Pure Inductive Logic

Jeff Paris

School of Mathematics, University of Manchester

– in collaboration with Jürgen Landes, Chris Nix, Alena Vencovská


gettyimages

Pure Inductive Logic Framework

Imagine an agent inhabiting a structure M for a first order language L with just finitely many relation symbols

$$P(x), P_1(x), P_2(x), R(x, y) \ldots \text{ etc.}$$

and countably constant symbols $a_1, a_2, a_3, \ldots$ which name every individual in the universe, and no function symbols nor equality.

This agent is assumed to have no further knowledge about M

Let SL denote the set of first order sentences of L.

Pure Inductive Logic Framework

Imagine an agent inhabiting a structure M for a first order language L with just finitely many relation symbols

$$P(x), P_1(x), P_2(x), R(x, y) \ldots \text{ etc.}$$

and countably constant symbols $a_1, a_2, a_3, \ldots$ which name every individual in the universe, and no function symbols nor equality.

This agent is assumed to have no further knowledge about M

Let SL denote the set of first order sentences of L.

Pure Inductive Logic Framework

Imagine an agent inhabiting a structure M for a first order language L with just finitely many relation symbols

$$P(x), P_1(x), P_2(x), R(x,y) \ldots \text{ etc.}$$

and countably constant symbols $a_1, a_2, a_3, \ldots$ which name every individual in the universe, and no function symbols nor equality.

This agent is assumed to have no further knowledge about M

Let SL denote the set of first order sentences of L.

Pure Inductive Logic Framework

Imagine an agent inhabiting a structure M for a first order language L with just finitely many relation symbols

$$P(x), P_1(x), P_2(x), R(x,y) \ldots \text{ etc.}$$

and countably constant symbols $a_1, a_2, a_3, \ldots$ which name every individual in the universe, and no function symbols nor equality.

This agent is assumed to have no further knowledge about M

Let SL denote the set of first order sentences of L.

We ask our agent to 'rationally' assign a probability $w(\theta)$ to $\theta \in SL$ being true in this ambient structure M.

Equivalently we're asking the agent to *pick* a 'rational' *probability function* w, where

$w : SL \to [0, 1]$ is a *probability function* on L if it satisfies

(P1) $\models \theta \;\Rightarrow\; w(\theta) = 1$

(P2) $\theta \models \neg\phi \;\Rightarrow\; w(\theta \vee \phi) = w(\theta) + w(\phi)$

(P3) $w(\exists x\, \psi(x)) = \lim_{n\to\infty} w\left(\bigvee_{i=1}^{n} \psi(a_i)\right)$

We ask our agent to 'rationally' assign a probability $w(\theta)$ to $\theta \in SL$ being true in this ambient structure M.

Equivalently we're asking the agent to *pick* a 'rational' *probability function* w, where

$w : SL \to [0,1]$ is a *probability function* on L if it satisfies

(P1) $\models \theta \;\Rightarrow\; w(\theta) = 1$

(P2) $\theta \models \neg\phi \;\Rightarrow\; w(\theta \vee \phi) = w(\theta) + w(\phi)$

(P3) $w(\exists x\, \psi(x)) = \lim_{n\to\infty} w\left(\bigvee_{i=1}^{n} \psi(a_i)\right)$

We ask our agent to 'rationally' assign a probability $w(\theta)$ to $\theta \in SL$ being true in this ambient structure M.

Equivalently we're asking the agent to *pick* a 'rational' *probability function* w, where

$w : SL \to [0,1]$ is a *probability function* on L if it satisfies

(P1) $\models \theta \;\Rightarrow\; w(\theta) = 1$

(P2) $\theta \models \neg\phi \;\Rightarrow\; w(\theta \vee \phi) = w(\theta) + w(\phi)$

(P3) $w(\exists x\, \psi(x)) = \lim_{n\to\infty} w\left(\bigvee_{i=1}^{n} \psi(a_i)\right)$

How should shehe do it?

By the application of 'rational principles' . . .

. . . based on Symmetry, Relevance, Irrelevance, Analogy, . . .

