An Introduction to Systems Biology from a Machine Learning Perspective

Neil D. Lawrence

Max Planck Institute Retreat, Ringberg Castle

5th May 2008

Outline

Outline

Talk Overview

- Feynman had experience of biology.
  "There's Plenty of Room at the Bottom: An Invitation to Enter a New Field of Physics" Feynman (1959).
- This talk
  "There's Still a lot of Room Here: An Invitation to Enter Another Sub-field of Machine Learning" Lawrence (2008).
- Systems biology: Interaction of Biological Components
- Differentiates from a *Reductionist* approach to biology that previously dominated.

Talk Overview

- Feynman had experience of biology.
  "There's Plenty of Room at the Bottom: An Invitation to Enter a New Field of Physics" Feynman (1959).
- This talk
  "There's Still a lot of Room Here: An Invitation to Enter Another Sub-field of Machine Learning" Lawrence (2008).
- Systems biology: Interaction of Biological Components
- Differentiates from a *Reductionist* approach to biology that previously dominated.

Talk Overview

- Feynman had experience of biology.
  "There's Plenty of Room at the Bottom: An Invitation to Enter a New Field of Physics" Feynman (1959).
- This talk
  "There's Still a lot of Room Here: An Invitation to Enter Another Sub-field of Machine Learning" Lawrence (2008).
- Systems biology: Interaction of Biological Components
- Differentiates from a *Reductionist* approach to biology that previously dominated.

Talk Overview

- Feynman had experience of biology.
  "There's Plenty of Room at the Bottom: An Invitation to Enter a New Field of Physics" Feynman (1959).
- This talk
  "There's Still a lot of Room Here: An Invitation to Enter Another Sub-field of Machine Learning" Lawrence (2008).
- Systems biology: Interaction of Biological Components
- Differentiates from a *Reductionist* approach to biology that previously dominated.

Reductionist Thinking

- Reductionism Garfinkel (1991) breaks down a system into its component parts.
  - The Max Planck society *is just* a collection of reseach institutes.
  - A Max Planck institute *is just* a collection of clever people.
  - Conclusion: to understand the Max Planck Society we must just understand clever people.
  - A human body *is just* a collection of biological cells.
  - A biological cell *is just* a collection of biochemical interactions.
  - Conclusion: To understand a human being we just need to understand what each gene does.

Reductionist Thinking

- Reductionism Garfinkel (1991) breaks down a system into its component parts.
  - The Max Planck society *is just* a collection of reseach institutes.
  - A Max Planck institute *is just* a collection of clever people.
  - Conclusion: to understand the Max Planck Society we must just understand clever people.
  - A human body *is just* a collection of biological cells.
  - A biological cell *is just* a collection of biochemical interactions.
  - Conclusion: To understand a human being we just need to understand what each gene does.

Reductionist Thinking

- Reductionism Garfinkel (1991) breaks down a system into its component parts.
  - The Max Planck society *is just* a collection of reseach institutes.
  - A Max Planck institute *is just* a collection of clever people.
  - Conclusion: to understand the Max Planck Society we must just understand clever people.
  - A human body *is just* a collection of biological cells.
  - A biological cell *is just* a collection of biochemical interactions.
  - Conclusion: To understand a human being we just need to understand what each gene does.

Reductionist Thinking

- Reductionism Garfinkel (1991) breaks down a system into its component parts.
  - The Max Planck society *is just* a collection of reseach institutes.
  - A Max Planck institute *is just* a collection of clever people.
  - Conclusion: to understand the Max Planck Society we must just understand clever people.
  - A human body *is just* a collection of biological cells.
  - A biological cell *is just* a collection of biochemical interactions.
  - Conclusion: To understand a human being we just need to understand what each gene does.

Reductionist Thinking

- Reductionism Garfinkel (1991) breaks down a system into its component parts.
  - The Max Planck society *is just* a collection of reseach institutes.
  - A Max Planck institute *is just* a collection of clever people.
  - Conclusion: to understand the Max Planck Society we must just understand clever people.
  - A human body *is just* a collection of biological cells.
  - A biological cell *is just* a collection of biochemical interactions.
  - Conclusion: To understand a human being we just need to understand what each gene does.

Reductionist Thinking

- Reductionism Garfinkel (1991) breaks down a system into its component parts.
  - The Max Planck society *is just* a collection of reseach institutes.
  - A Max Planck institute *is just* a collection of clever people.
  - Conclusion: to understand the Max Planck Society we must just understand clever people.
  - A human body *is just* a collection of biological cells.
  - A biological cell *is just* a collection of biochemical interactions.
  - Conclusion: To understand a human being we just need to understand what each gene does.

Reductionist Thinking

- Reductionism Garfinkel (1991) breaks down a system into its component parts.
  - The Max Planck society *is just* a collection of reseach institutes.
  - A Max Planck institute *is just* a collection of clever people.
  - Conclusion: to understand the Max Planck Society we must just understand clever people.
  - A human body *is just* a collection of biological cells.
  - A biological cell *is just* a collection of biochemical interactions.
  - Conclusion: To understand a human being we just need to understand what each gene does.

Systems Approach

- There is little point in reducing a system as far as its component parts (quarks and leptons ... strings (M-theory) ...) if the questions are at a higher level.
- Disease mechanisms may affect one gene, but finding a target for a cure involves the whole pathway.
- Study the system at the level in which we want to ask questions:
  - *e.g.* **Which proteins interact in the ERK/MAPK signalling pathway?**
    This is a critical pathway in cell proliferation. Of strong interest in cancer.

Systems Approach

- There is little point in reducing a system as far as its component parts (quarks and leptons ... strings (M-theory) ...) if the questions are at a higher level.
- Disease mechanisms may affect one gene, but finding a target for a cure involves the whole pathway.
- Study the system at the level in which we want to ask questions:
  - *e.g.* **Which proteins interact in the ERK/MAPK signalling pathway?**
    This is a critical pathway in cell proliferation. Of strong interest in cancer.

Systems Approach

- There is little point in reducing a system as far as its component parts (quarks and leptons ... strings (M-theory) ...) if the questions are at a higher level.
- Disease mechanisms may affect one gene, but finding a target for a cure involves the whole pathway.
- Study the system at the level in which we want to ask questions:
  - *e.g.* **Which proteins interact in the ERK/MAPK signalling pathway?**
    This is a critical pathway in cell proliferation. Of strong interest in cancer.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - Not enough information to describe the models. Need to turn to mechanistic models to help.
    - Lessons from large data set will still apply as system may be very complex.
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - Not enough information to describe the models. Need to turn to mechanistic models to help.
    - Lessons from large data set will still apply as system may be very complex.
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - Not enough information to describe the models, need to turn to mechanistic models to help.
    - Lessons from large data set will still apply as system may be very complex.
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - Not enough information to describe the models. Need to turn to mechanistic models to help.
    - Lessons from large data set will still apply as system may be very complex.
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - Not enough information to describe the model. Need to turn to mechanistic models to help.
    - Lessons from large data set will still apply as system may be very complex.
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - Not enough information to describe the model. Need to turn to mechanistic models to help.
    - Lessons from large data set will still apply as system may be very complex!
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - Not enough information to describe the model. Need to turn to mechanistic models to help.
    - Lessons from large data set will still apply as system may be very complex!
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - ⋆ Not enough information to describe the model. Need to turn to mechanistic models to help.
    - ⋆ Lessons from large data set will still apply as system may be very complex!
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - ⋆ Not enough information to describe the model. Need to turn to mechanistic models to help.
    - ⋆ Lessons from large data set will still apply as system may be very complex!
- A big focus for our research in Manchester: *inference in mechanistic models*.

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - ★ Not enough information to describe the model. Need to turn to mechanistic models to help.
    - ★ Lessons from large data set will still apply as system may be very complex!
- A big focus for our research in Manchester: *inference in mechanistic models.*

Machine Learning and Systems Biology

- Where does machine learning come in?
  - Models of interaction are not fully characterised. Use inference and learning to deal with unknowns.
- Is this what we normally do?
  - No — models are mechanistic in inspiration not black box.
  - However, perhaps it's what we *will* do in the future!
- My prediction: Machine learning in the future will have two major foci.
  - V. large data sets. e.g. prediction of relevant adverts.
  - Small data set relative to complexity of the system.
    - ★ Not enough information to describe the model. Need to turn to mechanistic models to help.
    - ★ Lessons from large data set will still apply as system may be very complex!
- A big focus for our research in Manchester: *inference in mechanistic models*.

Systems Biology

- Systems biology provides an oppurtunity to analyse complex systems by combining prior knowledge with data.
- Integrate data with knowledge of the chemical kinetics of the system.
- This talk:
  - Review of transcription.
  - Chemical kinetics in a simple synthetic biology system.
  - Inference of hidden variables in single input motifs.
  - Model selection through Bayes factors (given time!).

Systems Biology

- Systems biology provides an oppurtunity to analyse complex systems by combining prior knowledge with data.
- Integrate data with knowledge of the chemical kinetics of the system.
- This talk:
  - Review of transcription.
  - Chemical kinetics in a simple synthetic biology system.
  - Inference of hidden variables in single input motifs.
  - Model selection through Bayes factors (given time!).

Systems Biology

- Systems biology provides an oppurtunity to analyse complex systems by combining prior knowledge with data.
- Integrate data with knowledge of the chemical kinetics of the system.
- This talk:
  - Review of transcription.
  - Chemical kinetics in a simple synthetic biology system.
  - Inference of hidden variables in single input motifs.
  - Model selection through Bayes factors (given time!).

Systems Biology

- Systems biology provides an oppurtunity to analyse complex systems by combining prior knowledge with data.
- Integrate data with knowledge of the chemical kinetics of the system.
- This talk:
  - Review of transcription.
  - Chemical kinetics in a simple synthetic biology system.
  - Inference of hidden variables in single input motifs.
  - Model selection through Bayes factors (given time!).

Systems Biology

- Systems biology provides an oppurtunity to analyse complex systems by combining prior knowledge with data.
- Integrate data with knowledge of the chemical kinetics of the system.
- This talk:
  - Review of transcription.
  - Chemical kinetics in a simple synthetic biology system.
  - Inference of hidden variables in single input motifs.
  - Model selection through Bayes factors (given time!).

Systems Biology

- Systems biology provides an oppurtunity to analyse complex systems by combining prior knowledge with data.
- Integrate data with knowledge of the chemical kinetics of the system.
- This talk:
  - Review of transcription.
  - Chemical kinetics in a simple synthetic biology system.
  - Inference of hidden variables in single input motifs.
  - Model selection through Bayes factors (given time!).

Systems Biology

- Systems biology provides an oppurtunity to analyse complex systems by combining prior knowledge with data.
- Integrate data with knowledge of the chemical kinetics of the system.
- This talk:
  - Review of transcription.
  - Chemical kinetics in a simple synthetic biology system.
  - Inference of hidden variables in single input motifs.
  - Model selection through Bayes factors (given time!).

