

Experimental Design for Efficient Identification of Gene Regulatory Networks using Sparse Bayesian Models

Florian Steinke, Matthias Seeger, Koji Tsuda

Max Planck Institute for Biological Cybernetics
Tübingen, Germany

26 July 2007



MAX-PLANCK-GESELLSCHAFT

- 1 The Need for Experimental Design
- 2 System Identification of Genetic Regulation
- 3 Sparse Bayesian Linear Model
- 4 Experiments
- 5 Related Work. Conclusions

Why Experimental Design?

- **Large-scale** genome-wide experiments:
Affordable today in fully automatized labs
- Solve problems by **complete enumeration** or **random shooting**?
 - **Guaranteed** to run out of steam on hard problems
 - **Cutting-edge** experiments always hard/expensive
 - Even for large labs: (#Results)/\$ counts!
- **Sequential Optimal Design**
Plan next experiment based on all previous outcomes
⇒ Every **smart biologist** does that anyway!
- Can optimal design be **semi-automatized** on a dumb machine?
What general framework allows us to do that?

Why Experimental Design?

- **Large-scale** genome-wide experiments:
Affordable today in fully automatized labs
- Solve problems by **complete enumeration** or **random shooting**?
 - **Guaranteed** to run out of steam on hard problems
 - **Cutting-edge** experiments always hard/expensive
 - Even for large labs: (#Results)/\$ counts!
- **Sequential Optimal Design**
Plan next experiment based on all previous outcomes
⇒ Every **smart biologist** does that anyway!
- Can optimal design be **semi-automatized** on a dumb machine?
What general framework allows us to do that?

Why Experimental Design?

- **Large-scale** genome-wide experiments:
Affordable today in fully automatized labs
- Solve problems by **complete enumeration** or **random shooting**?
 - **Guaranteed** to run out of steam on hard problems
 - **Cutting-edge** experiments always hard/expensive
 - Even for large labs: (#Results)/\$ counts!
- **Sequential Optimal Design**
Plan next experiment based on all previous outcomes
⇒ Every **smart biologist** does that anyway!
- Can optimal design be **semi-automatized** on a dumb machine?
What general framework allows us to do that?

Why Experimental Design?

- **Large-scale** genome-wide experiments:
Affordable today in fully automatized labs
- Solve problems by **complete enumeration** or **random shooting**?
 - **Guaranteed** to run out of steam on hard problems
 - **Cutting-edge** experiments always hard/expensive
 - Even for large labs: (#Results)/\$ counts!
- **Sequential Optimal Design**
Plan next experiment based on all previous outcomes
⇒ Every **smart biologist** does that anyway!
- Can optimal design be **semi-automatized** on a dumb machine?
What general framework allows us to do that?

Bayesian Optimal Design

Smart Biologist

- Which variables could explain my data? How could dependencies look like?
- X look well-determined.
Did not learn much about Y
- I think: Exp. A (B) would tell me more about X (Y) now
⇒ Of course I do B !
- 1000s of X, Y . Combinatorial number of possible interactions ⇒ **Human intuition**

Bayesian Framework

- **Model design**
Observed, hidden variables.
Dependency model
- **Posterior uncertainty**
Reduced on X , but not on Y
- **Information Gain Scores**
 $S(A; \text{Data}) < S(B; \text{Data})$
⇒ OK> Should do B
- Run overnight, sift through raw data, (hopefully) **help** intuition along

Bayesian Optimal Design

Smart Biologist

- Which variables could explain my data? How could dependencies look like?
- **X** look well-determined. Did not learn much about **Y**
- I think: Exp. **A** (**B**) would tell me more about **X** (**Y**) now
⇒ Of course I do **B**!
- 1000s of **X**, **Y**. Combinatorial number of possible interactions ⇒ **Human intuition**

Bayesian Framework

- **Model design**
Observed, hidden variables.
Dependency model
- **Posterior uncertainty**
Reduced on **X**, but not on **Y**
- **Information Gain Scores**
 $S(A; \text{Data}) < S(B; \text{Data})$
⇒ OK> Should do **B**
- Run overnight, sift through raw data, (hopefully) **help** intuition along

Bayesian Optimal Design

Smart Biologist

- Which variables could explain my data? How could dependencies look like?
- **X** look well-determined. Did not learn much about **Y**
- I think: Exp. **A** (**B**) would tell me more about **X** (**Y**) now
⇒ Of course I do **B**!
- 1000s of **X**, **Y**. Combinatorial number of possible interactions ⇒ **Human intuition**