Example

Constant Exchangeability Principle, Ex

For $\theta(x_1, x_2, \dots, x_n)$ a formula of L not mentioning any constants

$$w(\theta(a_{i_1}, a_{i_2}, \dots, a_{i_n})) = w(\theta(a_{j_1}, a_{j_2}, \dots, a_{j_n}))$$

Similarly replacing a relation symbol R everywhere in $\phi \in SL$ by $\neg R$ should not change the probability (as in the coin toss example) – the *Strong Negation Principle*

By the application of 'rational principles' . . .

. . . based on Symmetry, Relevance, Irrelevance, Analogy, . . .

Example

Constant Exchangeability Principle, Ex

For $\theta(x_1, x_2, \ldots, x_n)$ a formula of L not mentioning any constants

$$w(\theta(a_{i_1}, a_{i_2}, \ldots, a_{i_n})) = w(\theta(a_{j_1}, a_{j_2}, \ldots, a_{j_n}))$$

Similarly replacing a relation symbol R everywhere in $\phi \in SL$ by $\neg R$ should not change the probability (as in the coin toss example) – the *Strong Negation Principle*

How should shehe do it?

By the application of ‘rational principles’ . . .

. . . based on Symmetry, Relevance, Irrelevance, Analogy, . . .

Example

Constant Exchangeability Principle, Ex

For $\theta(x_1, x_2, \ldots, x_n)$ a formula of L not mentioning any constants

$$w(\theta(a_{i_1}, a_{i_2}, \ldots, a_{i_n})) = w(\theta(a_{j_1}, a_{j_2}, \ldots, a_{j_n}))$$

Similarly replacing a relation symbol R everywhere in $\phi \in SL$ by $\neg R$ should not change the probability (as in the coin toss example) – the *Strong Negation Principle*

By the application of 'rational principles' . . .

. . . based on Symmetry, Relevance, Irrelevance, Analogy, . . .

Example

Constant Exchangeability Principle, Ex

For $\theta(x_1, x_2, \ldots, x_n)$ *a formula of L not mentioning any constants*

$$w(\theta(a_{i_1}, a_{i_2}, \ldots, a_{i_n})) = w(\theta(a_{j_1}, a_{j_2}, \ldots, a_{j_n}))$$

Similarly replacing a relation symbol R everywhere in $\phi \in SL$ by $\neg R$ should not change the probability (as in the coin toss example) – the *Strong Negation Principle*

By the application of 'rational principles' . . .

. . . based on Symmetry, Relevance, Irrelevance, Analogy, . . .

Example

Constant Exchangeability Principle, Ex

For $\theta(x_1, x_2, \ldots, x_n)$ *a formula of L not mentioning any constants*

$$w(\theta(a_{i_1}, a_{i_2}, \ldots, a_{i_n})) = w(\theta(a_{j_1}, a_{j_2}, \ldots, a_{j_n}))$$

Similarly replacing a relation symbol R everywhere in $\phi \in SL$ by $\neg R$ should not change the probability (as in the coin toss example) – the *Strong Negation Principle*

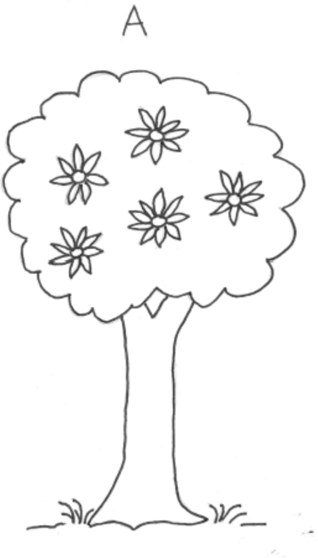
A
B

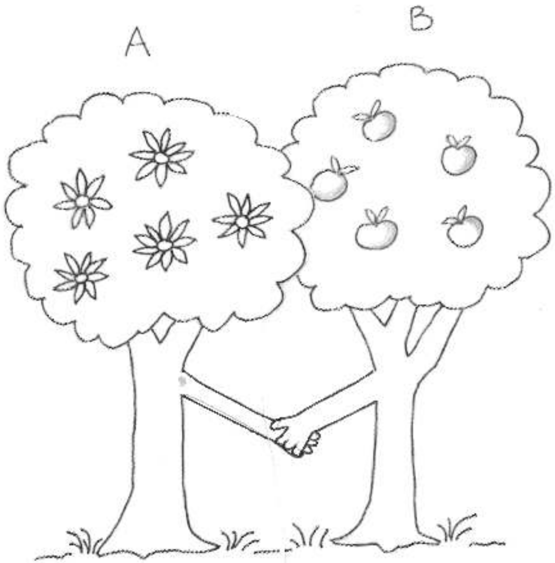
A
B

Such intuitions however are easily challenged, e.g.