Transcriptional regulation of gene expression

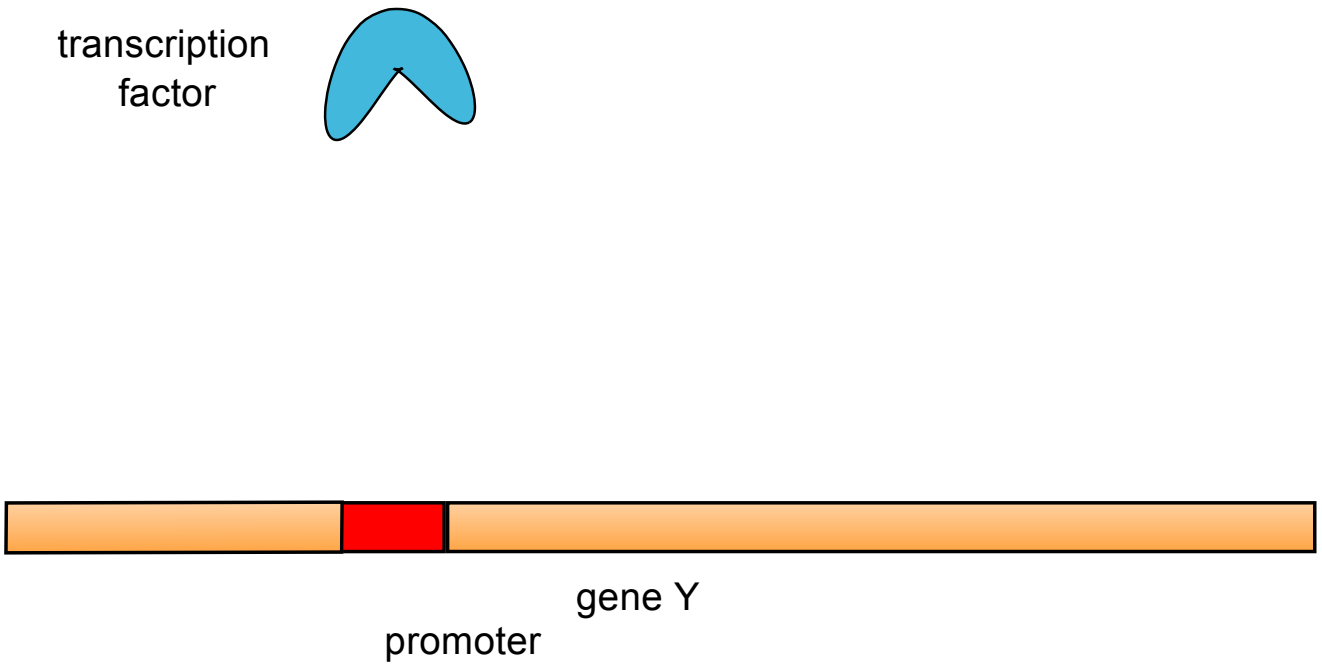

Transcriptional regulation of gene expression

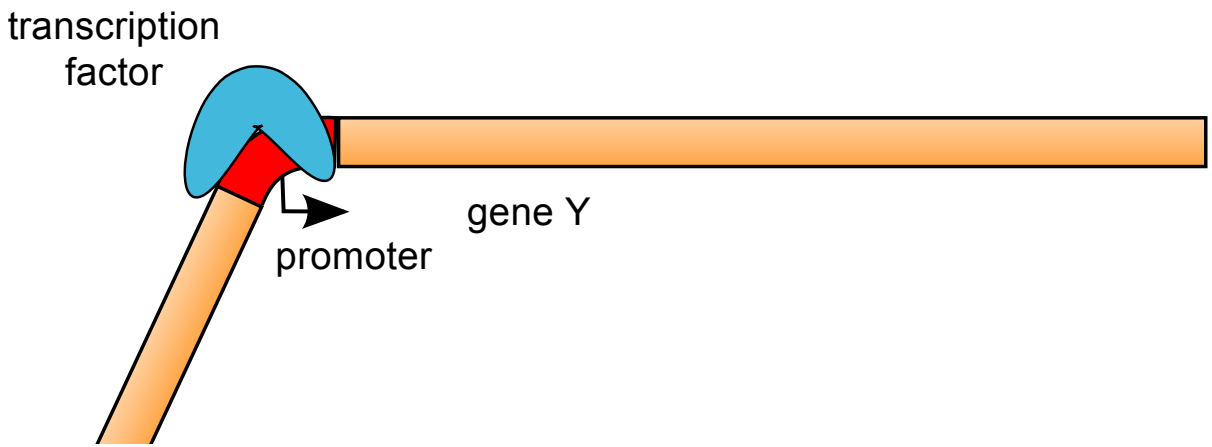

Transcriptional regulation of gene expression

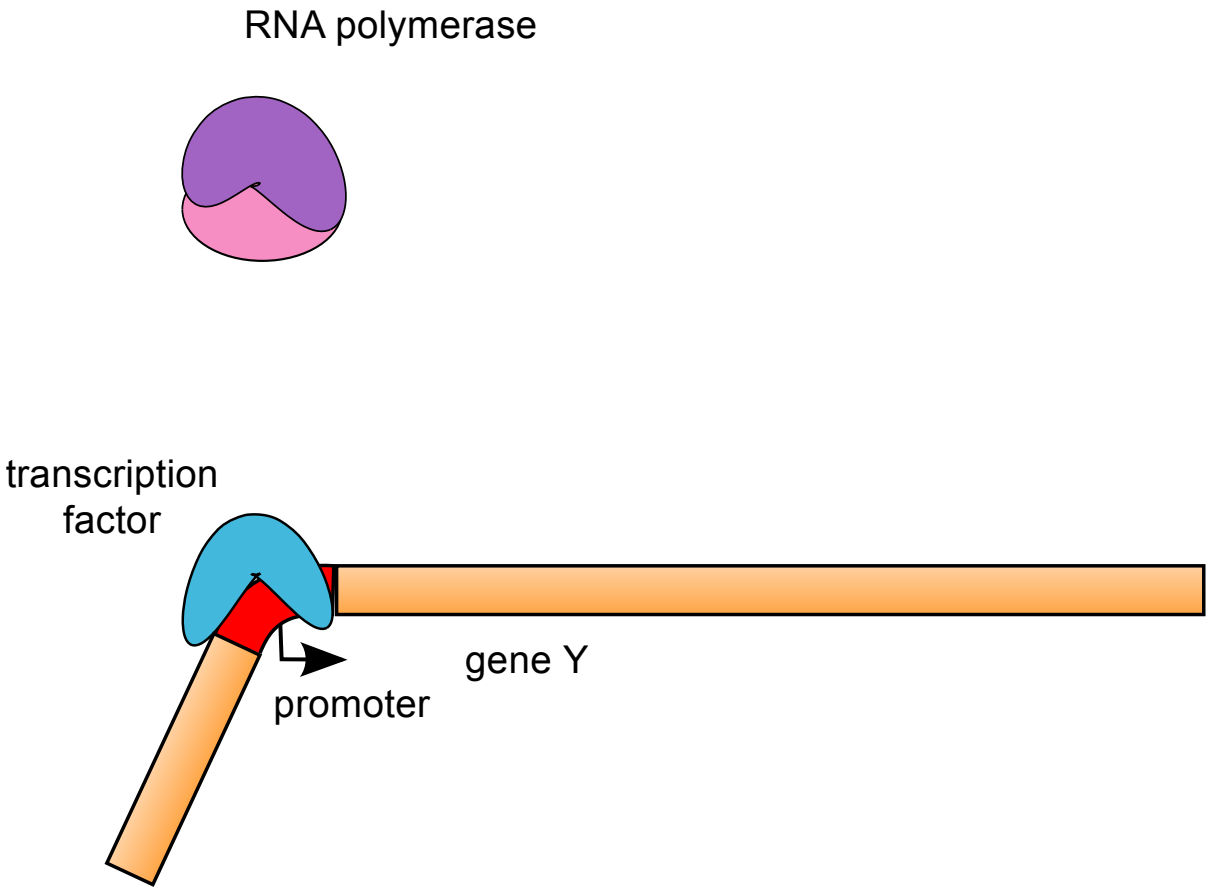

Transcriptional regulation of gene expression

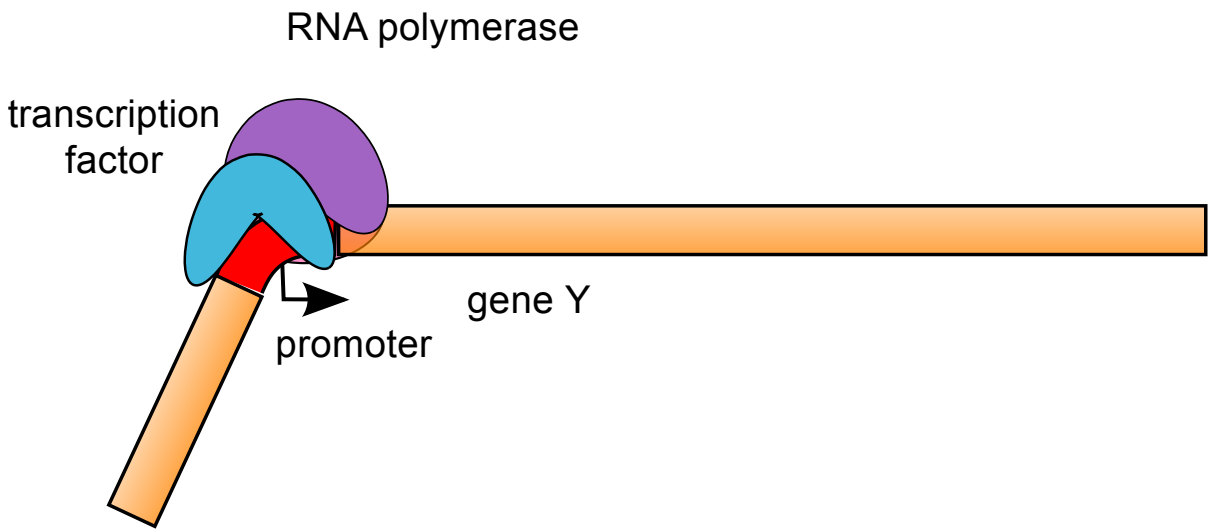

Transcriptional regulation of gene expression

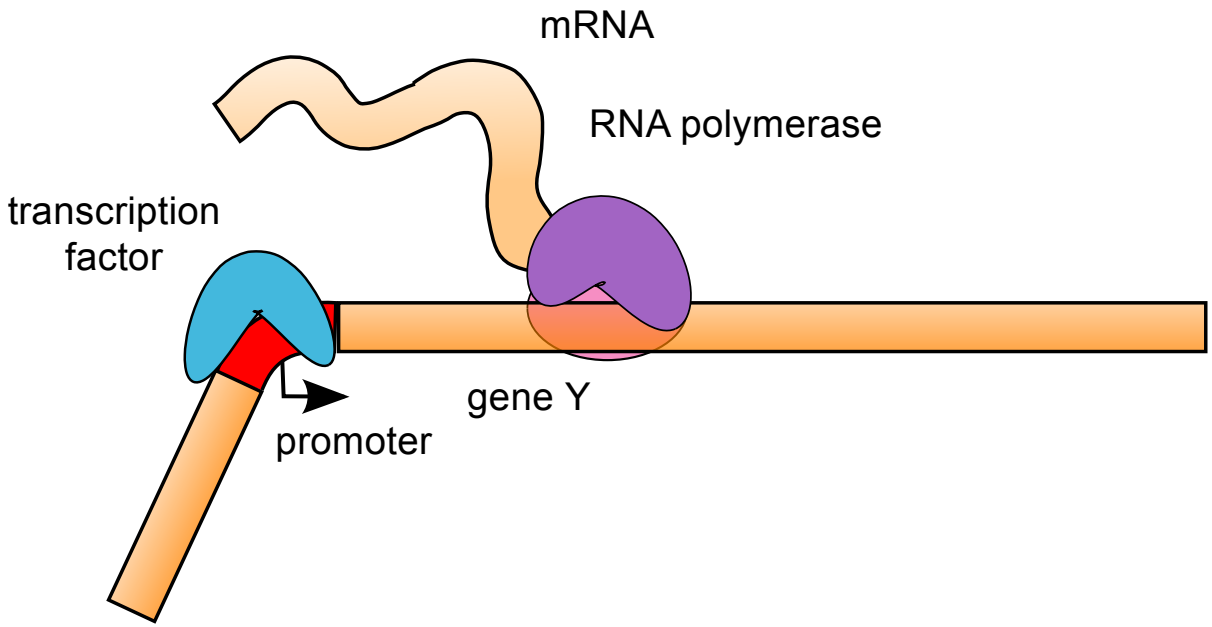

Transcriptional regulation of gene expression

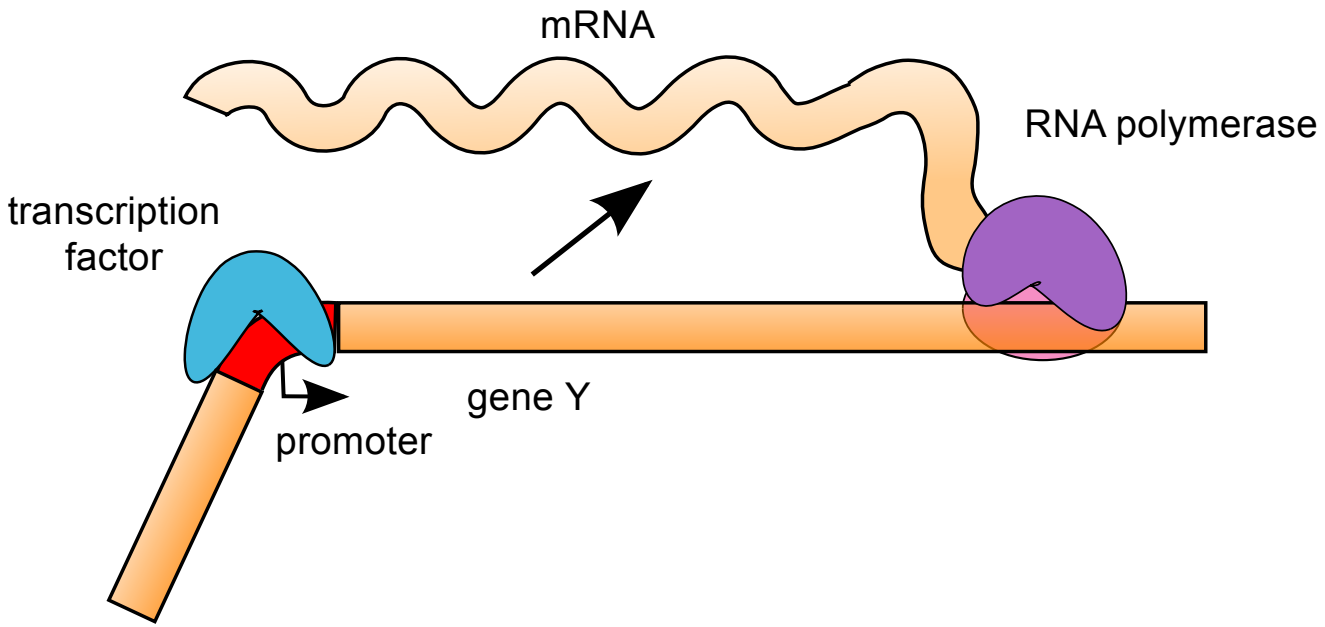

Transcriptional regulation of gene expression

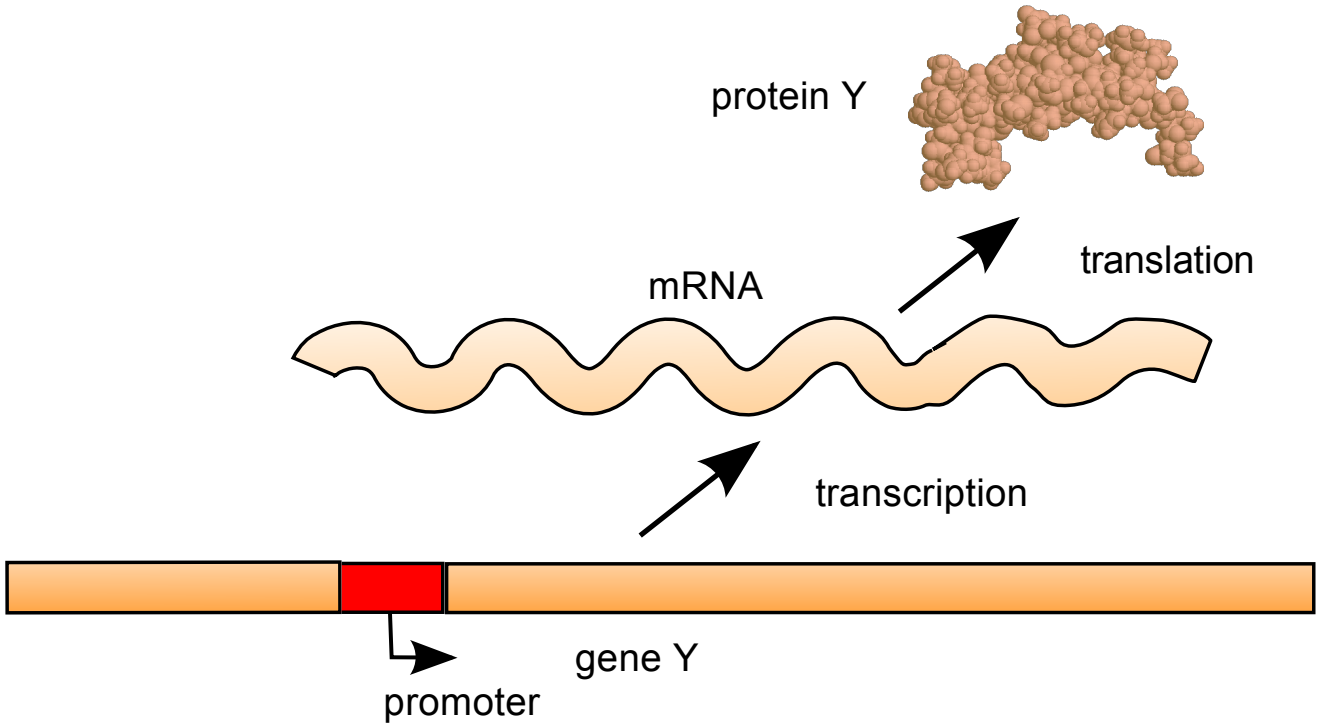

RNA Polymerase

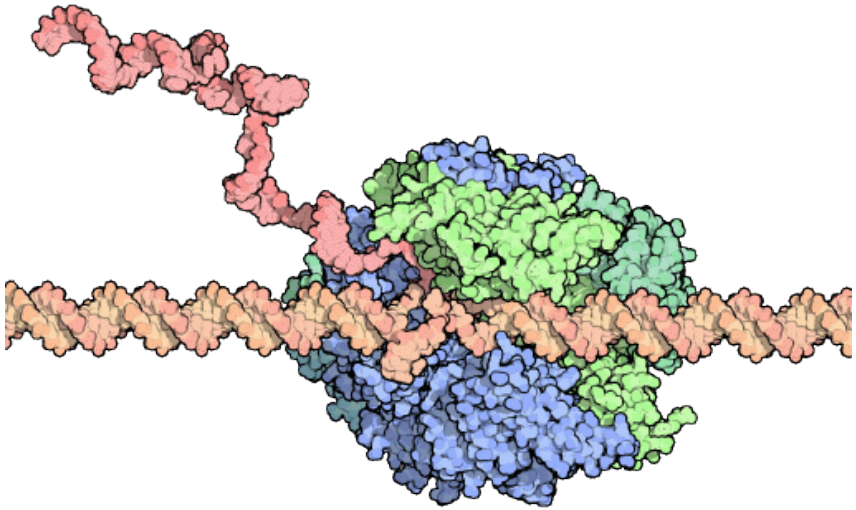

Figure: RNA Polymerase transcribing RNA from DNA. Image from "Molecule of the Month" at the protein data bank: `http://mgl.scripps.edu/people/goodsell/pdb/pdb98/pdb98_1.html`

Repression

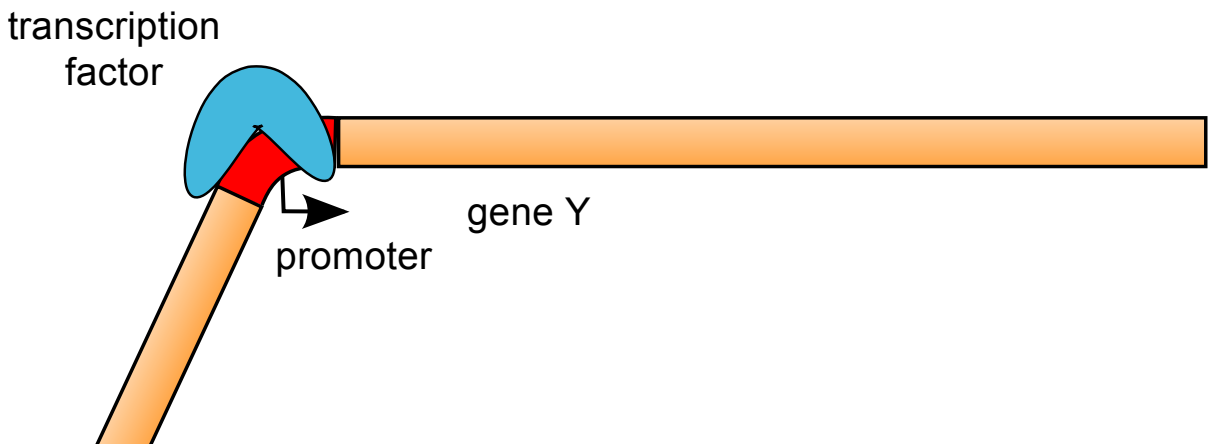

Repression

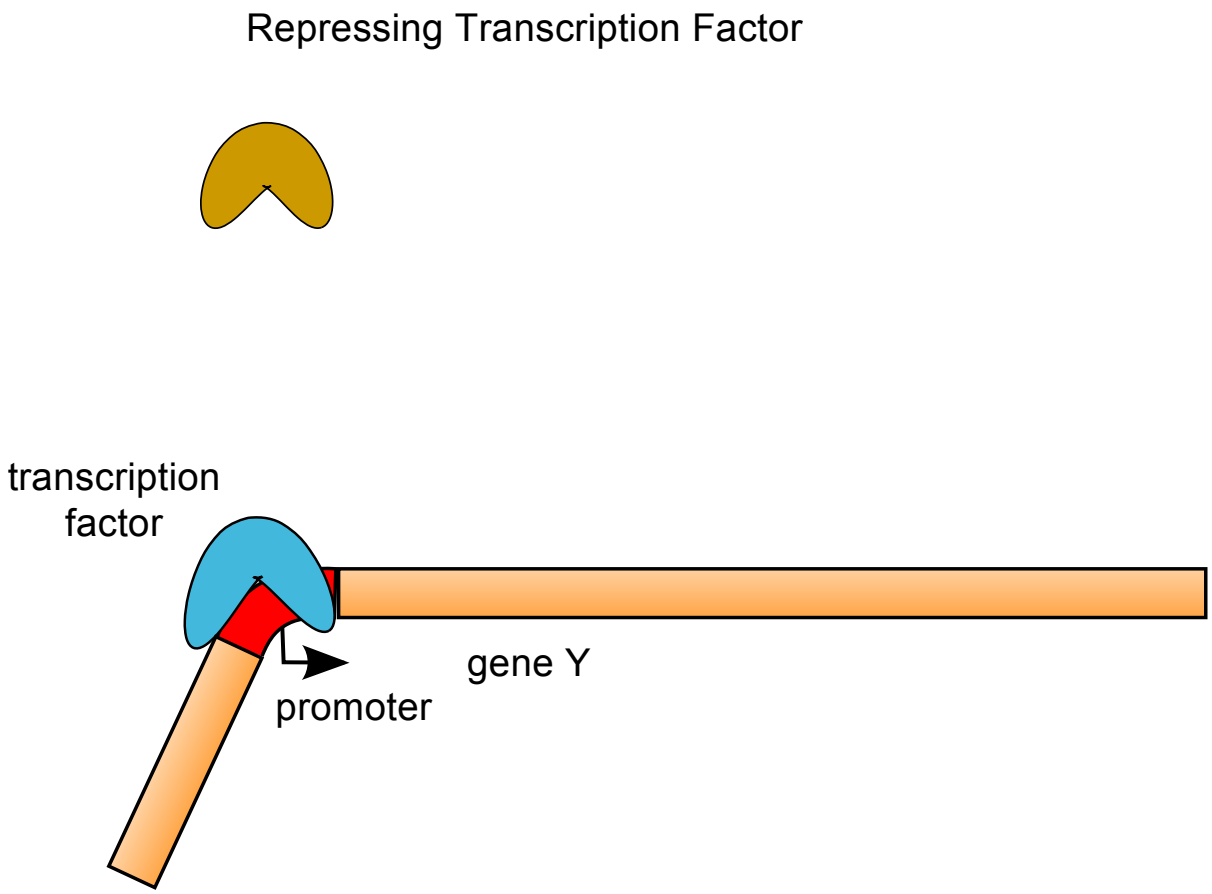
Repressing Transcription Factor
transcription
factor
gene Y
promoter