Bayesian Framework

- **Model design**
Observed, hidden variables. Dependency model
- **Posterior uncertainty**
Reduced on **X**, but not on **Y**
- **Information Gain Scores**
 $S(A; \text{Data}) < S(B; \text{Data})$
⇒ OK> Should do **B**
- Run overnight, sift through raw data, (hopefully) **help** intuition along

Bayesian Optimal Design

Smart Biologist

- Which variables could explain my data? How could dependencies look like?
- **X** look well-determined. Did not learn much about **Y**
- I think: Exp. **A** (**B**) would tell me more about **X** (**Y**) now
⇒ Of course I do **B**!
- 1000s of **X**, **Y**. Combinatorial number of possible interactions ⇒ **Human intuition**

Bayesian Framework

- **Model design**
Observed, hidden variables.
Dependency model
- **Posterior uncertainty**
Reduced on **X**, but not on **Y**
- **Information Gain Scores**
 $S(A; \text{Data}) < S(B; \text{Data})$
⇒ OK> Should do **B**
- Run overnight, sift through raw data, (hopefully) **help** intuition along

Gene Regulatory Networks

- **Genes can regulate other genes**

Protein from gene A can be transcription factor:
up-/down-regulates transcription of gene B .

Causal link $A \rightarrow B$ in **gene regulatory network**

- **Affordable Measurements**

m-RNA concentrations (micro-arrays), protein concentrations
 \leftrightarrow Expression levels $x_A(t)$, $x_B(t)$

- **System Identification**

Interventionist. Disturb system (without breaking it).
Learn structure from changes in measurements

- **Optimal Experimental Design**

For given model: **Short(est)** sequences of experiments leading to identification?

Gene Regulatory Networks

- **Genes can regulate other genes**

Protein from gene A can be transcription factor:
up-/down-regulates transcription of gene B .

Causal link $A \rightarrow B$ in **gene regulatory network**

- **Affordable Measurements**

m-RNA concentrations (micro-arrays), protein concentrations
 \leftrightarrow Expression levels $x_A(t)$, $x_B(t)$

- **System Identification**

Interventionist. Disturb system (without breaking it).
Learn structure from changes in measurements

- **Optimal Experimental Design**

For given model: **Short(est)** sequences of experiments leading to identification?

Gene Regulatory Networks

- **Genes can regulate other genes**

Protein from gene A can be transcription factor:
up-/down-regulates transcription of gene B .

Causal link $A \rightarrow B$ in **gene regulatory network**

- **Affordable Measurements**

m-RNA concentrations (micro-arrays), protein concentrations
 \leftrightarrow Expression levels $x_A(t)$, $x_B(t)$

- **System Identification**

Interventionist. Disturb system (without breaking it).
Learn structure from changes in measurements

- **Optimal Experimental Design**

For given model: **Short(est)** sequences of experiments leading to identification?

Gene Regulatory Networks

- **Genes can regulate other genes**

Protein from gene A can be transcription factor:
up-/down-regulates transcription of gene B .

Causal link $A \rightarrow B$ in **gene regulatory network**

- **Affordable Measurements**

m-RNA concentrations (micro-arrays), protein concentrations
 \leftrightarrow Expression levels $x_A(t)$, $x_B(t)$

- **System Identification**

Interventionist. Disturb system (without breaking it).
Learn structure from changes in measurements

- **Optimal Experimental Design**

For given model: **Short(est)** sequences of experiments leading to identification?

Linearized ODE Model

ODE Model

$$d\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t))dt + d\mathbf{W}(t)$$

$$E[\mathbf{x}(t)] \rightarrow \mathbf{x}_0 \quad (t \rightarrow \infty)$$

$\mathbf{x}(t)$ Expression levels n genes

$\mathbf{f}(\cdot)$ Non-linear model

\mathbf{x}_0 Unperturbed steady state

- 1 Linearize around steady state: $\mathbf{x}(t) \rightarrow \mathbf{x}(t) - \mathbf{x}_0$.