Given

$$R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)$$

which of $R(a_3, a_1), \neg R(a_3, a_1)$ *would you think the more likely?*

Such intuitions however are easily challenged, e.g.

Given

$$R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)$$

which of $R(a_3, a_1), \neg R(a_3, a_1)$ would you think the more likely?

Such intuitions however are easily challenged, e.g.

Given

$$R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)$$

which of $R(a_3, a_1), \neg R(a_3, a_1)$ *would you think the more likely?*

For simplicity assume that L has just a single binary relation symbol R.

A *state description for* $a_1, a_2, \ldots, a_n$ is a quantifier free sentence of the form

$$\bigwedge_{i,j=1}^{n} \pm R(a_i, a_j)$$

State descriptions are where it all happens in this subject because:-

Gaifman's Theorem

w is completely determined by its values on state descriptions.

For simplicity assume that L has just a single binary relation symbol R.

A *state description for* $a_1, a_2, \ldots, a_n$ is a quantifier free sentence of the form

$$\bigwedge_{i,j=1}^{n} \pm R(a_i, a_j)$$

State descriptions are where it all happens in this subject because:-

Gaifman's Theorem

w is completely determined by its values on state descriptions.

For simplicity assume that L has just a single binary relation symbol R.

A *state description for* $a_1, a_2, \ldots, a_n$ is a quantifier free sentence of the form

$$\bigwedge_{i,j=1}^{n} \pm R(a_i, a_j)$$

State descriptions are where it all happens in this subject because:-

Gaifman's Theorem

w is completely determined by its values on state descriptions.

For simplicity assume that L has just a single binary relation symbol R.

A *state description for* $a_1, a_2, \ldots, a_n$ is a quantifier free sentence of the form

$$\bigwedge_{i,j=1}^{n} \pm R(a_i, a_j)$$

State descriptions are where it all happens in this subject because:-

Gaifman's Theorem

w is completely determined by its values on state descriptions.

Into the Polyadic

For simplicity assume that L has just a single binary relation symbol R.

A *state description for* $a_1, a_2, \ldots, a_n$ is a quantifier free sentence of the form

$$\bigwedge_{i,j=1}^{n} \pm R(a_i, a_j)$$

State descriptions are where it all happens in this subject because:-

Gaifman's Theorem

w is completely determined by its values on state descriptions.

The *Completely Independent Probability Function* w_0 gives each of the $\pm R(a_i, a_j)$ probability 1/2 and treats them all as stochastically independent

E.g.

$$\mathrm{w}_0(R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)) = (1/2) \times (1/2) \times (1/2) = 1/8$$

Trouble is, to our earlier question

$$\mathrm{w}_0(R(a_3, a_1) \mid R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)) = 1/2 =$$
$$\mathrm{w}_0(\neg R(a_3, a_1) \mid R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3))$$

The Completely Independent Probability Function

The *Completely Independent Probability Function* w_0 gives each of the $\pm R(a_i, a_j)$ probability $1/2$ and treats them all as stochastically independent

E.g.

$$\mathrm{w}_0(R(a_1,a_2) \wedge R(a_2,a_1) \wedge \neg R(a_1,a_3)) = (1/2) \times (1/2) \times (1/2) = 1/8$$

Trouble is, to our earlier question

$$\mathrm{w}_0(R(a_3,a_1) \,|\, R(a_1,a_2) \wedge R(a_2,a_1) \wedge \neg R(a_1,a_3)) = 1/2 =$$
$$\mathrm{w}_0(\neg R(a_3,a_1) \,|\, R(a_1,a_2) \wedge R(a_2,a_1) \wedge \neg R(a_1,a_3))$$

The *Completely Independent Probability Function* w_0 gives each of the $\pm R(a_i, a_j)$ probability $1/2$ and treats them all as stochastically independent

E.g.