Repression

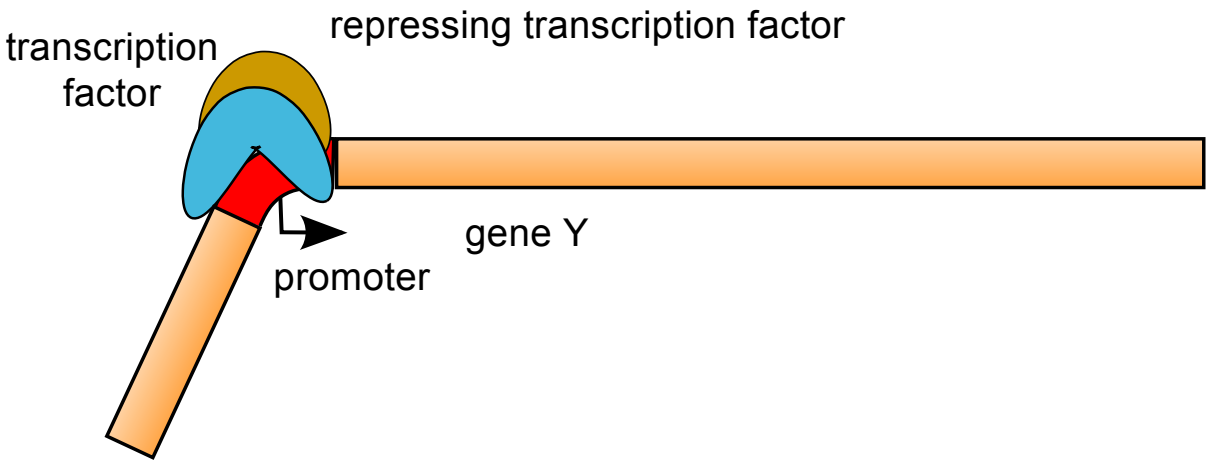
transcription
factor
repressing transcription factor
gene Y
promoter

Repression

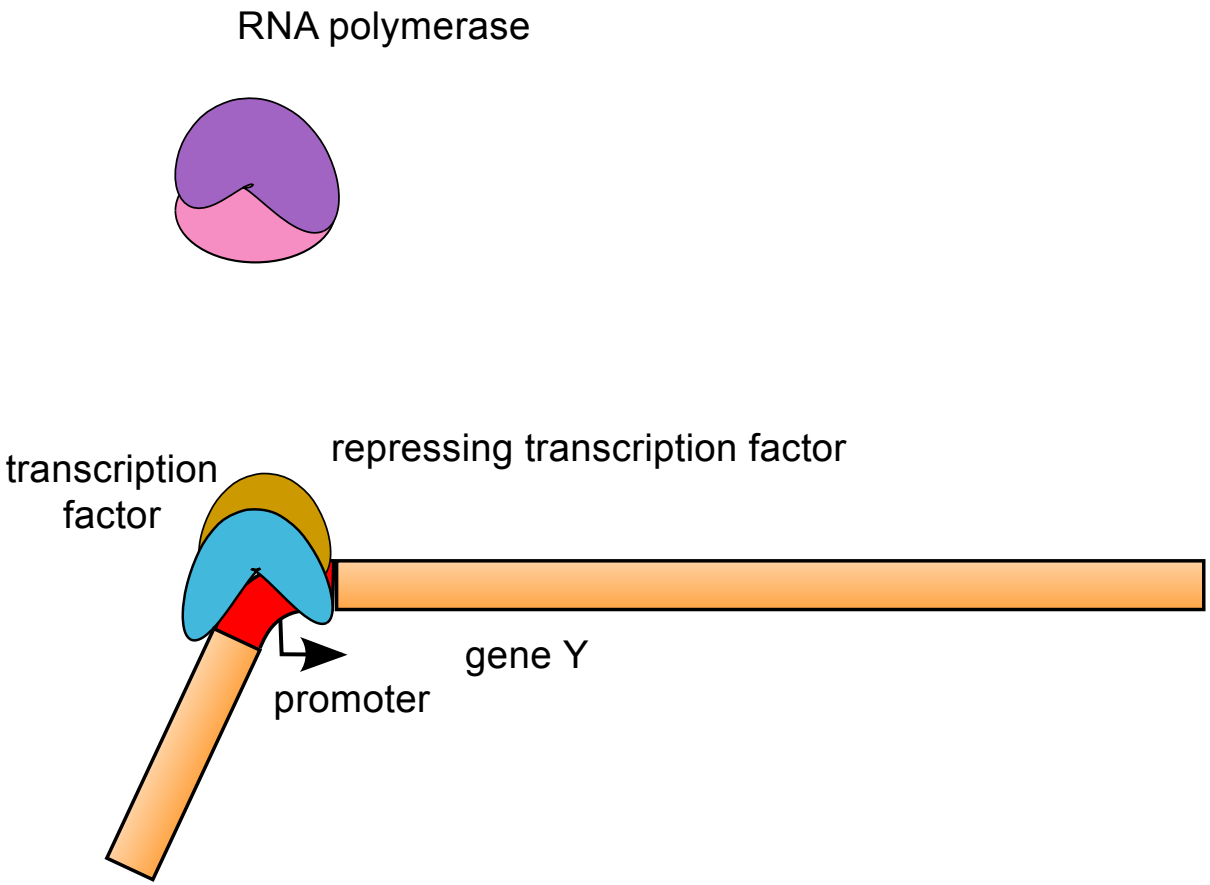
RNA polymerase
transcription factor
repressing transcription factor
gene Y
promoter

Repression

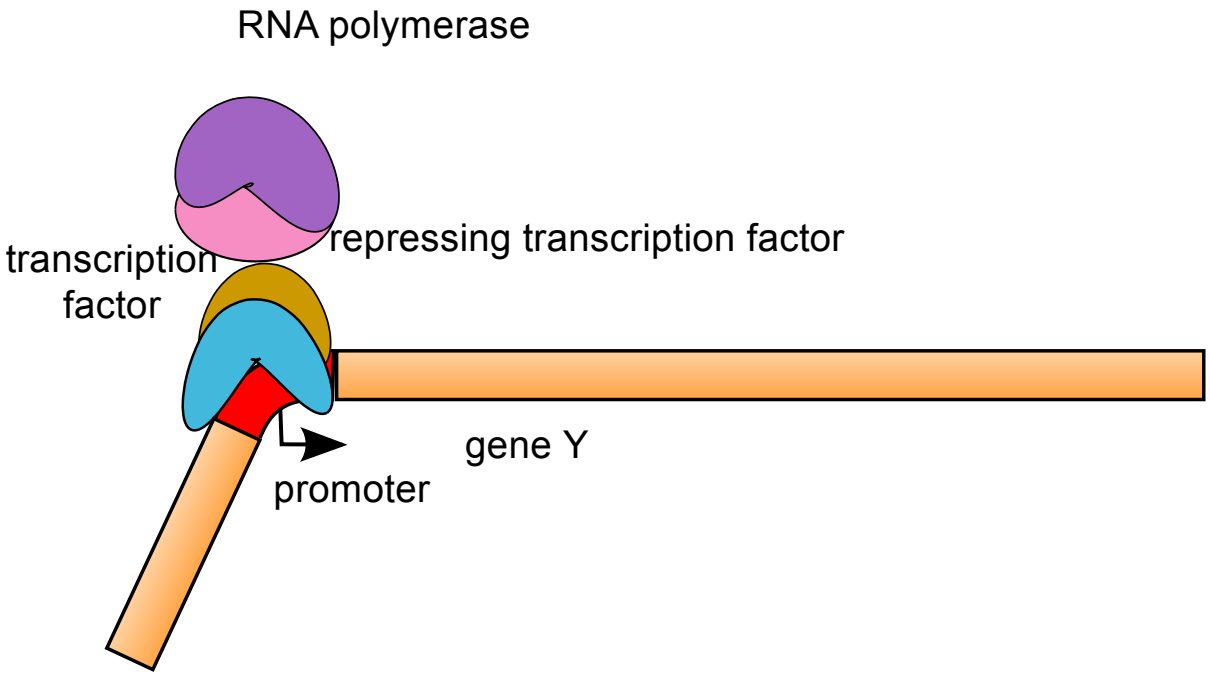
RNA polymerase
transcription factor
repressing transcription factor
gene Y
promoter

Repression

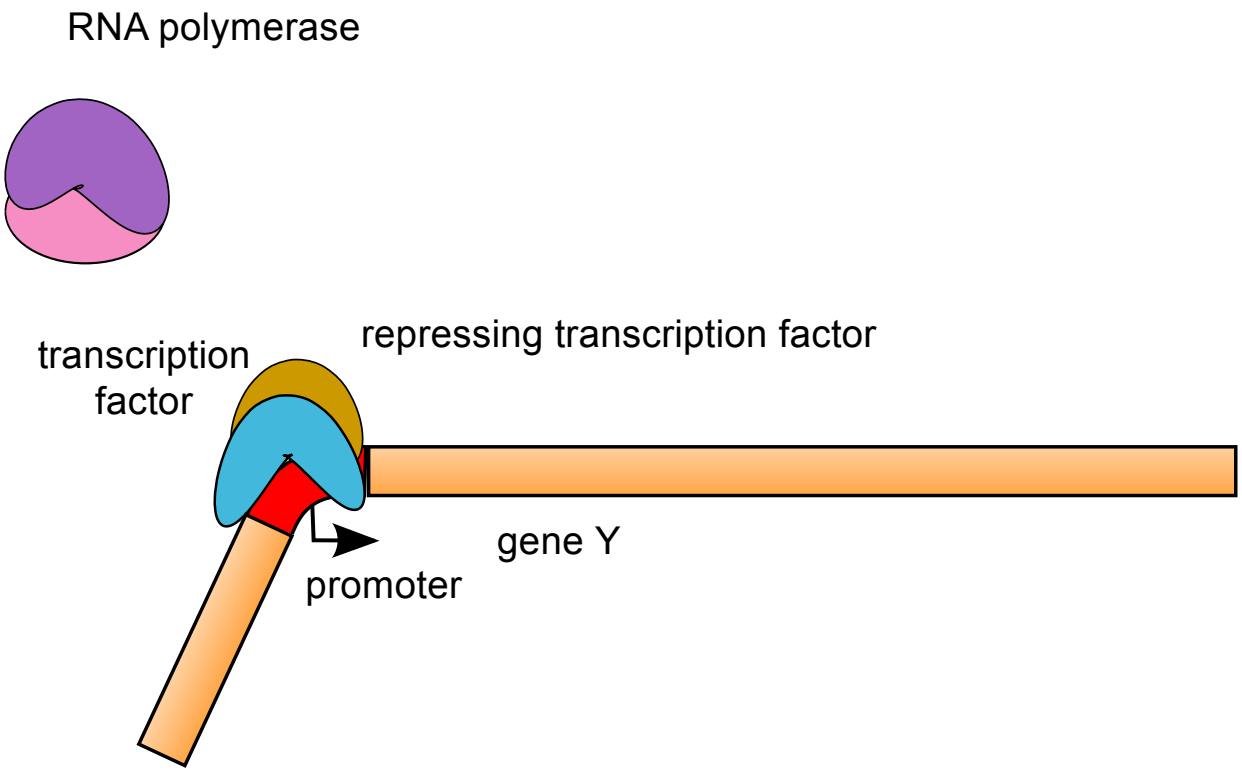
RNA polymerase
transcription factor
repressing transcription factor
gene Y
promoter

The Repressilator

- Real biology involves interaction of several systems.
- The repressilator is the first synthetic biology oscilator.
- Implemented in *E. coli* bacteria.
- How do we model such a system?

The Repressilator

- Real biology involves interaction of several systems.
- The repressilator is the first synthetic biology oscilator.
- Implemented in *E. coli* bacteria.
- How do we model such a system?

The Repressilator

- Real biology involves interaction of several systems.
- The repressilator is the first synthetic biology oscilator.
- Implemented in *E. coli* bacteria.
- How do we model such a system?

The Repressilator

- Real biology involves interaction of several systems.
- The repressilator is the first synthetic biology oscilator.
- Implemented in *E. coli* bacteria.
- How do we model such a system?

The Repressilator

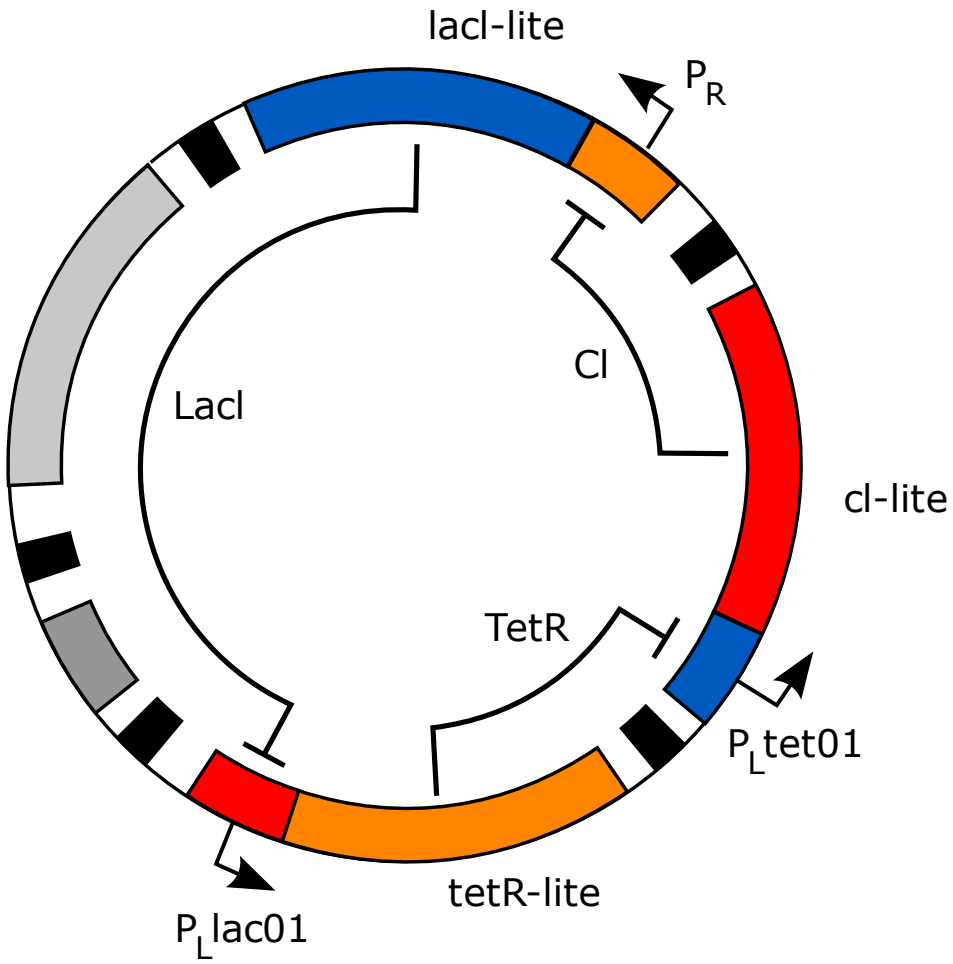

Figure: Repressilator Plasmid. (Elowitz and Leibler, 2000)

Bacteria Plasmids

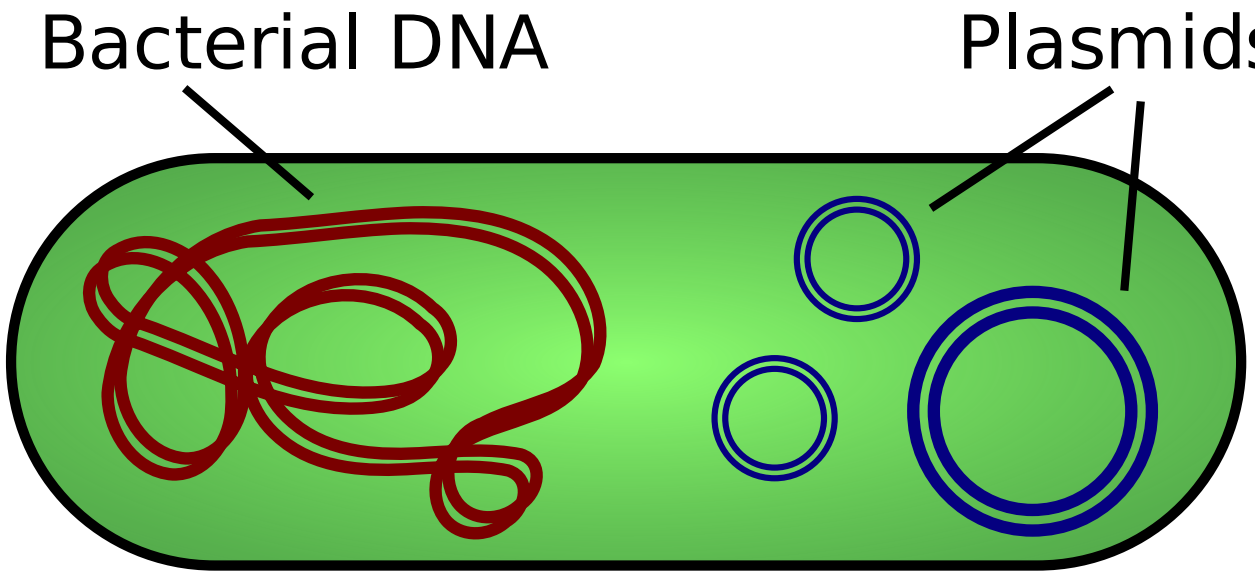

Figure: Schematic of a bacterium with plasmids (Image from wikimedia commons).