System matrix $\mathbf{A} = (df_i/dx_{0,j})_{ij}$

- 2 Disturb system by $\mathbf{u}(t) \equiv \mathbf{u}_*$, measure new steady state:

$$d\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t) - \mathbf{u}(t) + d\mathbf{W}(t), \quad \mathbf{x}_* = \lim_{t \rightarrow \infty} E[\mathbf{x}(t)]$$

- 3 Motivates linear model for measurements:

$$\mathbf{u}_* = \mathbf{A}\mathbf{x}_* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$$

Linearized ODE Model

ODE Model

$$d\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t))dt + d\mathbf{W}(t)$$

$$E[\mathbf{x}(t)] \rightarrow \mathbf{x}_0 \quad (t \rightarrow \infty)$$

$\mathbf{x}(t)$ Expression levels n genes

$\mathbf{f}(\cdot)$ Non-linear model

\mathbf{x}_0 Unperturbed steady state

- 1 Linearize around steady state: $\mathbf{x}(t) \rightarrow \mathbf{x}(t) - \mathbf{x}_0$.

System matrix $\mathbf{A} = (df_i/dx_{0,j})_{ij}$

- 2 Disturb system by $\mathbf{u}(t) \equiv \mathbf{u}_*$, measure new steady state:

$$d\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t) - \mathbf{u}(t) + d\mathbf{W}(t), \quad \mathbf{x}_* = \lim_{t \rightarrow \infty} E[\mathbf{x}(t)]$$

- 3 Motivates **linear model** for measurements:

$$\mathbf{u}_* = \mathbf{A}\mathbf{x}_* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$$

Linearized ODE Model

ODE Model

$$d\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t))dt + d\mathbf{W}(t)$$

$$E[\mathbf{x}(t)] \rightarrow \mathbf{x}_0 \quad (t \rightarrow \infty)$$

$\mathbf{x}(t)$ Expression levels n genes

$\mathbf{f}(\cdot)$ Non-linear model

\mathbf{x}_0 Unperturbed steady state

- 1 Linearize around steady state: $\mathbf{x}(t) \rightarrow \mathbf{x}(t) - \mathbf{x}_0$.

System matrix $\mathbf{A} = (df_i/dx_{0,j})_{ij}$

- 2 Disturb system by $\mathbf{u}(t) \equiv \mathbf{u}_*$, measure new steady state:

$$d\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t) - \mathbf{u}(t) + d\mathbf{W}(t), \quad \mathbf{x}_* = \lim_{t \rightarrow \infty} E[\mathbf{x}(t)]$$

- 3 Motivates **linear model** for measurements:

$$\mathbf{u}_* = \mathbf{A}\mathbf{x}_* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$$

Linearized ODE Model

ODE Model

$$d\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t))dt + d\mathbf{W}(t)$$

$$E[\mathbf{x}(t)] \rightarrow \mathbf{x}_0 \quad (t \rightarrow \infty)$$

$\mathbf{x}(t)$ Expression levels n genes

$\mathbf{f}(\cdot)$ Non-linear model

\mathbf{x}_0 Unperturbed steady state

- 1 Linearize around steady state: $\mathbf{x}(t) \rightarrow \mathbf{x}(t) - \mathbf{x}_0$.

System matrix $\mathbf{A} = (df_i/dx_{0,j})_{ij}$

- 2 Disturb system by $\mathbf{u}(t) \equiv \mathbf{u}_*$, measure new steady state:

$$d\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t) - \mathbf{u}(t) + d\mathbf{W}(t), \quad \mathbf{x}_* = \lim_{t \rightarrow \infty} E[\mathbf{x}(t)]$$

- 3 Motivates **linear model** for measurements:

$$\mathbf{u}_* = \mathbf{A}\mathbf{x}_* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$$

Bayesian Linear Model

- **Likelihood** $P(D|\mathbf{A}) = \prod_k N(\mathbf{u}_k | \mathbf{A}\mathbf{x}_k, \sigma^2 \mathbf{I})$. **Prior** $P(\mathbf{A})$

$$\text{Bayesian Posterior : } P(\mathbf{A}|D) \propto P(D|\mathbf{A})P(\mathbf{A})$$

Why not just (penalized) **maximum likelihood estimation**:

$$\hat{\mathbf{A}} = \operatorname{argmax} P(D|\mathbf{A})P(\mathbf{A}) ?$$

- **Estimation is not sufficient here**
Optimal design fundamentally needs uncertainty quantification
 \Rightarrow Posterior $P(\mathbf{A}|D)$ is just that
- Decisions are needed after **many fewer** than n experiments.
 \Rightarrow “Objective” classical estimation theory breaks down
- Besides: Is \mathbf{A} really completely unknown ... ?