$$\mathrm{w}_0(R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)) = (1/2) \times (1/2) \times (1/2) = 1/8$$

Trouble is, to our earlier question

$$\mathrm{w}_0(R(a_3, a_1) \mid R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)) = 1/2 =$$
$$\mathrm{w}_0(\neg R(a_3, a_1) \mid R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3))$$

The Completely Independent Probability Function

The *Completely Independent Probability Function* w_0 gives each of the $\pm R(a_i, a_j)$ probability $1/2$ and treats them all as stochastically independent

E.g.

$$\mathrm{w}_0(R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)) = (1/2) \times (1/2) \times (1/2) = 1/8$$

Trouble is, to our earlier question

$$\mathrm{w}_0(R(a_3, a_1) \mid R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)) = 1/2 = \mathrm{w}_0(\neg R(a_3, a_1) \mid R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3))$$

Spectrum Exchangeability

Given a state description $\Theta(a_1, a_2, \ldots, a_n)$ define the equivalence relation $\sim_\Theta$ on $\{a_1, \ldots, a_n\}$ by

$$a_i \sim_\Theta a_j \iff \Theta(a_1, a_2, \ldots, a_n) \wedge a_i = a_j \text{ is consistent}$$

equivalently iff a_i, a_j are indistinguishable on the basis of $\Theta(a_1, \ldots, a_n)$.

The *spectrum* of $\Theta(a_1, \ldots, a_n)$ is the multiset of sizes of the equivalence classes according to $\sim_\Theta$.

Spectrum Exchangeability

Given a state description $\Theta(a_1, a_2, \ldots, a_n)$ define the equivalence relation $\sim_\Theta$ on $\{a_1, \ldots, a_n\}$ by

$$a_i \sim_\Theta a_j \iff \Theta(a_1, a_2, \ldots, a_n) \wedge a_i = a_j \text{ is consistent}$$

equivalently iff a_i, a_j are indistinguishable on the basis of $\Theta(a_1, \ldots, a_n)$.

The *spectrum* of $\Theta(a_1, \ldots, a_n)$ is the multiset of sizes of the equivalence classes according to $\sim_\Theta$.

Spectrum Exchangeability

Given a state description $\Theta(a_1, a_2, \ldots, a_n)$ define the equivalence relation $\sim_\Theta$ on $\{a_1, \ldots, a_n\}$ by

$$a_i \sim_\Theta a_j \iff \Theta(a_1, a_2, \ldots, a_n) \wedge a_i = a_j \text{ is consistent}$$

equivalently iff a_i, a_j are indistinguishable on the basis of $\Theta(a_1, \ldots, a_n)$.

The *spectrum* of $\Theta(a_1, \ldots, a_n)$ is the multiset of sizes of the equivalence classes according to $\sim_\Theta$.

Given a state description $\Theta(a_1, a_2, \ldots, a_n)$ define the equivalence relation $\sim_\Theta$ on $\{a_1, \ldots, a_n\}$ by

$$a_i \sim_\Theta a_j \iff \Theta(a_1, a_2, \ldots, a_n) \wedge a_i = a_j \text{ is consistent}$$

equivalently iff a_i, a_j are indistinguishable on the basis of $\Theta(a_1, \ldots, a_n)$.

The *spectrum* of $\Theta(a_1, \ldots, a_n)$ is the multiset of sizes of the equivalence classes according to $\sim_\Theta$.

Example

Suppose $\Theta(a_1, a_2, a_3, a_4)$ is the conjunction of

$$\begin{array}{cccc} R(a_1, a_1) & \neg R(a_1, a_2) & R(a_1, a_3) & R(a_1, a_4) \\ R(a_2, a_1) & \neg R(a_2, a_2) & R(a_2, a_3) & \neg R(a_2, a_4) \\ R(a_3, a_1) & \neg R(a_3, a_2) & R(a_3, a_3) & R(a_3, a_4) \\ R(a_4, a_1) & R(a_4, a_2) & R(a_4, a_3) & R(a_4, a_4) \end{array}$$