Repressilator

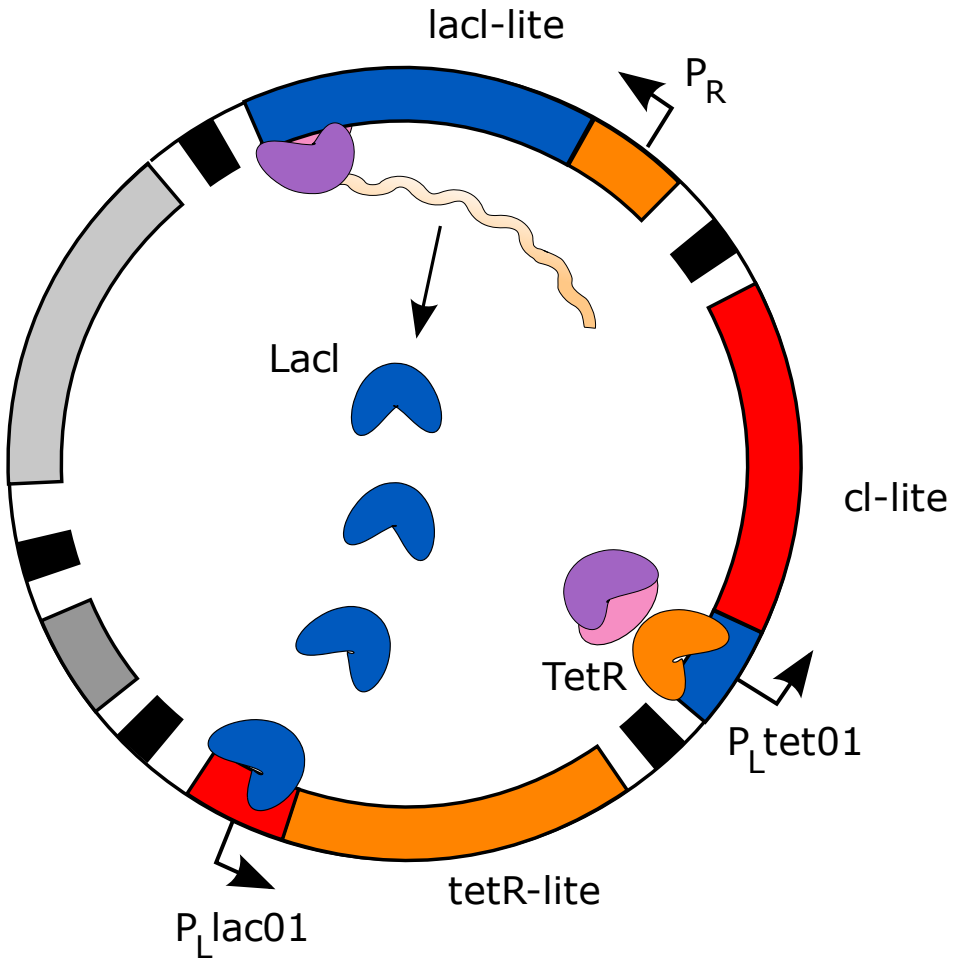

Figure: Repressilator schematic

Repressilator

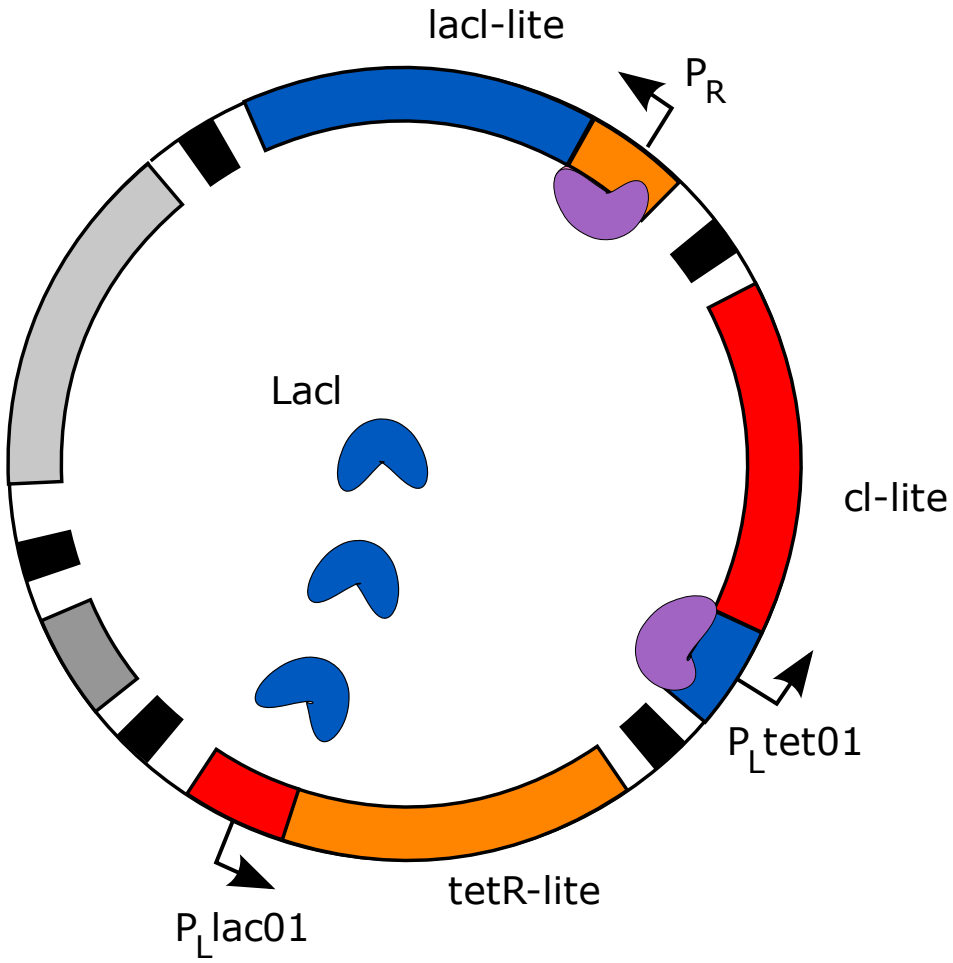

Figure: Repressilator schematic

Repressilator

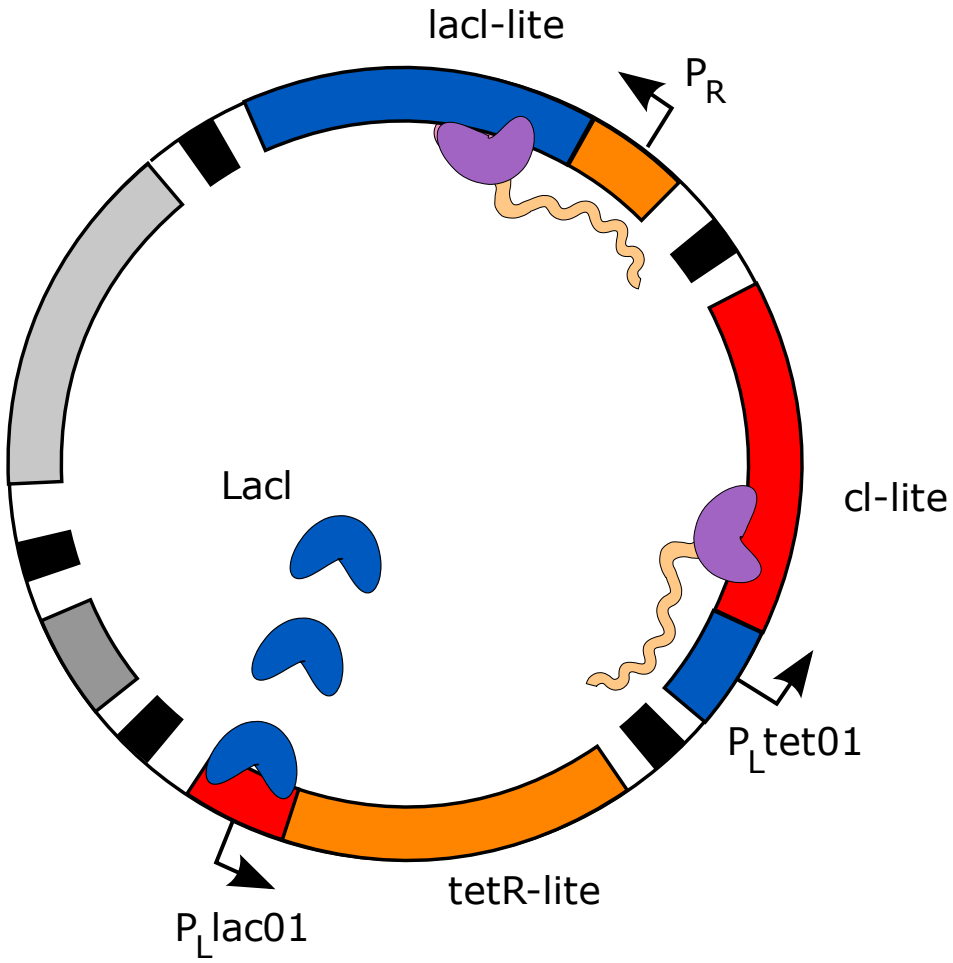

Figure: Repressilator schematic

Repressilator

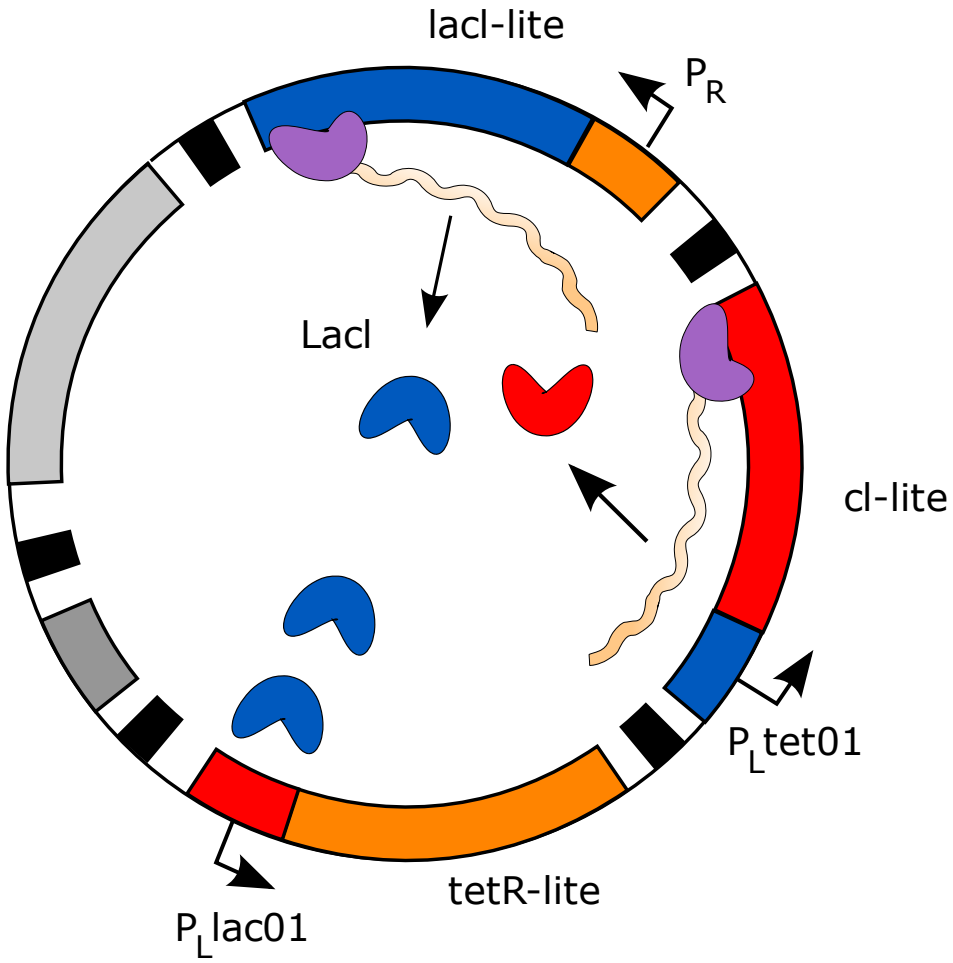

Figure: Repressilator schematic

Repressilator

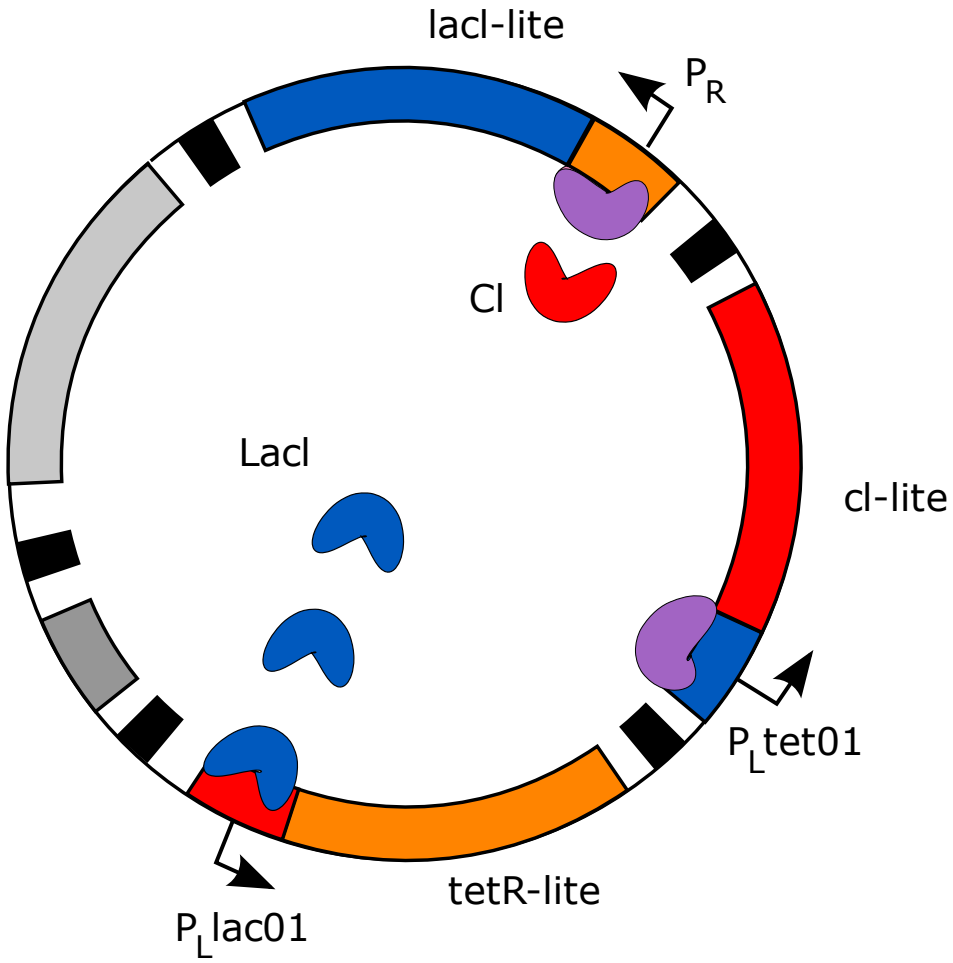

Figure: Repressilator schematic

Repressilator

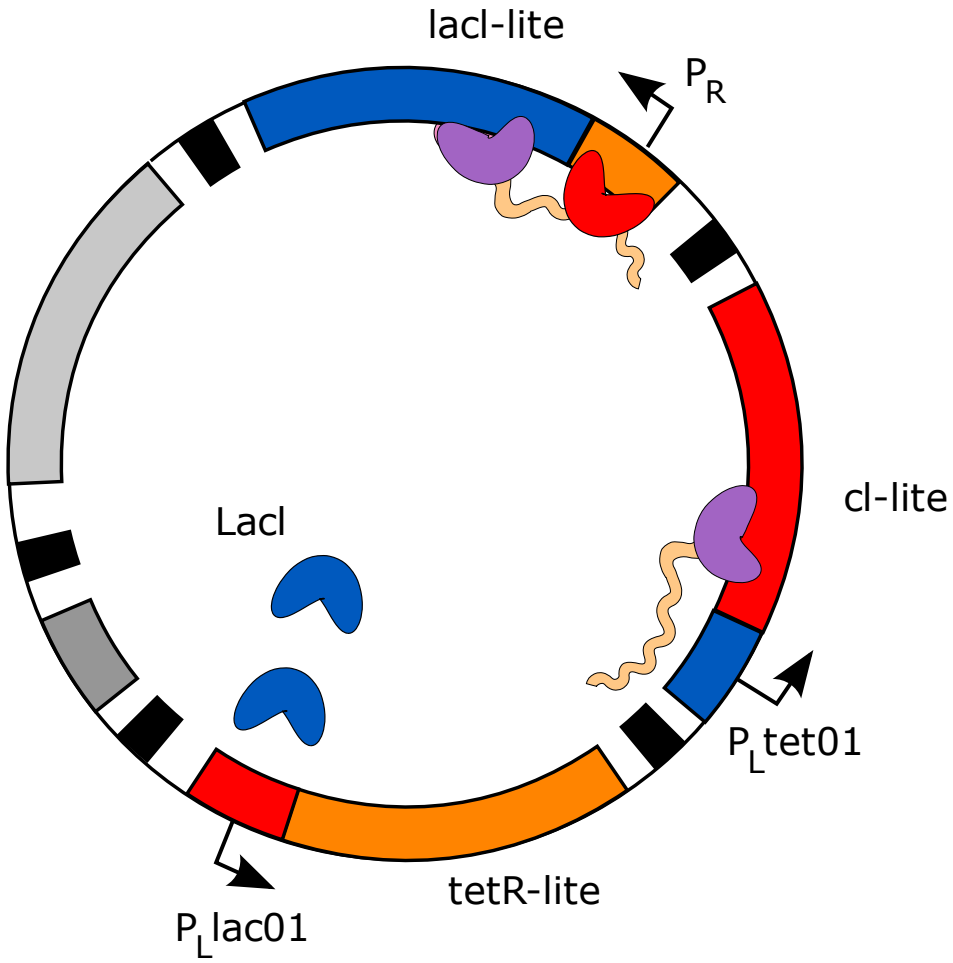

Figure: Repressilator schematic

Repressilator

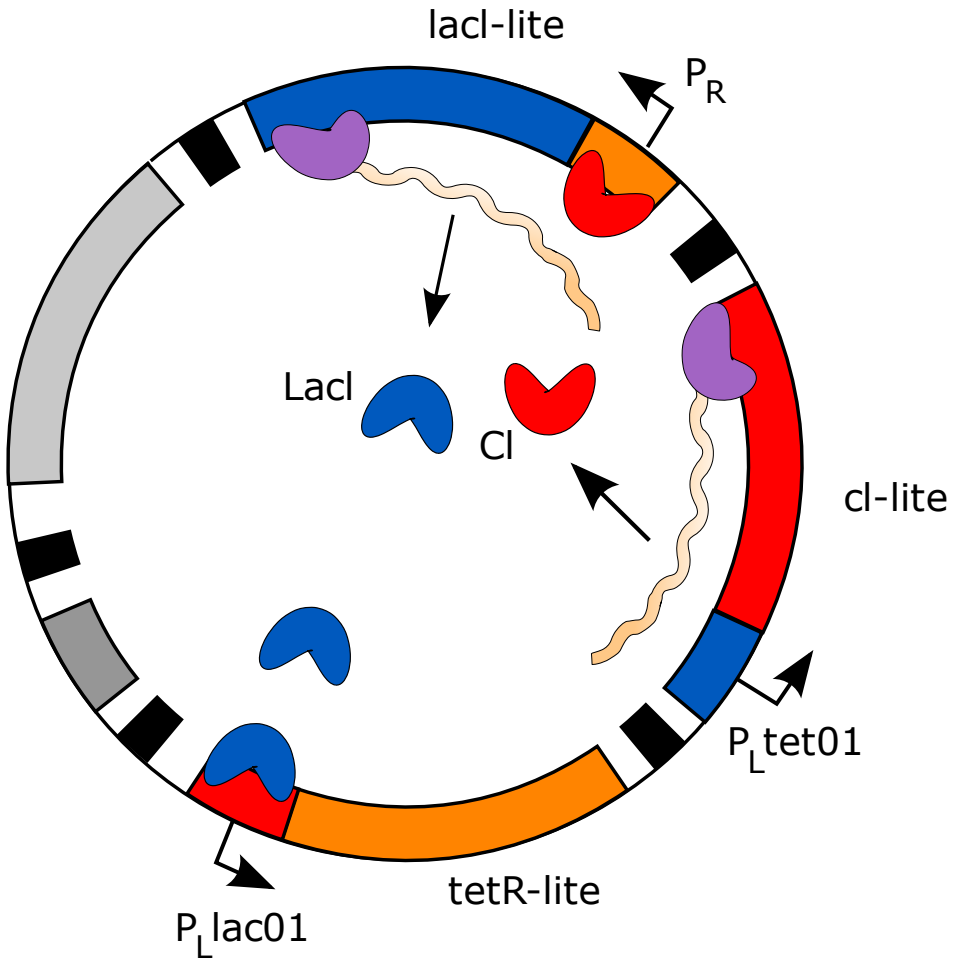

Figure: Repressilator schematic

Repressilator

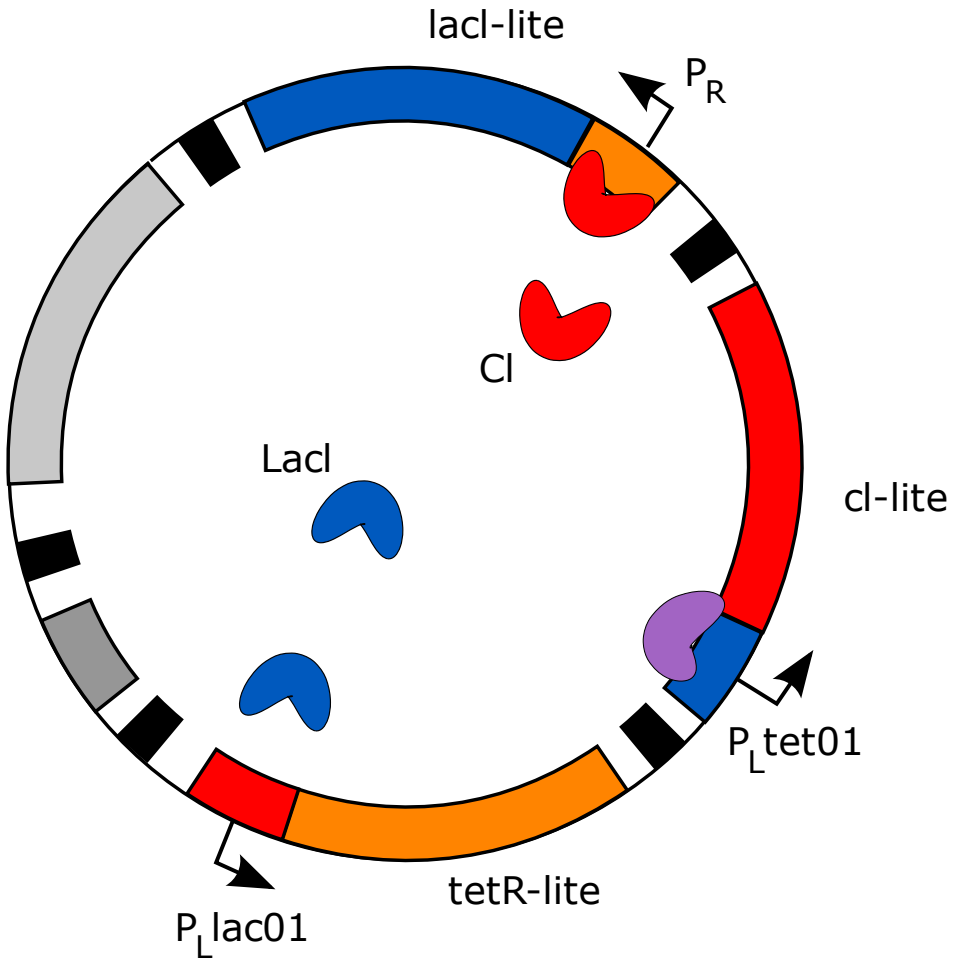

Figure: Repressilator schematic

Repressilator

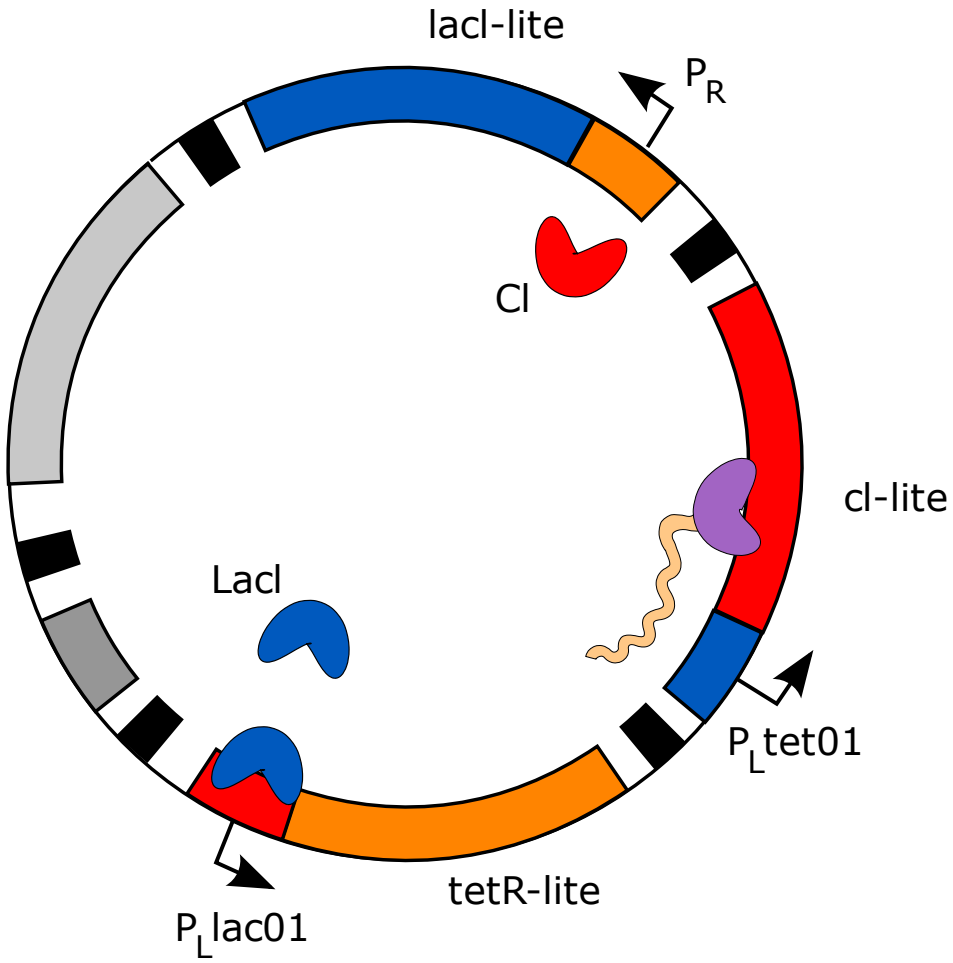

Figure: Repressilator schematic

Repressilator

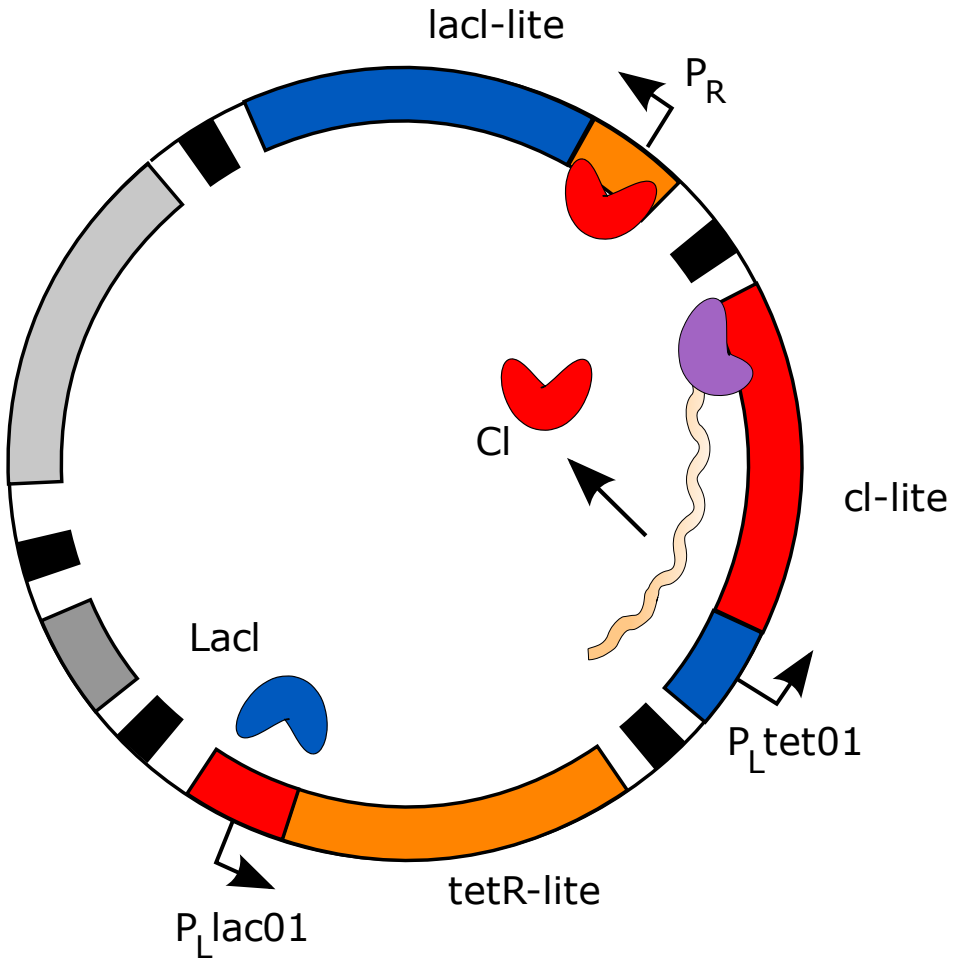

Figure: Repressilator schematic

Repressilator

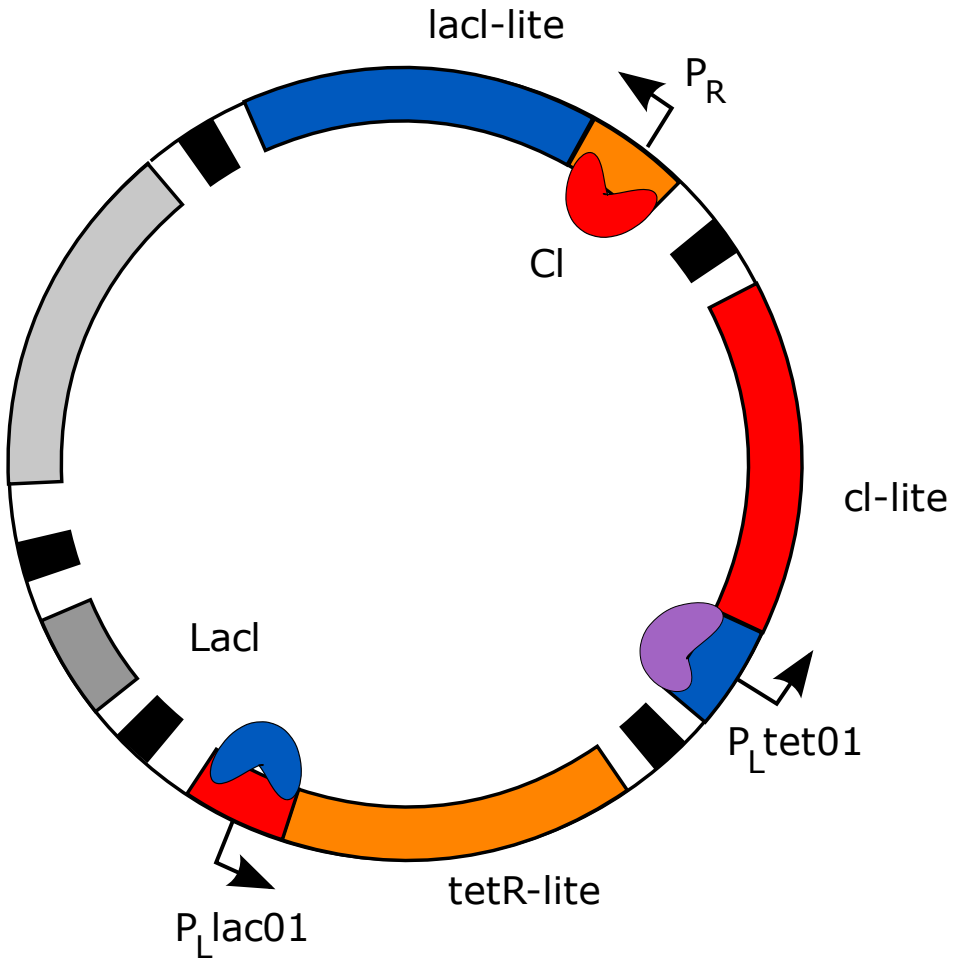

Figure: Repressilator schematic

Repressilator

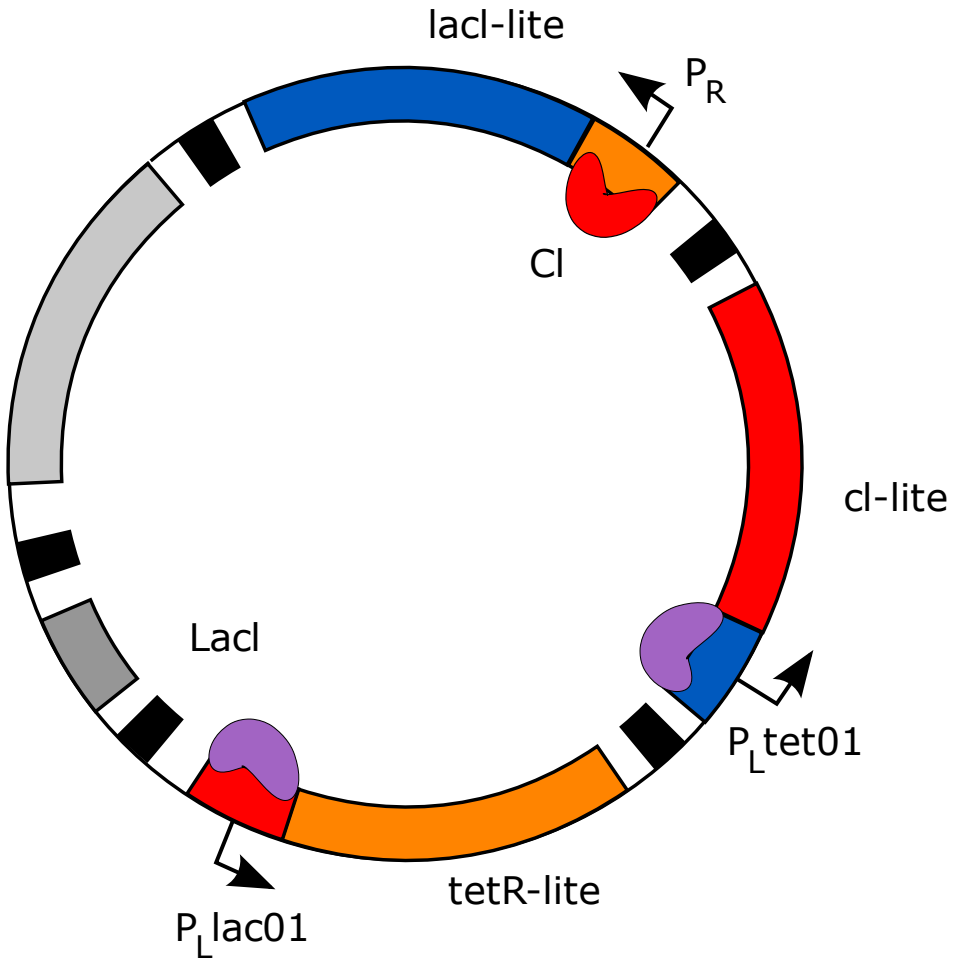

Figure: Repressilator schematic

Repressilator

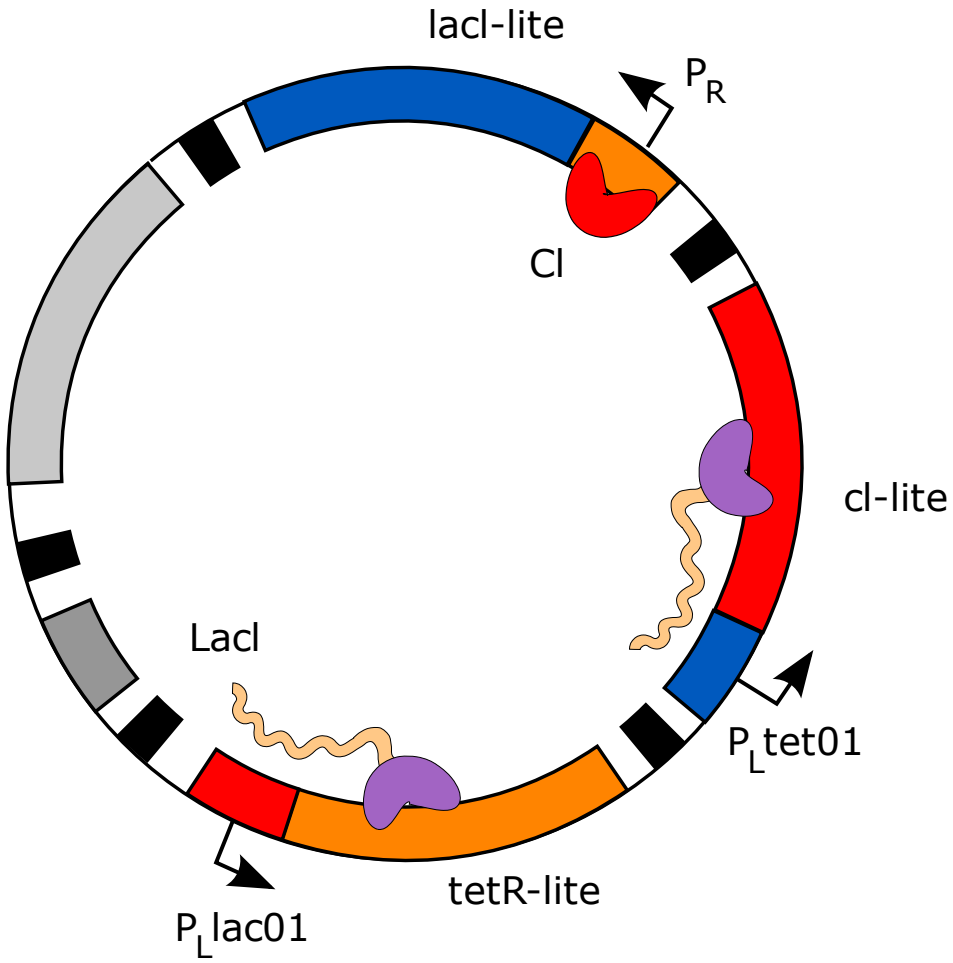

Figure: Repressilator schematic

Repressilator

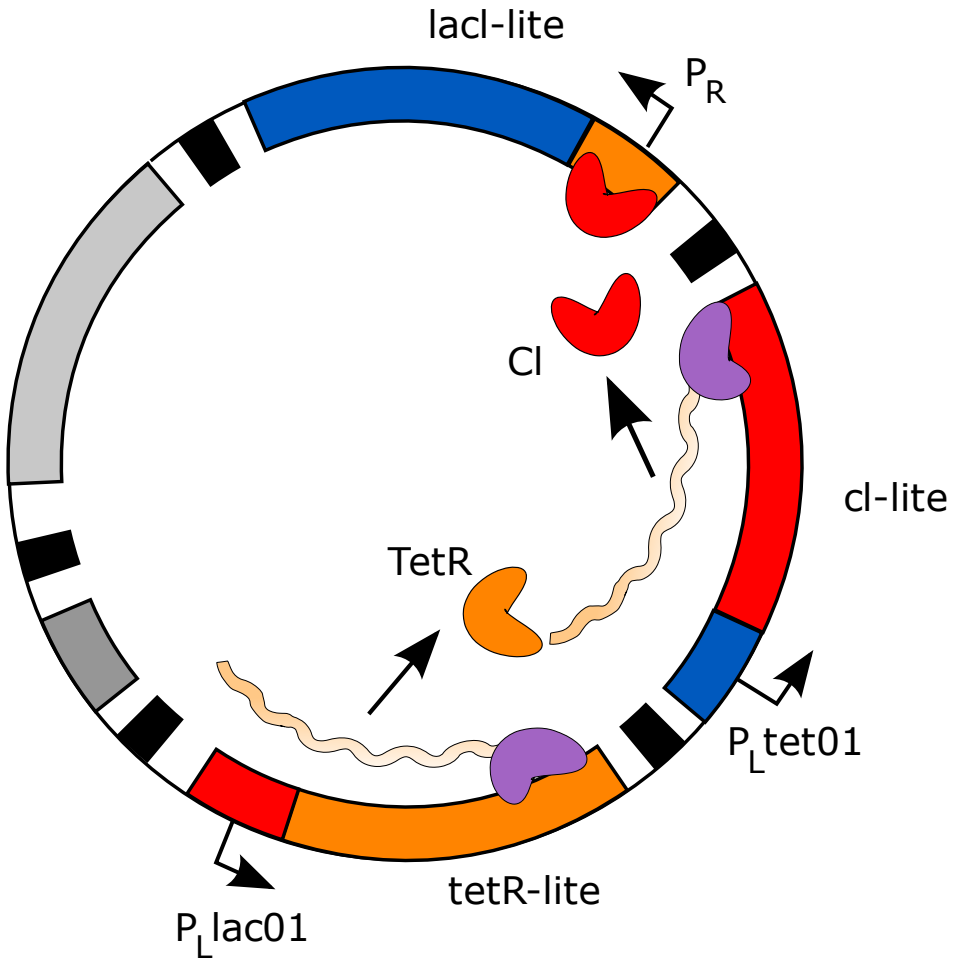

Figure: Repressilator schematic

Repressilator

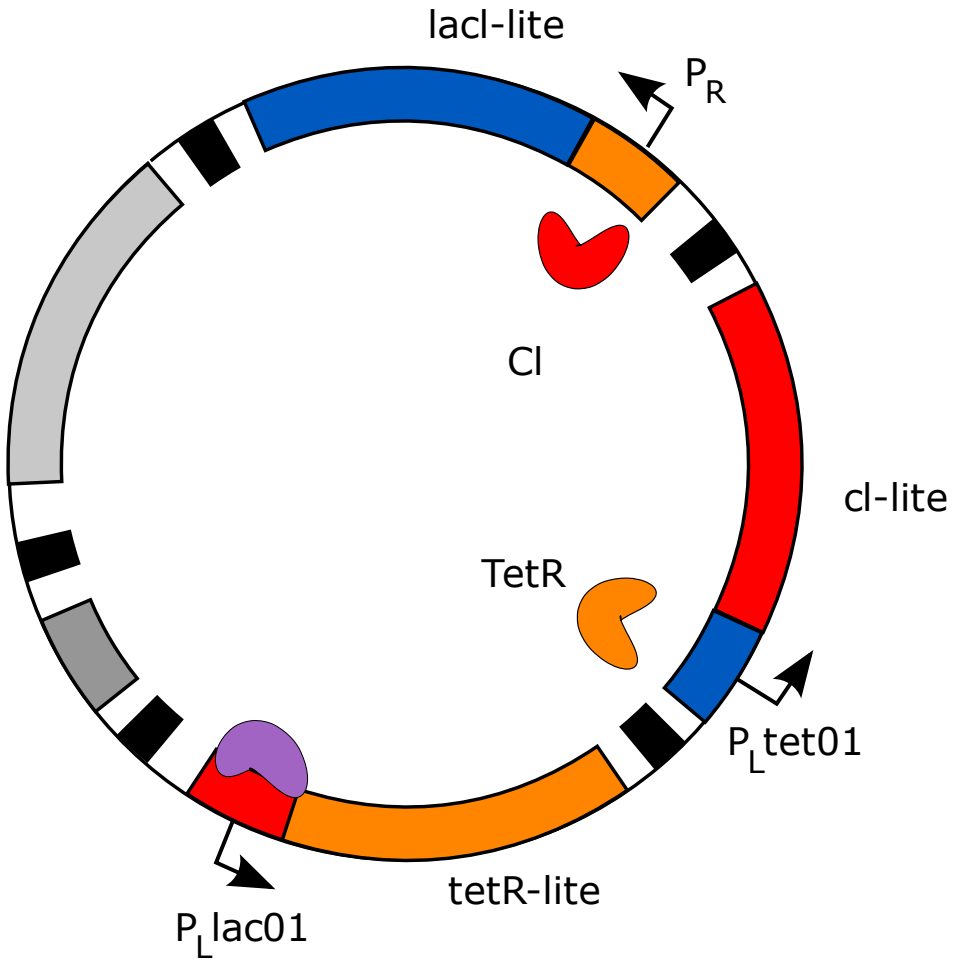

Figure: Repressilator schematic

Repressilator

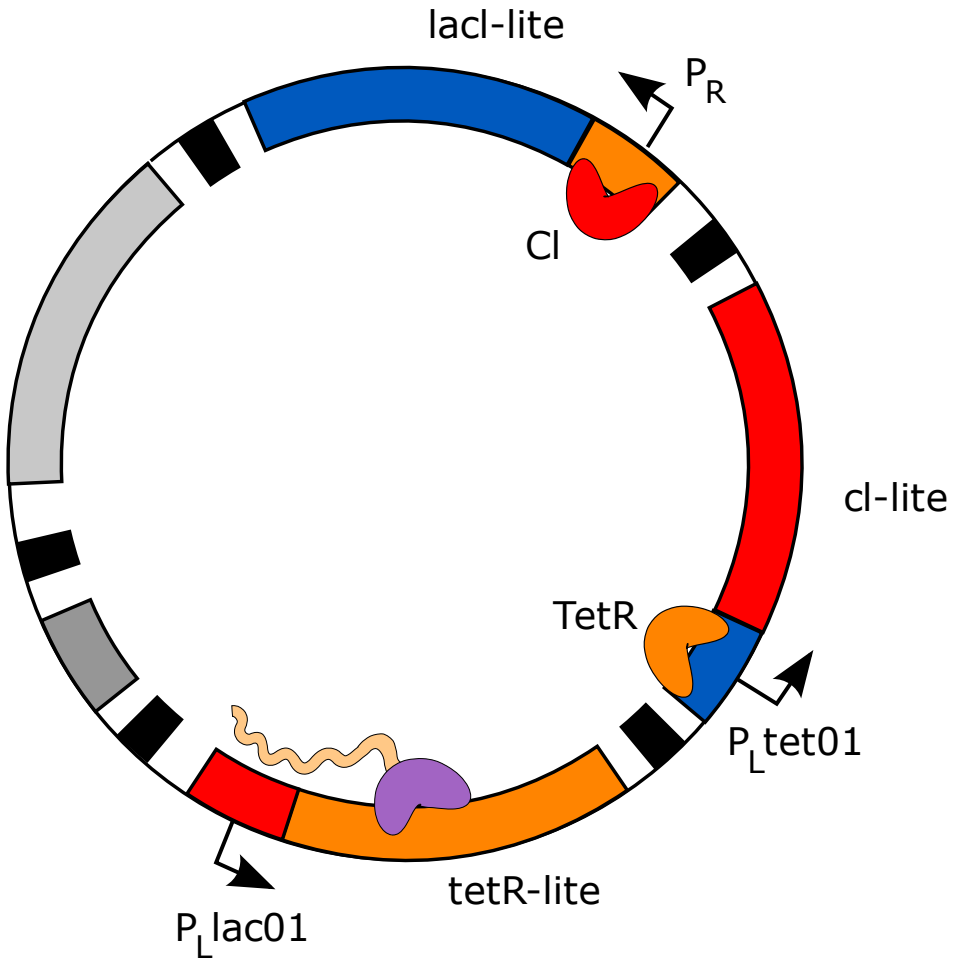

Figure: Repressilator schematic

Repressilator

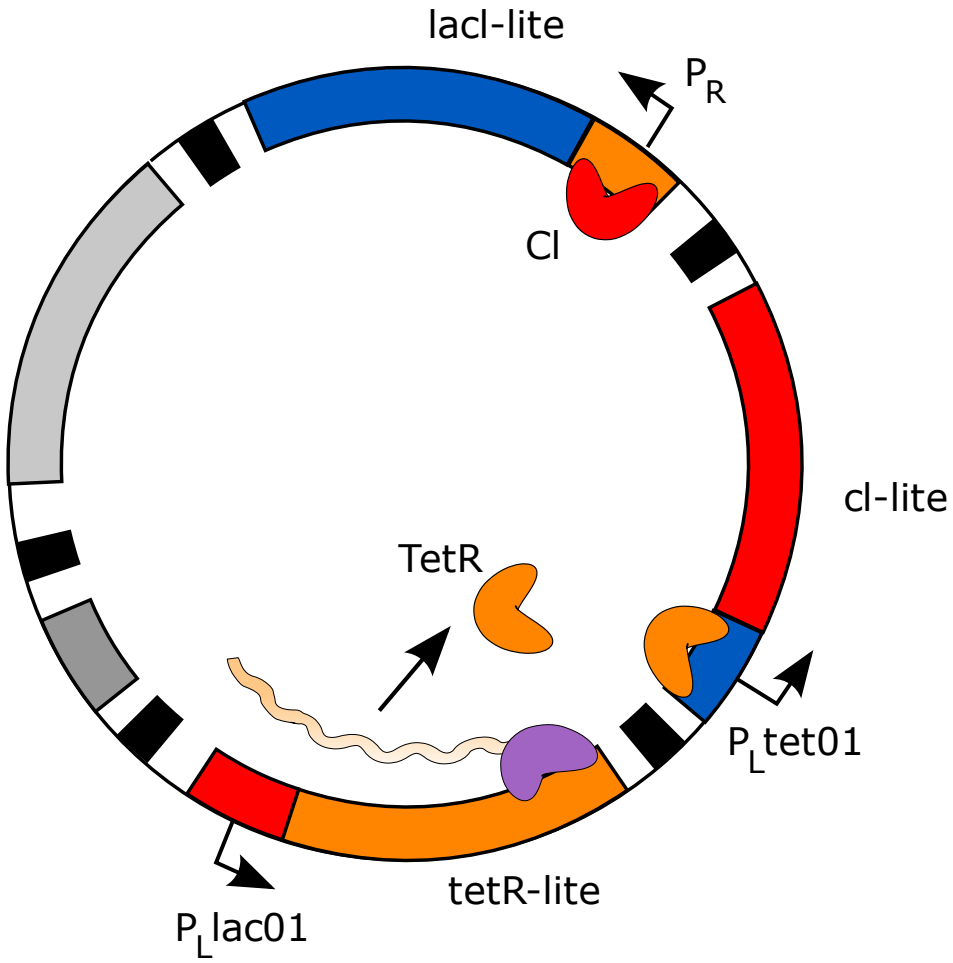

Figure: Repressilator schematic

Repressilator

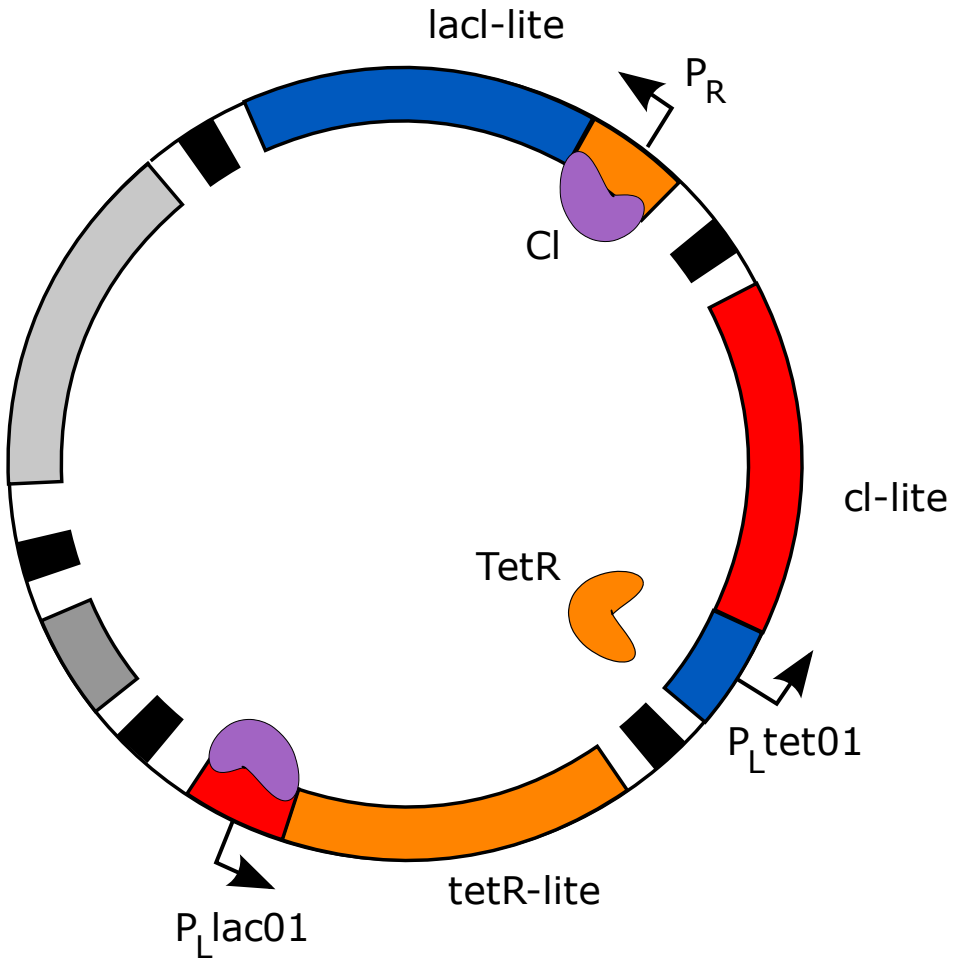

Figure: Repressilator schematic

Repressilator

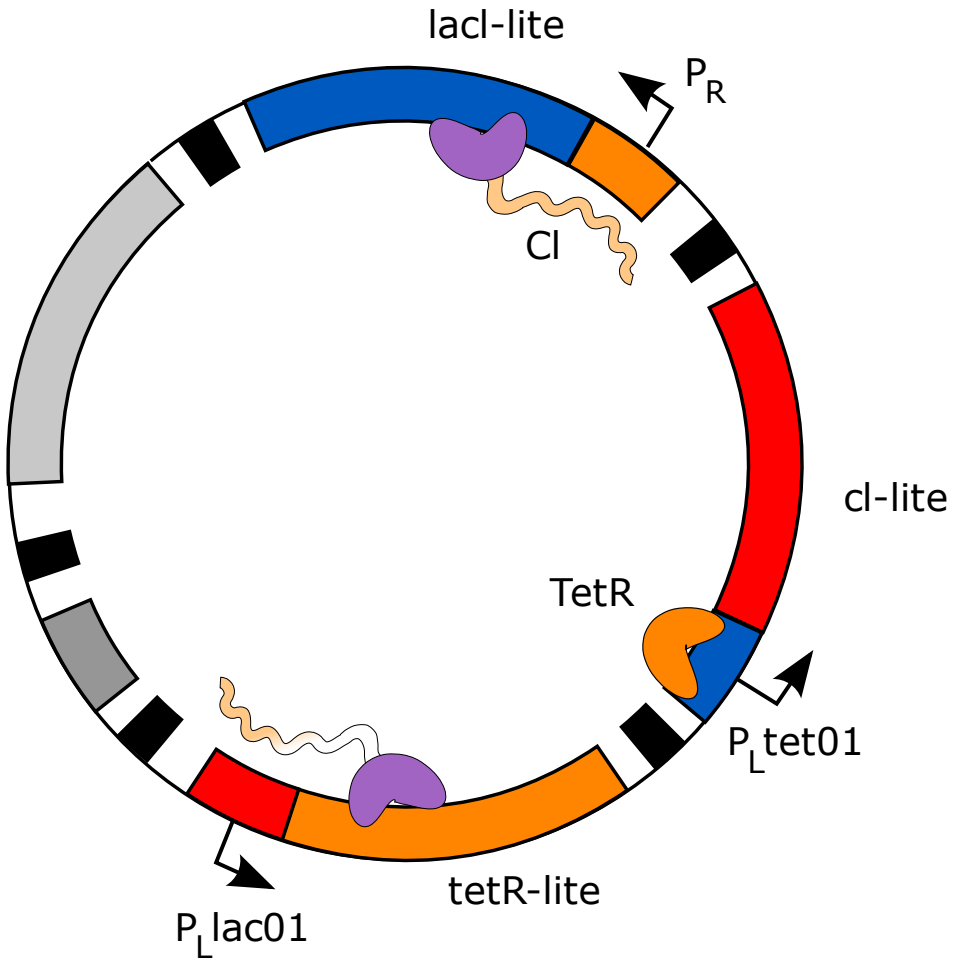

Figure: Repressilator schematic

Repressilator

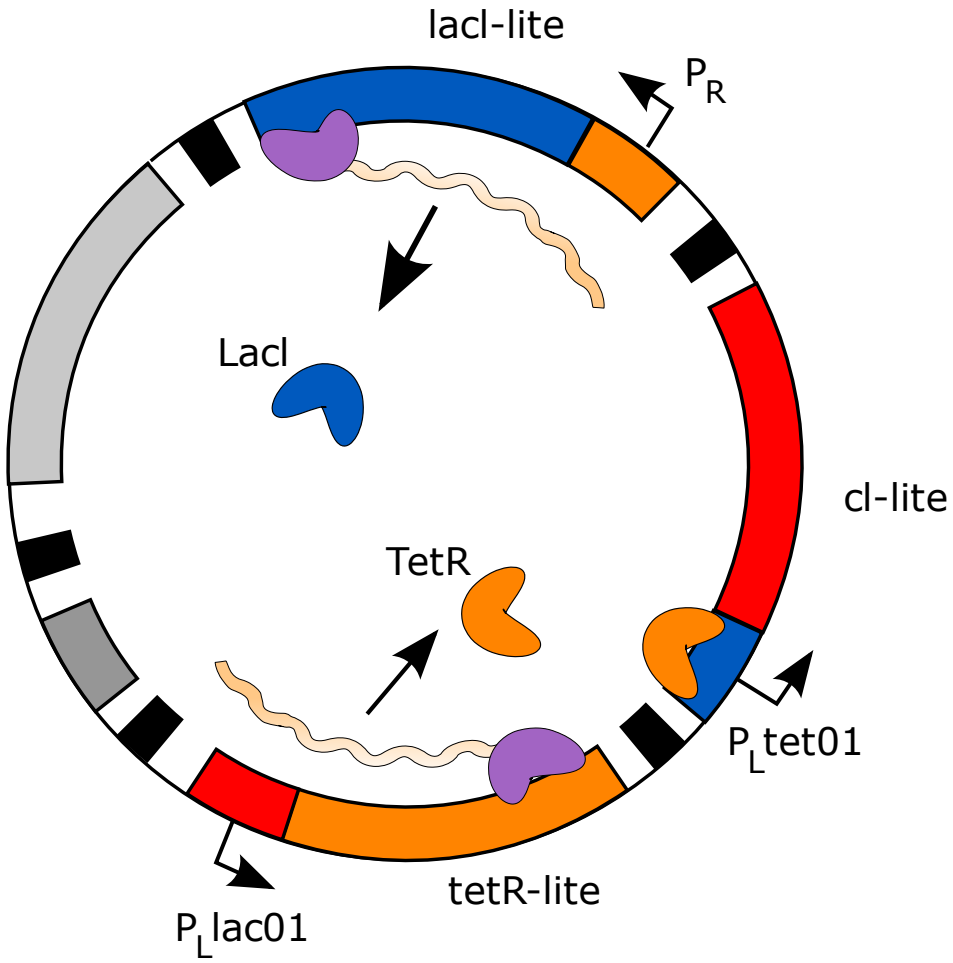

Figure: Repressilator schematic

Repressilator

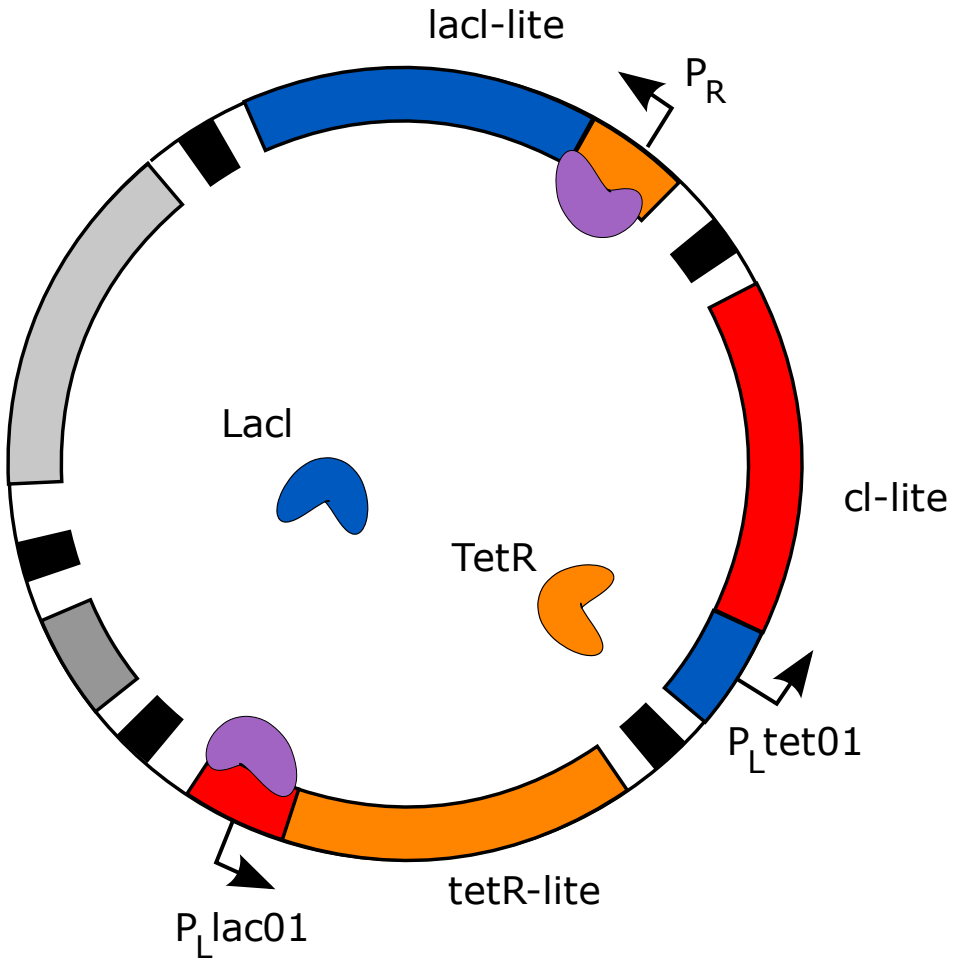

Figure: Repressilator schematic

Repressilator

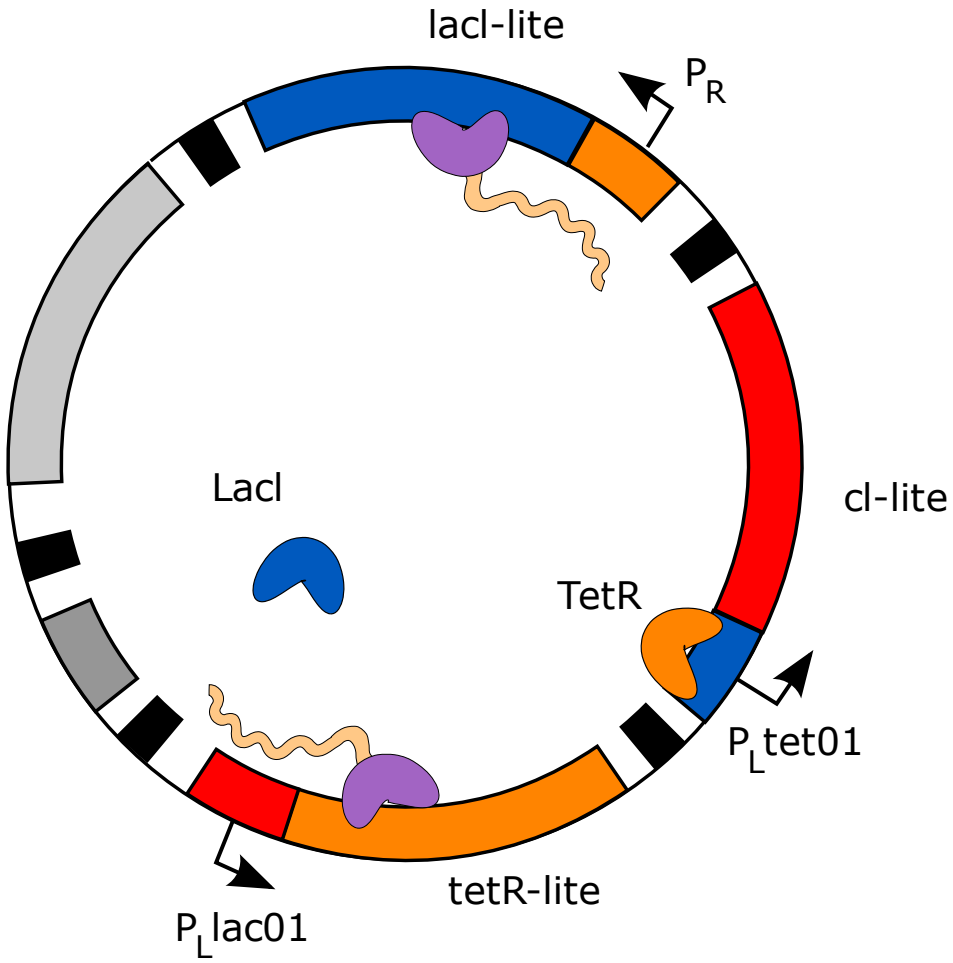

Figure: Repressilator schematic

Repressilator

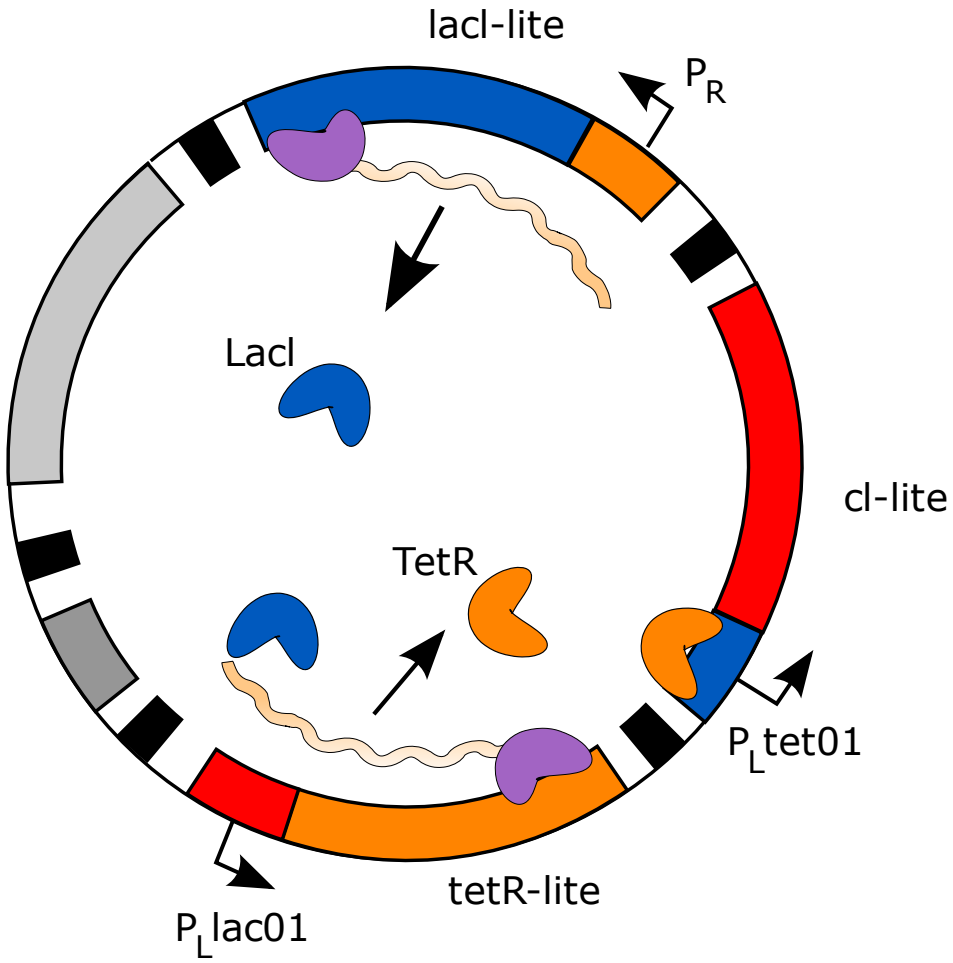

Figure: Repressilator schematic

Repressilator Results

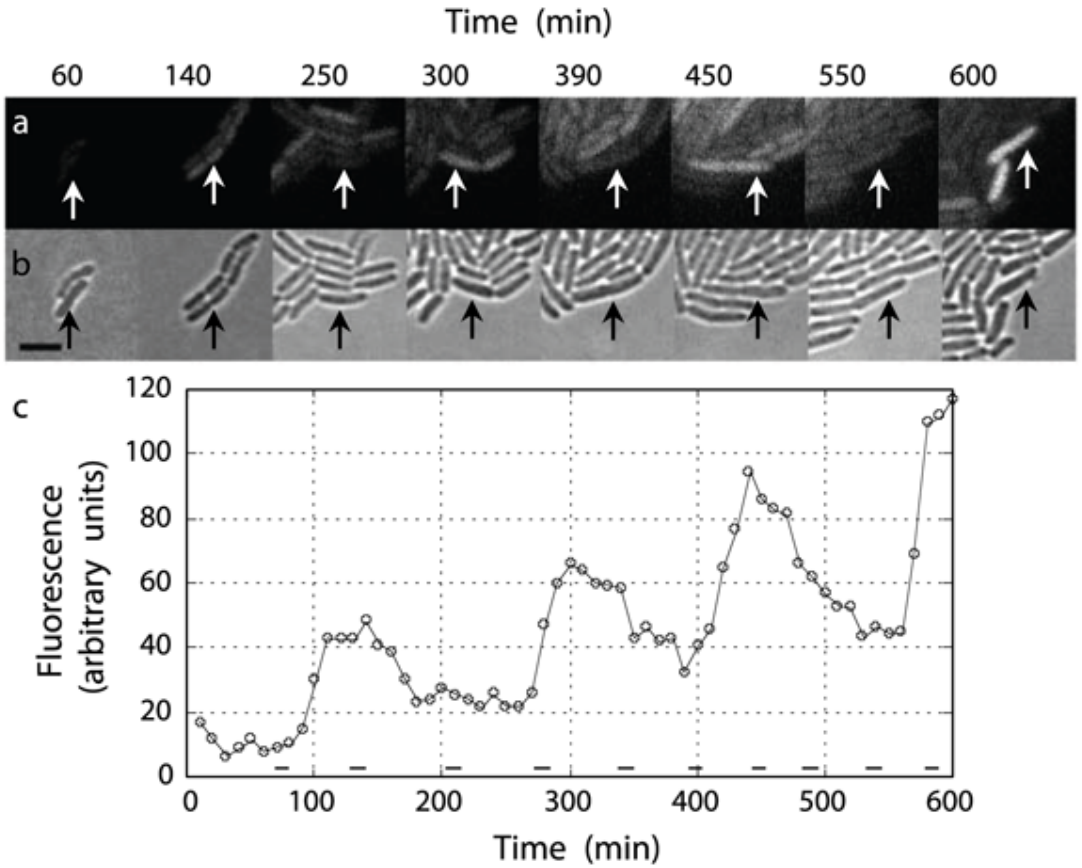

Figure: Observations of GFP. Source `http://en.wikipedia.org/wiki/Image:Repressilator_observations_1.png`

Outline

Stochastic Process and Mass Action Kinetics

Further reading: Wilkinson (2006, Chapters 1 and 6)

- Mass action kinetics — reaction occurs when relevant molecules *collide*.
- Probability of any given reaction, i, occuring in a given instant interval of time $\mathrm{d}t$ is given by $h_i \mathrm{d}t + o(\mathrm{d}t)$.
  - Where h_i is a rate law or hazard function. It is dependent on the current state of the system and c_i a stochastic rate constant.
- Represent a reaction in the form

$$X_1 + X_2 \rightarrow X_3 + X_4$$

where X_1 and X_2 are the *reactants* and X_3 and X_4 are the *products*. Denote numbers of each species by x_1, x_2, x_3 and x_4. State of the system given by vector $\mathbf{x}$.

