decision support tool,” in 2014 IEEE 27th International Symposium on Computer-Based Medical Systems, 2014, pp. 28–33.

Content suppressed due to copyright constraints

2014

Temporal modeling of preeclampsia diagnosis

M. Velikova, J. T. van Scheltinga, P. J. F. Lucas, and M. Spaanderman, “Exploiting causal functional relationships in Bayesian network modelling for personalised healthcare,” Int. J. Approx. Reason., vol. 55, no. 1, pp. 59–73, Jan. 2014.

Take away message

M1: During inference and decision support, uncertainty needs to be reduced.

S1: Better focus on the variables that reduce uncertainty the most (e.g. when sugesting a test).

M2: Bayesian models (e.g. networks) are intrinsically modeling uncertainty and can map biostatistics.

S2: Consider Bayesian networks (or other probabilistic methods) as models to support clinical decision.

Uncertainty in Modeling

A toy example

Uncertainty in modeling

- You have access to a data set obtained from a cohort of suspected SARS patients, with one of the available variables being “Fever”.
- You learn from your data that “Fever” is associated with SARS.

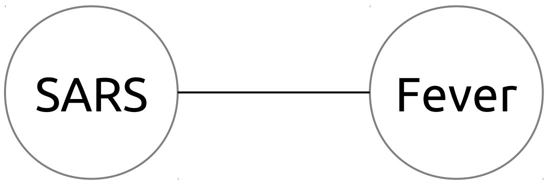

- Based on expert-knowledge you turn the association into causation.

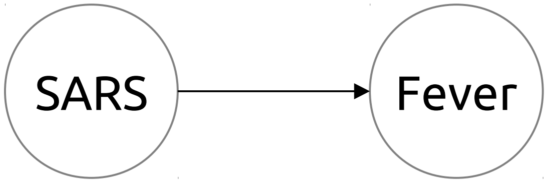

Uncertainty in modeling

- But the problem lingers:
  - what does “Fever” mean?
  - is it really observed?
- Although unlikely, you may have a reading of less than 37.5º and still have fever (e.g. if controlled with ibuprofen) or a reading of more than 37.5º without actuallly having fever.
- So, we should not reduce that uncertainty during modeling, rather include it in the model:

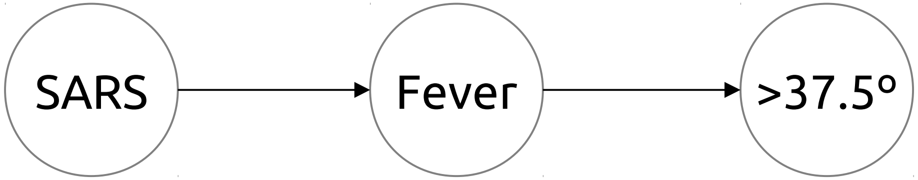

Take away message

M1: During inference and decision support, uncertainty needs to be reduced.

S1: Better focus on the variables that reduce uncertainty the most (e.g. when sugesting a test).

M2: Bayesian models (e.g. networks) are intrinsically modeling uncertainty and can map biostatistics.

S2: Consider Bayesian networks (or other probabilistic methods) as models to support clinical decision.

M3: If what you observe is what you record, it should also be what you model.

S3: Better search for the actual meaning (e.g. model temp above 37.5 instead of / along with fever).

Uncertainty in Modeling

A simple but real example

- There are cases where the knowledge discovery process needs to be merged with expert-based modeling and associations gathered from traditional meta-analysis.
- Imagine modeling the association between pneumonia and HIV infeccion, using a Bayesian net.

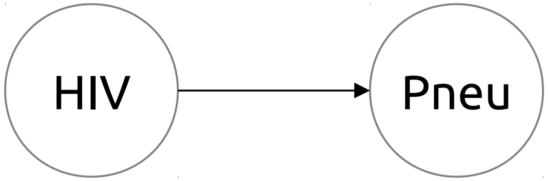

- The MD presents you a meta-analysis where this association is assessed and confirmed.
- So you can even use the meta-analysis risk assessment to compute the conditional probabilities of your Bayesian net (expert knowledge).

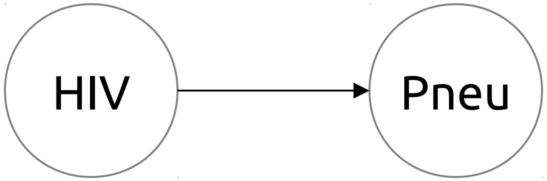

- You now have access to a database and, after the knowledge discovery process, it reveals the same association, so you consider merging the two data sources.
- But the variable HIV in your data is, in fact, given by the application of a standard test (for ilustrative purposes, lets consider PCR with 98% sensitivity and 99% specificity).
- So what you end up learning is the association between pneumonia and a positive PCR test result, which is an uncertain expression of HIV (precision may be below 10% for low disease prevalences)...