Then the equivalence classes are $\{a_1, a_3\}$, $\{a_2\}$, $\{a_4\}$ and the spectrum is

$$\{2, 1, 1\}$$

Spectrum Exchangeability, Sx

If the state descriptions $\Theta(a_1, \ldots, a_n), \Phi(a_1, \ldots, a_n)$ *have the same spectrum then*

$$w(\Theta(a_1, \ldots, a_n)) = w(\Phi(a_1, \ldots, a_n))$$

Spectrum Exchangeability, Sx

If the state descriptions $\Theta(a_1, \ldots, a_n), \Phi(a_1, \ldots, a_n)$ *have the same spectrum then*

$$w(\Theta(a_1, \ldots, a_n)) = w(\Phi(a_1, \ldots, a_n))$$

So the conjunctions of

$$\begin{array}{ccc} R(a_1,a_1) & \neg R(a_1,a_2) & R(a_1,a_3) \\ R(a_2,a_1) & \neg R(a_2,a_2) & R(a_2,a_3) \\ R(a_3,a_1) & \neg R(a_3,a_2) & R(a_3,a_3) \end{array}$$

and

$$\begin{array}{ccc} \neg R(a_1,a_1) & \neg R(a_1,a_2) & R(a_1,a_3) \\ \neg R(a_2,a_1) & \neg R(a_2,a_2) & R(a_2,a_3) \\ R(a_3,a_1) & R(a_3,a_2) & R(a_3,a_3) \end{array}$$

get the same probability as both have spectrum $\{2,1\}$

The Promised Land (?)

Given

$$R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)$$

which of $R(a_3, a_1), \neg R(a_3, a_1)$ *would you think the more likely?*

Sx implies that the $\neg R(a_3, a_1)$ is at least as likely as $R(a_3, a_1)$ (so 'analogy' wins out)

The Promised Land (?)

Given

$$R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)$$

which of $R(a_3, a_1), \neg R(a_3, a_1)$ *would you think the more likely?*

Sx implies that the $\neg R(a_3, a_1)$ is at least as likely as $R(a_3, a_1)$ (so 'analogy' wins out)

Given

$$R(a_1, a_2) \wedge R(a_2, a_1) \wedge \neg R(a_1, a_3)$$

which of $R(a_3, a_1), \neg R(a_3, a_1)$ *would you think the more likely?*

Sx implies that the $\neg R(a_3, a_1)$ is at least as likely as $R(a_3, a_1)$ (so 'analogy' wins out)

Conformity

Consider the two 'unary relations' $R(a_1, x)$ and $R(x, x)$ of L. Which of the two 'state descriptions'

$$R(a_1, a_1) \wedge R(a_1, a_2) \wedge \neg R(a_1, a_3) \wedge R(a_1, a_4)$$

$$R(a_1, a_1) \wedge R(a_2, a_2) \wedge \neg R(a_3, a_3) \wedge R(a_4, a_4)$$

should we think the more likely?

The intuition is that there is no rational reason why $R(a_1, x)$ and $R(x, x)$ should, in isolation, differ

Hence the above 'state descriptions' should get the same probability.

Assuming Sx they do!

Conformity

Consider the two 'unary relations' $R(a_1, x)$ and $R(x, x)$ of L. Which of the two 'state descriptions'

$$R(a_1, a_1) \wedge R(a_1, a_2) \wedge \neg R(a_1, a_3) \wedge R(a_1, a_4)$$

$$R(a_1, a_1) \wedge R(a_2, a_2) \wedge \neg R(a_3, a_3) \wedge R(a_4, a_4)$$

should we think the more likely?

The intuition is that there is no rational reason why $R(a_1, x)$ and $R(x, x)$ should, in isolation, differ

Hence the above 'state descriptions' should get the same probability.

Assuming Sx they do!

Conformity

Consider the two 'unary relations' $R(a_1, x)$ and $R(x, x)$ of L. Which of the two 'state descriptions'

$$R(a_1, a_1) \wedge R(a_1, a_2) \wedge \neg R(a_1, a_3) \wedge R(a_1, a_4)$$

$$R(a_1, a_1) \wedge R(a_2, a_2) \wedge \neg R(a_3, a_3) \wedge R(a_4, a_4)$$

should we think the more likely?

The intuition is that there is no rational reason why $R(a_1, x)$ and $R(x, x)$ should, in isolation, differ

Hence the above 'state descriptions' should get the same probability.