Stochastic Process and Mass Action Kinetics

Further reading: Wilkinson (2006, Chapters 1 and 6)

- Mass action kinetics — reaction occurs when relevant molecules *collide*.
- Probability of any given reaction, i, occuring in a given instant interval of time $\mathrm{d}t$ is given by $h_i \mathrm{d}t + o(\mathrm{d}t)$.
  - Where h_i is a rate law or hazard function. It is dependent on the current state of the system and c_i a stochastic rate constant.
- Represent a reaction in the form

$$X_1 + X_2 \rightarrow X_3 + X_4$$

where X_1 and X_2 are the *reactants* and X_3 and X_4 are the *products*. Denote numbers of each species by x_1, x_2, x_3 and x_4. State of the system given by vector $\mathbf{x}$.

Stochastic Process and Mass Action Kinetics

Further reading: Wilkinson (2006, Chapters 1 and 6)

- Mass action kinetics — reaction occurs when relevant molecules *collide*.
- Probability of any given reaction, i, occuring in a given instant interval of time $\mathrm{d}t$ is given by $h_i \mathrm{d}t + o(\mathrm{d}t)$.
  - Where h_i is a rate law or hazard function. It is dependent on the current state of the system and c_i a stochastic rate constant.
- Represent a reaction in the form

$$X_1 + X_2 \rightarrow X_3 + X_4$$

where X_1 and X_2 are the *reactants* and X_3 and X_4 are the *products*. Denote numbers of each species by x_1, x_2, x_3 and x_4. State of the system given by vector $\mathbf{x}$.

Stochastic Process and Mass Action Kinetics

Further reading: Wilkinson (2006, Chapters 1 and 6)

- Mass action kinetics — reaction occurs when relevant molecules *collide*.
- Probability of any given reaction, i, occuring in a given instant interval of time $\mathrm{d}t$ is given by $h_i \mathrm{d}t + o(\mathrm{d}t)$.
  - Where h_i is a rate law or hazard function. It is dependent on the current state of the system and c_i a stochastic rate constant.
- Represent a reaction in the form

$$X_1 + X_2 \rightarrow X_3 + X_4$$

where X_1 and X_2 are the *reactants* and X_3 and X_4 are the *products*. Denote numbers of each species by x_1, x_2, x_3 and x_4. State of the system given by vector $\mathbf{x}$.

Stochastic Mass Action Kinetics

- Zeroth order:

$$0 \rightarrow X_1$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i \mathrm{d}t$

- First order (*e.g.* decay):

$$X_1 \rightarrow 0$$

probability of this rection in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 \mathrm{d}t$

- Second order:

$$X_1 + X_2 \rightarrow X_3,$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 x_2 \mathrm{d}t$.

- For individual reaction, waiting time, τ_i, is sampled from $p(\tau_i) = h_i \exp(-h_i \tau_i)$.

Stochastic Mass Action Kinetics

- Zeroth order:

$$0 \rightarrow X_1$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i \mathrm{d}t$

- First order (*e.g.* decay):

$$X_1 \rightarrow 0$$

probability of this rection in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 \mathrm{d}t$

- Second order:

$$X_1 + X_2 \rightarrow X_3,$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 x_2 \mathrm{d}t$.

- For individual reaction, waiting time, τ_i, is sampled from $p(\tau_i) = h_i \exp(-h_i \tau_i)$.

Stochastic Mass Action Kinetics

- Zeroth order:

$$0 \rightarrow X_1$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i \mathrm{d}t$

- First order (*e.g.* decay):

$$X_1 \rightarrow 0$$

probability of this rection in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 \mathrm{d}t$

- Second order:

$$X_1 + X_2 \rightarrow X_3,$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 x_2 \mathrm{d}t$.

- For individual reaction, waiting time, τ_i, is sampled from $p(\tau_i) = h_i \exp(-h_i \tau_i)$.

Stochastic Mass Action Kinetics

- Zeroth order:

$$0 \rightarrow X_1$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i \mathrm{d}t$

- First order (*e.g.* decay):

$$X_1 \rightarrow 0$$

probability of this rection in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 \mathrm{d}t$

- Second order:

$$X_1 + X_2 \rightarrow X_3,$$

probability of this reaction in interval $\mathrm{d}t$ is $h_i \mathrm{d}t = c_i x_1 x_2 \mathrm{d}t$.

- For individual reaction, waiting time, τ_i, is sampled from $p(\tau_i) = h_i \exp(-h_i \tau_i)$.

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}$, $\tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}$, $\tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}$, $\tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}$, $\tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}, \quad \tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}, \quad \tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}, \quad \tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}, \quad \tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}, \quad \tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Combining Reactions

- Typical system has multiple reactions at the same time.
- The hazard is a "rate" parameter — if there were no other reactions waiting time until the reaction would be given by an exponential.
- In practice there are other reactions and associated hazards, $\mathbf{h} = \{h_j\}_{j=1}^{M}$.
- Each reaction can affect all other hazard functions, $h_i(\mathbf{x}, c_i)$.
- Sample from the system (Gillespie's *first reaction* method):
  1. Sample time of next reaction from all reactions: $\{\tau_i\}_{i=1}^{M}, \quad \tau_i \sim h_i \exp(-h_i \tau_i)$.
  2. Find next reaction $\mu = \text{argmin}_i \tau_i$.
  3. Update state of system, $\mathbf{x}$, according to rule for that reaction.
  4. Recompute vector of hazards, $\mathbf{h}$.
  5. Repeat

Gillespie's Direct Method

- Previous sampling scheme: M random numbers (1 for each reaction).
- Exploit properties of exponential:
  - τ_j is the minimum value from $\{\tau_i\}_{i=1}^M$ sampled from different exponentials with rates $\{h_i\}_{i=1}^M$.
  - This implies: $\tau_j \sim h_0 \exp(-h_0 \tau_j)$ where $h_0 = \sum_{i=1}^M h_i$ and is known as the *combined reaction hazard*.