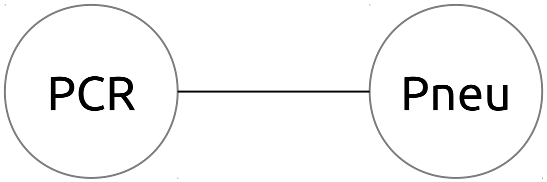

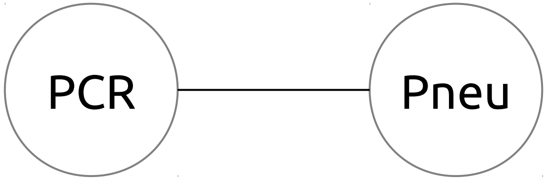

- But you have information on the association between the standard test and HIV infeccion...
  (remember that PCR has 98% sensitivity and 99% specificity)

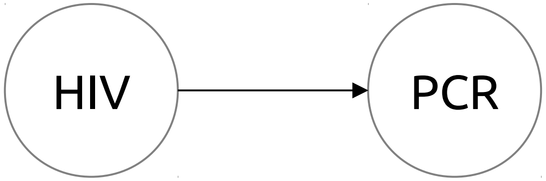

- So the model seems a bit more accurate now...

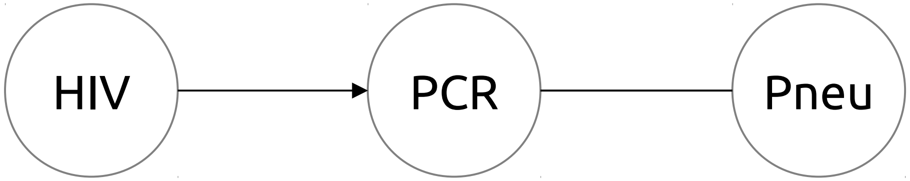

Uncertainty in modeling with expert knowledge

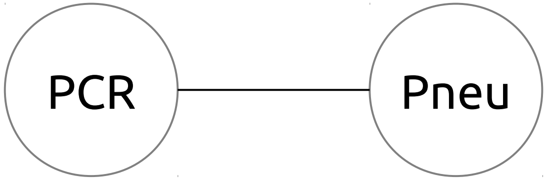

- But you have information on the association between the standard test and HIV infeccion...
  (remember that PCR has 98% sensitivity and 99% specificity)

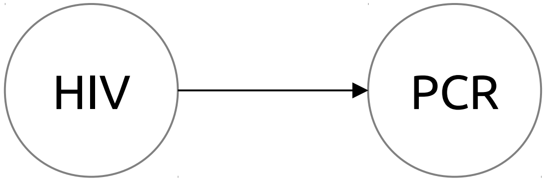

- So the model seems a bit more accurate now...

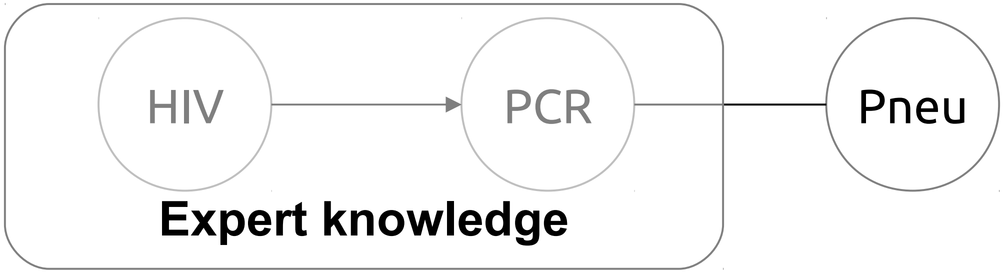

Uncertainty in modeling with expert knowledge

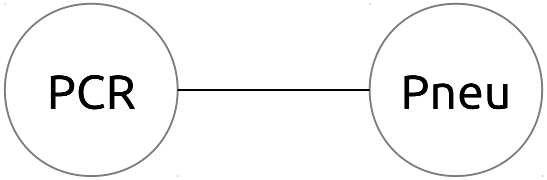

- But you have information on the association between the standard test and HIV infeccion...
  (remember that PCR has 98% sensitivity and 99% specificity)

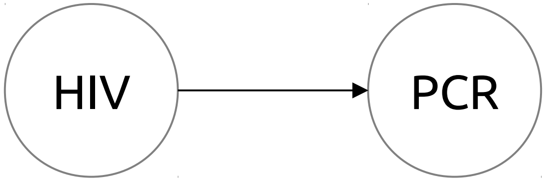

- So the model seems a bit more accurate now...