Bayesian Linear Model

- **Likelihood** $P(D|\mathbf{A}) = \prod_k N(\mathbf{u}_k | \mathbf{A}\mathbf{x}_k, \sigma^2 \mathbf{I})$. **Prior** $P(\mathbf{A})$

$$\text{Bayesian Posterior : } P(\mathbf{A}|D) \propto P(D|\mathbf{A})P(\mathbf{A})$$

Why not just (penalized) **maximum likelihood estimation**:

$$\hat{\mathbf{A}} = \operatorname{argmax} P(D|\mathbf{A})P(\mathbf{A}) ?$$

- **Estimation is not sufficient here**

Optimal design fundamentally needs uncertainty quantification

⇒ Posterior $P(\mathbf{A}|D)$ is just that

- Decisions are needed after **many fewer** than n experiments.
⇒ “Objective” classical estimation theory breaks down
- Besides: Is \mathbf{A} really completely unknown ... ?

Bayesian Linear Model

- **Likelihood** $P(D|\mathbf{A}) = \prod_k N(\mathbf{u}_k | \mathbf{A}\mathbf{x}_k, \sigma^2 \mathbf{I})$. **Prior** $P(\mathbf{A})$

$$\text{Bayesian Posterior : } P(\mathbf{A}|D) \propto P(D|\mathbf{A})P(\mathbf{A})$$

Why not just (penalized) **maximum likelihood estimation**:

$$\hat{\mathbf{A}} = \operatorname{argmax} P(D|\mathbf{A})P(\mathbf{A}) ?$$

- **Estimation is not sufficient here**
Optimal design fundamentally needs uncertainty quantification
 \Rightarrow Posterior $P(\mathbf{A}|D)$ is just that
- Decisions are needed after **many fewer** than n experiments.
 \Rightarrow “Objective” classical estimation theory breaks down
- Besides: Is \mathbf{A} really completely unknown ... ?

Bayesian Linear Model

- **Likelihood** $P(D|\mathbf{A}) = \prod_k N(\mathbf{u}_k | \mathbf{A}\mathbf{x}_k, \sigma^2 \mathbf{I})$. **Prior** $P(\mathbf{A})$

$$\text{Bayesian Posterior : } P(\mathbf{A}|D) \propto P(D|\mathbf{A})P(\mathbf{A})$$

Why not just (penalized) **maximum likelihood estimation**:

$$\hat{\mathbf{A}} = \operatorname{argmax} P(D|\mathbf{A})P(\mathbf{A}) ?$$

- **Estimation is not sufficient here**
Optimal design fundamentally needs uncertainty quantification
 \Rightarrow Posterior $P(\mathbf{A}|D)$ is just that
- Decisions are needed after **many fewer** than n experiments.
 \Rightarrow “Objective” classical estimation theory breaks down
- Besides: Is \mathbf{A} really completely unknown ... ?

A Sparsity Prior Distribution

- All biological regulatory networks are **sparsely connected**
⇒ **A** should have many very small entries
- Encoding sparsity of **A** is a must!
⇒ **Sparsity-enforcing prior distribution** $P(\mathbf{A})$

Laplace Prior

$$P(\mathbf{A}) = \prod_{ij} P(a_{ij}), \quad P(a_{ij}) = \frac{\tau}{2} e^{-\tau |a_{ij}|}$$

A Sparsity Prior Distribution

- All biological regulatory networks are **sparsely connected**
⇒ **A** should have many very small entries
- Encoding sparsity of **A** is a must!
⇒ **Sparsity-enforcing prior distribution** $P(\mathbf{A})$




Laplace Prior

$$P(\mathbf{A}) = \prod_{ij} P(a_{ij}), \quad P(a_{ij}) = \frac{\tau}{2} e^{-\tau |a_{ij}|}$$

A Sparsity Prior Distribution

Laplace Prior

$$P(\mathbf{A}) = \prod_{ij} P(a_{ij}), \quad P(a_{ij}) = \frac{\tau}{2} e^{-\tau |a_{ij}|}$$



Approximate Inference: Rough Idea

- Bayesian posterior for one row \mathbf{a} of \mathbf{A}

$$P(\mathbf{a}|D) \propto P(D|\mathbf{a}) \prod_i P(a_i)$$

Hard “just” because $P(a_i)$ are **not** Gaussian

- **Moment matching** idea: $P(D|\mathbf{a})P(a_i)$ **not** Gaussian either. Gaussian with **same moments** have form $P(D|\mathbf{a})\tilde{P}(a_i|b_i, \pi_i)$.