  - *i.e.* in each small time interval probability of any reaction is $h_0 \mathrm{d}t$.
- The probability of it having arisen from the jth reaction is given by

$$\frac{h_j}{h_0}$$

cf superposition of Poisson processes.

Gillespie's Direct Method

- Previous sampling scheme: M random numbers (1 for each reaction).
- Exploit properties of exponential:
  - τ_j is the minimum value from $\{\tau_i\}_{i=1}^M$ sampled from different exponentials with rates $\{h_i\}_{i=1}^M$.
  - This implies: $\tau_j \sim h_0 \exp(-h_0\tau_j)$ where $h_0 = \sum_{i=1}^M h_i$ and is known as the *combined reaction hazard.*
  - *i.e.* in each small time interval probability of any reaction is $h_0 \mathrm{d}t$.
- The probability of it having arisen from the jth reaction is given by

$$\frac{h_j}{h_0}$$

cf superposition of Poisson processes.

Gillespie's Direct Method

- Previous sampling scheme: M random numbers (1 for each reaction).
- Exploit properties of exponential:
  - τ_j is the minimum value from $\{\tau_i\}_{i=1}^M$ sampled from different exponentials with rates $\{h_i\}_{i=1}^M$.
  - This implies: $\tau_j \sim h_0 \exp(-h_0 \tau_j)$ where $h_0 = \sum_{i=1}^M h_i$ and is known as the *combined reaction hazard.*
  - *i.e.* in each small time interval probability of any reaction is $h_0 \mathrm{d}t$.
- The probability of it having arisen from the jth reaction is given by

$$\frac{h_j}{h_0}$$

cf superposition of Poisson processes.

Gillespie's Direct Method

- Previous sampling scheme: M random numbers (1 for each reaction).
- Exploit properties of exponential:
  - τ_j is the minimum value from $\{\tau_i\}_{i=1}^M$ sampled from different exponentials with rates $\{h_i\}_{i=1}^M$.
  - This implies: $\tau_j \sim h_0 \exp(-h_0 \tau_j)$ where $h_0 = \sum_{i=1}^M h_i$ and is known as the *combined reaction hazard.*
  - *i.e.* in each small time interval probability of any reaction is $h_0 \mathrm{d}t$.
- The probability of it having arisen from the jth reaction is given by

$$\frac{h_j}{h_0}$$

cf superposition of Poisson processes.

Gillespie's Direct Method

- Previous sampling scheme: M random numbers (1 for each reaction).
- Exploit properties of exponential:
  - τ_j is the minimum value from $\{\tau_i\}_{i=1}^M$ sampled from different exponentials with rates $\{h_i\}_{i=1}^M$.
  - This implies: $\tau_j \sim h_0 \exp(-h_0 \tau_j)$ where $h_0 = \sum_{i=1}^M h_i$ and is known as the *combined reaction hazard.*
  - *i.e.* in each small time interval probability of any reaction is $h_0 \mathrm{d}t$.
- The probability of it having arisen from the jth reaction is given by

$$\frac{h_j}{h_0}$$

cf superposition of Poisson processes.

Gillespie's Direct Method

- Previous sampling scheme: M random numbers (1 for each reaction).
- Exploit properties of exponential:
  - τ_j is the minimum value from $\{\tau_i\}_{i=1}^M$ sampled from different exponentials with rates $\{h_i\}_{i=1}^M$.