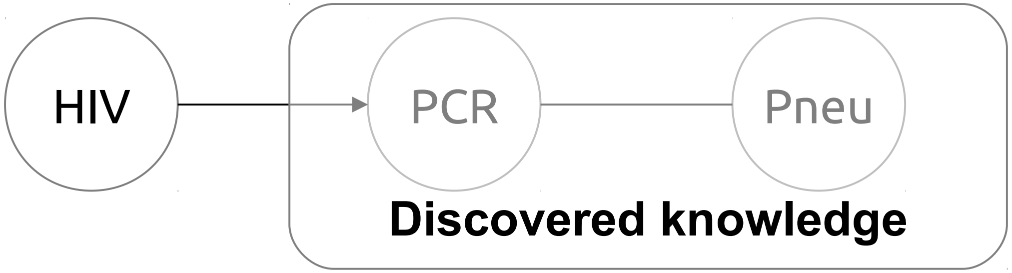

Uncertainty in modeling with expert knowledge

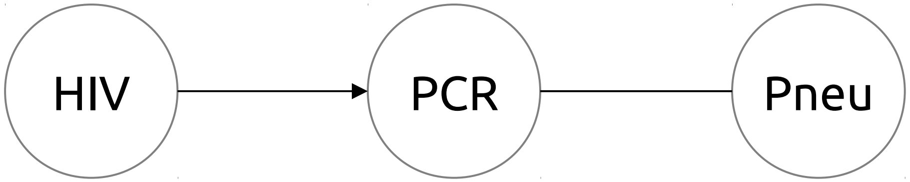

But your expert opinion tells you that is not the PCR test that is associated with pneumonia; it's the HIV infeccion, so it should look like this, instead:

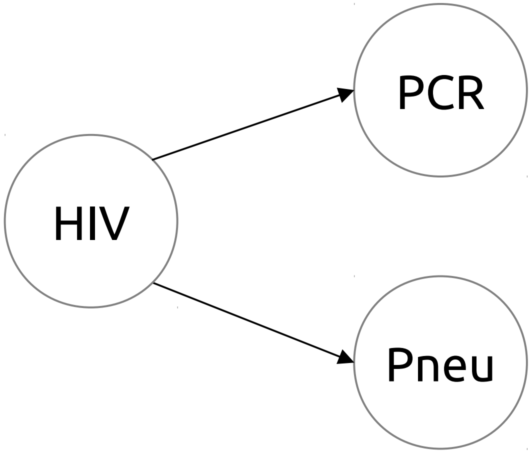

Uncertainty in modeling with expert knowledge

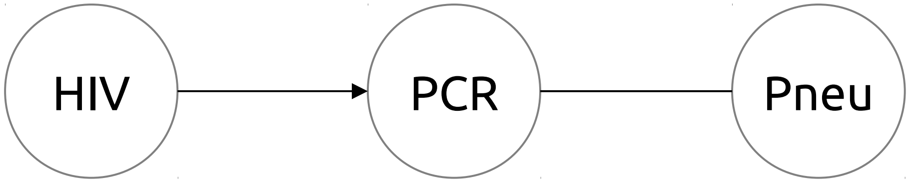

But your expert opinion tells you that is not the PCR test that is associated with pneumonia; it's the HIV infeccion, so it should look like this, instead:

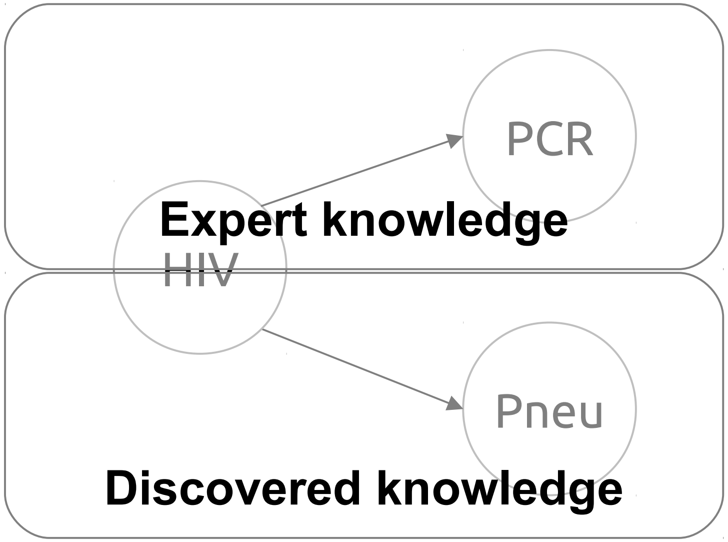

If what you observe is what you record, it should also be what you model.