$$P(\mathbf{a}|D) \approx Q(\mathbf{a}) \propto P(D|\mathbf{a}) \prod_i \tilde{P}(a_i|b_i, \pi_i)$$

- **Expectation Propagation**: iterates moment matching over i : Update variational parameters b_i, π_i s.t.:

$$Q_{old}(\mathbf{a})P(a_i)/\tilde{P}(a_i) \longleftrightarrow Q_{new}(\mathbf{a}) \text{ [same moments]}$$

Approximate Inference: Rough Idea

- Bayesian posterior for one row \mathbf{a} of \mathbf{A}

$$P(\mathbf{a}|D) \propto P(D|\mathbf{a}) \prod_i P(a_i)$$

Hard “just” because $P(a_i)$ are **not** Gaussian

- **Moment matching** idea: $P(D|\mathbf{a})P(a_i)$ **not** Gaussian either. Gaussian with **same moments** have form $P(D|\mathbf{a})\tilde{P}(a_i|b_i, \pi_i)$.

$$P(\mathbf{a}|D) \approx Q(\mathbf{a}) \propto P(D|\mathbf{a}) \prod_i \tilde{P}(a_i|b_i, \pi_i)$$

- **Expectation Propagation**: iterates moment matching over i :
Update variational parameters b_i, π_i s.t.:

$$Q_{old}(\mathbf{a})P(a_i)/\tilde{P}(a_i) \longleftrightarrow Q_{new}(\mathbf{a}) \text{ [same moments]}$$

Approximate Inference: Rough Idea

- Bayesian posterior for one row \mathbf{a} of \mathbf{A}

$$P(\mathbf{a}|D) \propto P(D|\mathbf{a}) \prod_i P(a_i)$$

Hard “just” because $P(a_i)$ are **not** Gaussian

- Moment matching** idea: $P(D|\mathbf{a})P(a_i)$ **not** Gaussian either. Gaussian with **same moments** have form $P(D|\mathbf{a})\tilde{P}(a_i|b_i, \pi_i)$.

$$P(\mathbf{a}|D) \approx Q(\mathbf{a}) \propto P(D|\mathbf{a}) \prod_i \tilde{P}(a_i|b_i, \pi_i)$$

- Expectation Propagation**: iterates moment matching over i :
Update variational parameters b_i, π_i s.t.:

$$Q_{old}(\mathbf{a})P(a_i)/\tilde{P}(a_i) \longleftrightarrow Q_{new}(\mathbf{a}) \text{ [same moments]}$$

Bayesian Experimental Design

Information Gain Score

$$S(\mathbf{u}_*, \mathbf{x}_* | D) = D[Q'(\mathbf{A} | D \cup \{(\mathbf{u}_*, \mathbf{x}_*)\}) \parallel Q(\mathbf{A} | D)]$$

$D[Q' \parallel Q]$: Information gained in $Q \rightarrow Q'$.

Efficient exact computation for Gaussians Q, Q'

- But outcome \mathbf{x}_* **unknown** before experiment \mathbf{u}_* done!?
 \Rightarrow Use **expected score** under current knowledge $Q(\mathbf{x}_* | D, \mathbf{u}_*)$.
 Exact sampling: $\mathbf{A} \sim Q(\cdot | D)$, $\mathbf{x}_* = \mathbf{A}^{-1} \mathbf{u}_*$
- Score many candidates \mathbf{u}_* very efficiently:
 Pick maximizer of $E_{Q(\mathbf{x}_* | D, \mathbf{u}_*)}[S(\mathbf{u}_*, \mathbf{x}_* | D)]$

Bayesian Experimental Design

Information Gain Score

$$S(\mathbf{u}_*, \mathbf{x}_* | D) = D[Q'(\mathbf{A} | D \cup \{(\mathbf{u}_*, \mathbf{x}_*)\}) \parallel Q(\mathbf{A} | D)]$$

$D[Q' \parallel Q]$: Information gained in $Q \rightarrow Q'$.