Assuming Sx they do!

Conformity

Consider the two 'unary relations' $R(a_1, x)$ and $R(x, x)$ of L. Which of the two 'state descriptions'

$$R(a_1, a_1) \wedge R(a_1, a_2) \wedge \neg R(a_1, a_3) \wedge R(a_1, a_4)$$

$$R(a_1, a_1) \wedge R(a_2, a_2) \wedge \neg R(a_3, a_3) \wedge R(a_4, a_4)$$

should we think the more likely?

The intuition is that there is no rational reason why $R(a_1, x)$ and $R(x, x)$ should, in isolation, differ

Hence the above 'state descriptions' should get the same probability.

Assuming Sx they do!

Consider the two 'unary relations' $R(a_1, x)$ and $R(x, x)$ of L. Which of the two 'state descriptions'

$$R(a_1, a_1) \wedge R(a_1, a_2) \wedge \neg R(a_1, a_3) \wedge R(a_1, a_4)$$

$$R(a_1, a_1) \wedge R(a_2, a_2) \wedge \neg R(a_3, a_3) \wedge R(a_4, a_4)$$

should we think the more likely?

The intuition is that there is no rational reason why $R(a_1, x)$ and $R(x, x)$ should, in isolation, differ

Hence the above 'state descriptions' should get the same probability.

Assuming Sx they do!

Inseparability

Suppose that w satisfies Sx and is not equal to w_0.

Then, given a state description $\Theta(a_1, a_2, \ldots, a_n)$ in which a_1, a_2 are indistinguishable (i.e. $a_1 \sim_\Theta a_2$) there is a non-zero probability according to w that they will remain forever indistinguishable.

BUT the probability according to w that a_1, a_2 will be forever indistinguishable but be distinguishable from each of $a_3, a_4, a_5, \ldots$ is zero

Suppose that w satisfies Sx and is not equal to w_0.

Then, given a state description $\Theta(a_1, a_2, \ldots, a_n)$ in which a_1, a_2 are indistinguishable (i.e. $a_1 \sim_\Theta a_2$) there is a non-zero probability according to w that they will remain forever indistinguishable.

BUT the probability according to w that a_1, a_2 will be forever indistinguishable but be distinguishable from each of $a_3, a_4, a_5, \ldots$ is zero

Inseparability

Suppose that w satisfies Sx and is not equal to w_0.

Then, given a state description $\Theta(a_1, a_2, \ldots, a_n)$ in which a_1, a_2 are indistinguishable (i.e. $a_1 \sim_\Theta a_2$) there is a non-zero probability according to w that they will remain forever indistinguishable.

BUT the probability according to w that a_1, a_2 will be forever indistinguishable but be distinguishable from each of $a_3, a_4, a_5, \ldots$ is zero

Inseparability

Suppose that w satisfies Sx and is not equal to w_0.

Then, given a state description $\Theta(a_1, a_2, \ldots, a_n)$ in which a_1, a_2 are indistinguishable (i.e. $a_1 \sim_\Theta a_2$) there is a non-zero probability according to w that they will remain forever indistinguishable.

BUT the probability according to w that a_1, a_2 will be forever indistinguishable but be distinguishable from each of $a_3, a_4, a_5, \ldots$ is zero

Suppose that w satisfies Sx and $\Theta(\vec{a})$ is the state description of $L' \subset L$ satisfied by $\vec{a}$. Then according to w the most probable state description(s) of L satisfied by $\vec{a}$ have the same spectrum as $\Theta(\vec{a})$.

Sx looks the business . . . but . . .

What is the rational justification for Sx ?

Restricted to unary languages Sx can be justified by 'symmetry'

But can Sx be justified by 'symmetry' in the polyadic?

What is the rational justification for Sx ?

Restricted to unary languages Sx can be justified by ‘symmetry’

But can Sx be justified by ‘symmetry’ in the polyadic?

What is the rational justification for Sx ?

Restricted to unary languages Sx can be justified by ‘symmetry’

But can Sx be justified by ‘symmetry’ in the polyadic?

Sx looks the business . . . but . . .

What is the rational justification for Sx ?

Restricted to unary languages Sx can be justified by ‘symmetry’

But can Sx be justified by ‘symmetry’ in the polyadic?