Take away message

M1: During inference and decision support, uncertainty needs to be reduced.

S1: Better focus on the variables that reduce uncertainty the most (e.g. when sugesting a test).

M2: Bayesian models (e.g. networks) are intrinsically modeling uncertainty and can map biostatistics.

S2: Consider Bayesian networks (or other probabilistic methods) as models to support clinical decision.

M3: If what you observe is what you record, it should also be what you model.

S3: Better search for the actual meaning (e.g. model temp above 37.5 instead of / along with fever).

M4: During modeling and knowledge discovery, uncertainty needs to be formalized, not ignored.

S4: Better not dismiss variables' association that include uncertainty (e.g. do not assume PCR=HIV)

Thank you!

Acknowledgements

Cristina Granja

Altamiro Costa-Pereira

Cláudia Camila Dias

Liliana Leite

Cristina Costa-Santos

João Gama

References

D. Owens and H. Sox, “Biomedical decision making: probabilistic clinical reasoning,” in Biomedical Informatics, Chapter 3, Springer Verlag, 2006, pp. 80–132.

D. L. Sackett, W. M. Rosenberg, J. A. Gray, R. B. Haynes, and W. S. Richardson, “Evidence based medicine: what it is and what it isn’t.,” BMJ, vol. 312, no. 7023, pp. 71–2, Jan. 1996.

D. Bowers, A. House, and D. Owens, Understanding clinical papers. 2006.

A. Petrie and C. Sabin, Medical statistics at a glance. Blackwell, 2009, p. 180 pages.

P. Lucas, “Bayesian analysis, pattern analysis, and data mining in health care.,” Curr. Opin. Crit. Care, vol. 10, no. 5, pp. 399–403, Oct. 2004.

P. J. F. Lucas, L. C. van der Gaag, and A. Abu-Hanna, “Bayesian networks in biomedicine and health-care,” Artif. Intell. Med., vol. 30, no. 3, pp. 201–14, 2004.

T. M. Mitchell, Machine Learning. McGraw-Hill, 1997.

D. Poole, A. Mackworth, and R. Goebel, Computational Intelligence: A Logical Approach. Oxford University Press, 1998.

References (examples of Bayesian network applications)

M. Velikova, J. T. van Scheltinga, P. J. F. Lucas, and M. Spaanderman, “Exploiting causal functional relationships in Bayesian network modelling for personalised healthcare,” Int. J. Approx. Reason., vol. 55, no. 1, pp. 59–73, Jan. 2014.

L. Leite, C. Costa-Santos, and P. P. Rodrigues, “Can we avoid unnecessary polysomnographies in the diagnosis of Obstructive Sleep Apnea? A Bayesian network decision support tool,” in 2014 IEEE 27th International Symposium on Computer-Based Medical Systems, 2014, pp. 28–33.

C. C. Dias, C. Granja, A. Costa-Pereira, J. Gama, and P. P. Rodrigues, “Using probabilistic graphical models to enhance the prognosis of health-related quality of life in adult survivors of critical illness,” in 2014 IEEE 27th International Symposium on Computer-Based Medical Systems, 2014, pp. 56–61.

C.-R. Nicandro, M.-M. Efrén, A.-A. María Yaneli, M.-D.-C.-M. Enrique, A.-M. Héctor Gabriel, P.-C. Nancy, G.-H. Alejandro, H.-R. Guillermo de Jesús, and B.-M. Rocío Erandi, “Evaluation of the diagnostic power of thermography in breast cancer using Bayesian network classifiers.,” Comput. Math. Methods Med., vol. 2013, p. 264246, Jan. 2013.

L. A. Celi, L. C. Hinske, G. Alterovitz, and P. Szolovits, “An artificial intelligence tool to predict fluid requirement in the intensive care unit: a proof-of-concept study,” Crit. Care, vol. 12, no. 6, p. R151, Jan. 2008.

C. A. M. Schurink, P. J. F. Lucas, I. M. Hoepelman, and M. J. M. Bonten, “Computer-assisted decision support for the diagnosis and treatment of infectious diseases in intensive care units.,” Lancet Infect. Dis., vol. 5, no. 5, pp. 305–12, May 2005.

G. . Sakellaropoulos and G. . Nikiforidis, “Prognostic performance of two expert systems based on Bayesian belief networks,” Decis. Support Syst., vol. 27, no. 4, pp. 431–442, Jan. 2000.