Efficient exact computation for Gaussians Q, Q'

- But outcome \mathbf{x}_* **unknown** before experiment \mathbf{u}_* done!
 \Rightarrow Use **expected score** under current knowledge $Q(\mathbf{x}_* | D, \mathbf{u}_*)$.
 Exact sampling: $\mathbf{A} \sim Q(\cdot | D)$, $\mathbf{x}_* = \mathbf{A}^{-1} \mathbf{u}_*$
- Score many candidates \mathbf{u}_* very efficiently:
 Pick maximizer of $E_{Q(\mathbf{x}_* | D, \mathbf{u}_*)}[S(\mathbf{u}_*, \mathbf{x}_* | D)]$

Bayesian Experimental Design

Information Gain Score

$$S(\mathbf{u}_*, \mathbf{x}_* | D) = D[Q'(\mathbf{A} | D \cup \{(\mathbf{u}_*, \mathbf{x}_*)\}) \parallel Q(\mathbf{A} | D)]$$

$D[Q' \parallel Q]$: Information gained in $Q \rightarrow Q'$.

Efficient exact computation for Gaussians Q, Q'

- But outcome \mathbf{x}_* **unknown** before experiment \mathbf{u}_* done!
 \Rightarrow Use **expected score** under current knowledge $Q(\mathbf{x}_* | D, \mathbf{u}_*)$.
 Exact sampling: $\mathbf{A} \sim Q(\cdot | D)$, $\mathbf{x}_* = \mathbf{A}^{-1} \mathbf{u}_*$
- Score many candidates \mathbf{u}_* very efficiently:
 Pick maximizer of $E_{Q(\mathbf{x}_* | D, \mathbf{u}_*)}[S(\mathbf{u}_*, \mathbf{x}_* | D)]$

Our Approach As Black Box

- **Robust, efficient** code will be released:
Predictable running time. **Easy to use** for non-experts
- Free parameters σ^2, τ :
Bayesian **automatic selection**, given related task data
- Applies to **time series data** just as well (if linear model does)
- Encompasses **generalized linear models**:
 - Non-Gaussian noise (outliers)
 - Discrete or point process observations

Our Approach As Black Box

- **Robust, efficient** code will be released:
Predictable running time. **Easy to use** for non-experts
- Free parameters σ^2, τ :
Bayesian **automatic selection**, given related task data
- Applies to **time series data** just as well (if linear model does)
- Encompasses **generalized linear models**:
 - Non-Gaussian noise (outliers)
 - Discrete or point process observations

Our Approach As Black Box

- **Robust, efficient** code will be released:
Predictable running time. **Easy to use** for non-experts
- Free parameters σ^2, τ :
Bayesian **automatic selection**, given related task data
- Applies to **time series data** just as well (if linear model does)
- Encompasses **generalized linear models**:
 - Non-Gaussian noise (outliers)
 - Discrete or point process observations


Our Approach As Black Box

- **Robust, efficient** code will be released:
Predictable running time. **Easy to use** for non-experts
- Free parameters σ^2, τ :
Bayesian **automatic selection**, given related task data
- Applies to **time series data** just as well (if linear model does)
- Encompasses **generalized linear models**:
 - Non-Gaussian noise (outliers)
 - Discrete or point process observations

Experimental Setup

Common practice: validate on data from realistic simulation.

- Sample small-world network, $n = 50$ genes



- Model with Hill-type kinetics, parameters randomly drawn (similar to Kholodenko *et.al.*, 02)


$$f_i(\mathbf{x}) = -V_{di} \frac{x_i}{d_i + x_i} + V_{si} \prod_{j \in \mathcal{A}_i} \frac{1 + A_{ij} \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}}{1 + \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}} \prod_{j \in \mathcal{I}_i} \frac{1}{1 + \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}}$$

- Pool of 200 μ_s (unit norm; 3 non-zeros, sparsity for biological relevance) randomly drawn
- Noise variance σ^2 estimated from simpler random networks. Prior precision τ set by heuristic

Experimental Setup

Common practice: validate on data from realistic simulation.

- Sample small-world network, $n = 50$ genes



- Model with Hill-type kinetics, parameters randomly drawn (similar to Kholodenko *et.al.*, 02)


$$f_i(\mathbf{x}) = -V_{di} \frac{x_i}{d_i + x_i} + V_{si} \prod_{j \in \mathcal{A}_i} \frac{1 + A_{ij} \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}}{1 + \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}} \prod_{j \in \mathcal{I}_i} \frac{1}{1 + \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}}$$

- Pool of 200 \mathbf{u}_* (unit norm; 3 non-zeros, sparsity for biological relevance) randomly drawn
- Noise variance σ^2 estimated from simpler random networks. Prior precision τ set by heuristic

Experimental Setup

Common practice: validate on data from realistic simulation.

- Sample small-world network, $n = 50$ genes



- Model with Hill-type kinetics, parameters randomly drawn (similar to Kholodenko *et.al.*, 02)

$$f_i(\mathbf{x}) = -V_{di} \frac{x_i}{d_i + x_i} + V_{si} \prod_{j \in \mathcal{A}_i} \frac{1 + A_{ij} \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}}{1 + \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}} \prod_{j \in \mathcal{I}_i} \frac{1}{1 + \left(\frac{x_j}{\kappa_{ij}} \right)^{n_{ij}}}$$

- Pool of 200 \mathbf{u}_* (unit norm; 3 non-zeros, sparsity for biological relevance) randomly drawn
- Noise variance σ^2 estimated from simpler random networks. Prior precision τ set by heuristic

Decision and Evaluation

- Network from joint posterior $Q(\mathbf{A})$?
Rank edges $i \leftarrow j$ by $Q(\{|a_{ij}| > 0.1\})$
- ROC curve: false positive rate \rightarrow true positive rate.
iAUC: area under ROC curve, up to # FPs = # edges.
Random ranking has iAUC = 0.02
- About 25% edges have value ≈ 0 in true \mathbf{A} (at steady state), not detectable by linearized model. Excluded from iAUC computation


Decision and Evaluation

- Network from joint posterior $Q(\mathbf{A})$?
Rank edges $i \leftarrow j$ by $Q(\{|a_{ij}| > 0.1\})$
- ROC curve: false positive rate \rightarrow true positive rate.
iAUC: area under ROC curve, up to # FPs = # edges.
Random ranking has iAUC = 0.02
- About 25% edges have value ≈ 0 in true \mathbf{A} (at steady state), not detectable by linearized model. Excluded from iAUC computation

Decision and Evaluation


- Network from joint posterior $Q(\mathbf{A})$?
Rank edges $i \leftarrow j$ by $Q(\{|a_{ij}| > 0.1\})$
- ROC curve: false positive rate \rightarrow true positive rate.
iAUC: area under ROC curve, up to # FPs = # edges.
Random ranking has iAUC = 0.02
- About 25% edges have value ≈ 0 in true \mathbf{A} (at steady state), not detectable by linearized model. Excluded from iAUC computation

Results




Comparison Tegnér *et.al.*

- Tegnér *et.al.* (PNAS 03): most cited work on experimental design for network identification.
- We do not use quantizations: our method works better and is 2 orders of magnitude faster
- They require node in-degree ≤ 3 (unrealistic in scale-free networks), we do not [comparison done on such graphs]




Comparison Tegnér *et.al.*

- Tegnér *et.al.* (PNAS 03): most cited work on experimental design for network identification.
- We do not use quantizations: our method works better and is 2 orders of magnitude faster
- They require node in-degree ≤ 3 (unrealistic in scale-free networks), we do not [comparison done on such graphs]



Comparison Tegnér *et.al.*

- Tegnér *et.al.* (PNAS 03): most cited work on experimental design for network identification.
- We do not use quantizations: our method works better and is 2 orders of magnitude faster
- They require node in-degree ≤ 3 (unrealistic in scale-free networks), we do not [comparison done on such graphs]



Related Work

- Much work on disturbed linearized ODE models.
Estimation, no inference, no experimental design (except Tegnér *et.al.*)
- Sparse Bayesian Learning (Tipping, 01; Rogers, Girolami, 05)
No experimental design. Uses non-log-concave Student- t prior.
EP more general than SBL
- Markov Chain Monte Carlo (Park, Casella, 05)
Much slower than our method (too slow for large-scale experimental design). Hard to assess convergence even for experts

Related Work

- Much work on disturbed linearized ODE models.
Estimation, no inference, no experimental design (except Tegnér *et.al.*)
- Sparse Bayesian Learning (Tipping, 01; Rogers, Girolami, 05)
No experimental design. Uses non-log-concave Student- t prior.
EP more general than SBL
- Markov Chain Monte Carlo (Park, Casella, 05)
Much slower than our method (too slow for large-scale experimental design). Hard to assess convergence even for experts

Related Work

- Much work on disturbed linearized ODE models.
Estimation, no inference, no experimental design (except Tegnér *et.al.*)
- Sparse Bayesian Learning (Tipping, 01; Rogers, Girolami, 05)
No experimental design. Uses non-log-concave Student- t prior.
EP more general than SBL
- Markov Chain Monte Carlo (Park, Casella, 05)
Much slower than our method (too slow for large-scale experimental design). Hard to assess convergence even for experts

Conclusions

- **Fast accurate** approximate inference, experimental design in disturbed linearized ODE setup
- Network sparsity is **key prior** assumption. Experimental design can lead to **large savings**
- Can be used with **time-course** measurements just as well
- Robust, easy-to-use method. **Code** with Matlab interface will be released
- Linearized ODE approach is limited:
 - Small, controlled u_* to stay in linearity region (experimental techniques?), but large u_* for better SNR
 - No saturation, Michaelis-Menten, *etc*
- Bayesian inference and experimental design for (simple) non-linear ODEs of biochemical kinetics?

Conclusions

- **Fast accurate** approximate inference, experimental design in disturbed linearized ODE setup
- Network sparsity is **key prior** assumption. Experimental design can lead to **large savings**
- Can be used with **time-course** measurements just as well
- Robust, easy-to-use method. **Code** with Matlab interface will be released
- Linearized ODE approach is limited:
 - Small, controlled u , to stay in linearity region (experimental techniques?), but large u , for better SNR
 - No saturation, Michaelis-Menten, *etc*
- Bayesian inference and experimental design for (simple) non-linear ODEs of biochemical kinetics?

Conclusions

- **Fast accurate** approximate inference, experimental design in disturbed linearized ODE setup
- Network sparsity is **key prior** assumption. Experimental design can lead to **large savings**
- Can be used with **time-course** measurements just as well
- Robust, easy-to-use method. **Code** with Matlab interface will be released
- Linearized ODE approach is limited:
 - Small, controlled u , to stay in linearity region (experimental techniques?), but large u , for better SNR
 - No saturation, Michaelis-Menten, *etc*
- Bayesian inference and experimental design for (simple) non-linear ODEs of biochemical kinetics?

Conclusions

- **Fast accurate** approximate inference, experimental design in disturbed linearized ODE setup
- Network sparsity is **key prior** assumption. Experimental design can lead to **large savings**
- Can be used with **time-course** measurements just as well
- Robust, easy-to-use method. **Code** with Matlab interface will be released
- Linearized ODE approach is limited:
 - Small, controlled u , to stay in linearity region (experimental techniques?), but large u , for better SNR
 - No saturation, Michaelis-Menten, *etc*
- Bayesian inference and experimental design for (simple) non-linear ODEs of biochemical kinetics?

Conclusions

- **Fast accurate** approximate inference, experimental design in disturbed linearized ODE setup
- Network sparsity is **key prior** assumption. Experimental design can lead to **large savings**
- Can be used with **time-course** measurements just as well
- Robust, easy-to-use method. **Code** with Matlab interface will be released
- Linearized ODE approach is limited:
 - Small, controlled \mathbf{u}_* to stay in linearity region (experimental techniques?), but large \mathbf{u}_* for better SNR
 - No saturation, Michaelis-Menten, *etc*
- Bayesian inference and experimental design for (simple) non-linear ODEs of biochemical kinetics?

Conclusions

- **Fast accurate** approximate inference, experimental design in disturbed linearized ODE setup
- Network sparsity is **key prior** assumption. Experimental design can lead to **large savings**
- Can be used with **time-course** measurements just as well
- Robust, easy-to-use method. **Code** with Matlab interface will be released
- Linearized ODE approach is limited:
 - Small, controlled \mathbf{u}_* to stay in linearity region (experimental techniques?), but large \mathbf{u}_* for better SNR
 - No saturation, Michaelis-Menten, *etc*
- Bayesian inference and experimental design for (simple) non-linear ODEs of biochemical kinetics?

Conclusions (II)

- Other applications of sparse (generalized) linear models, in systems biology and beyond (natural image statistics, neural spike coding, adaptive control, *etc*)
- Applications to dynamical or nonparametric models?
- Submitted for journal publication
- Details:
M. Seeger, F. Steinke, K. Tsuda
Bayesian Inference and Optimal Design in the Sparse Linear Model, AI and Statistics 2007
www.kyb.tuebingen.mpg.de/bs/people/seeger
- Useful for your work? **Do not hesitate to get in touch**