  - This implies: $\tau_j \sim h_0 \exp(-h_0 \tau_j)$ where $h_0 = \sum_{i=1}^M h_i$ and is known as the *combined reaction hazard.*
  - *i.e.* in each small time interval probability of any reaction is $h_0 \mathrm{d}t$.
- The probability of it having arisen from the jth reaction is given by

$$\frac{h_j}{h_0}$$

cf superposition of Poisson processes.

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \rightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \rightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \rightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \rightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \rightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \longrightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \longrightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp(-h_0 \tau_\mu)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \longrightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O(M)$.
- Can do in $O(\log M)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Sampling the System

Gillespie Direct Method

1. Compute the hazards, $\mathbf{h}$.
2. Sample time of next reaction from $\tau_\mu \sim h_0 \exp\left(-h_0 \tau_\mu\right)$
3. Determine which reaction it was: sample μ from a multinomial with probabilities given by $\frac{h_j}{h_0}$.
4. Update the state of the system, $\mathbf{x}$.
5. Increment time $t \longrightarrow t + \tau_\mu$.
6. Repeat until simulation time complete.

- This is $O\left(M\right)$.
- Can do in $O\left(\log M\right)$ — use a *dependency* graph to determine when things need calculation Gibson and Bruck (2000).

Repressilator Simulation

Translation:

$$\text{mlac} \longrightarrow \text{Lac} + \text{mlac}$$

First order reaction of mRNA from *lac* gene to protein plus mRNA from *lac* gene.

Repressilator Simulation

mRNA decay:

$$\text{mlac} \longrightarrow 0$$

First order reaction of mRNA from *lac* gene.

Repressilator Simulation

Protein decay:

$$\text{Lac} \rightarrow 0$$

First order reaction of Lac protein.

Repressilator Simulation

Transcription:

$$\text{lac} + \text{RNAP} \rightarrow \text{mlac} + \text{RNAP} + \text{lac}$$

Second order reaction of *lac* gene and RNA polymerase to *lac* mRNA, *lac* gene and RNA polymerase.

Repressilator Simulation

Protein (TF) bound to promoter:

$$\mathrm{CI} + \mathrm{lac} \rightarrow \mathrm{CI} \cdot \mathrm{lac}$$

Second order reaction, TF protein (CI) from another gene binds to *lac* promoter (represented by the gene). This prevents transcription.

Repressilator Simulation

Protein unbinds from promoter:

$$\mathrm{CI} \cdot \mathrm{lac} \rightarrow \mathrm{CI} + \mathrm{lac}$$

First order reaction, TF protein and *lac* promoter region unbind, allowing transcription to take place.

Other Implementation Details

- The effect of each reaction is stored in a matrix $\mathbf{S}$, the stoichiometry matrix.
- A row of this matrix is added to the state vector, $\mathbf{x}$, to account for effects from each reaction.

Simulation Result

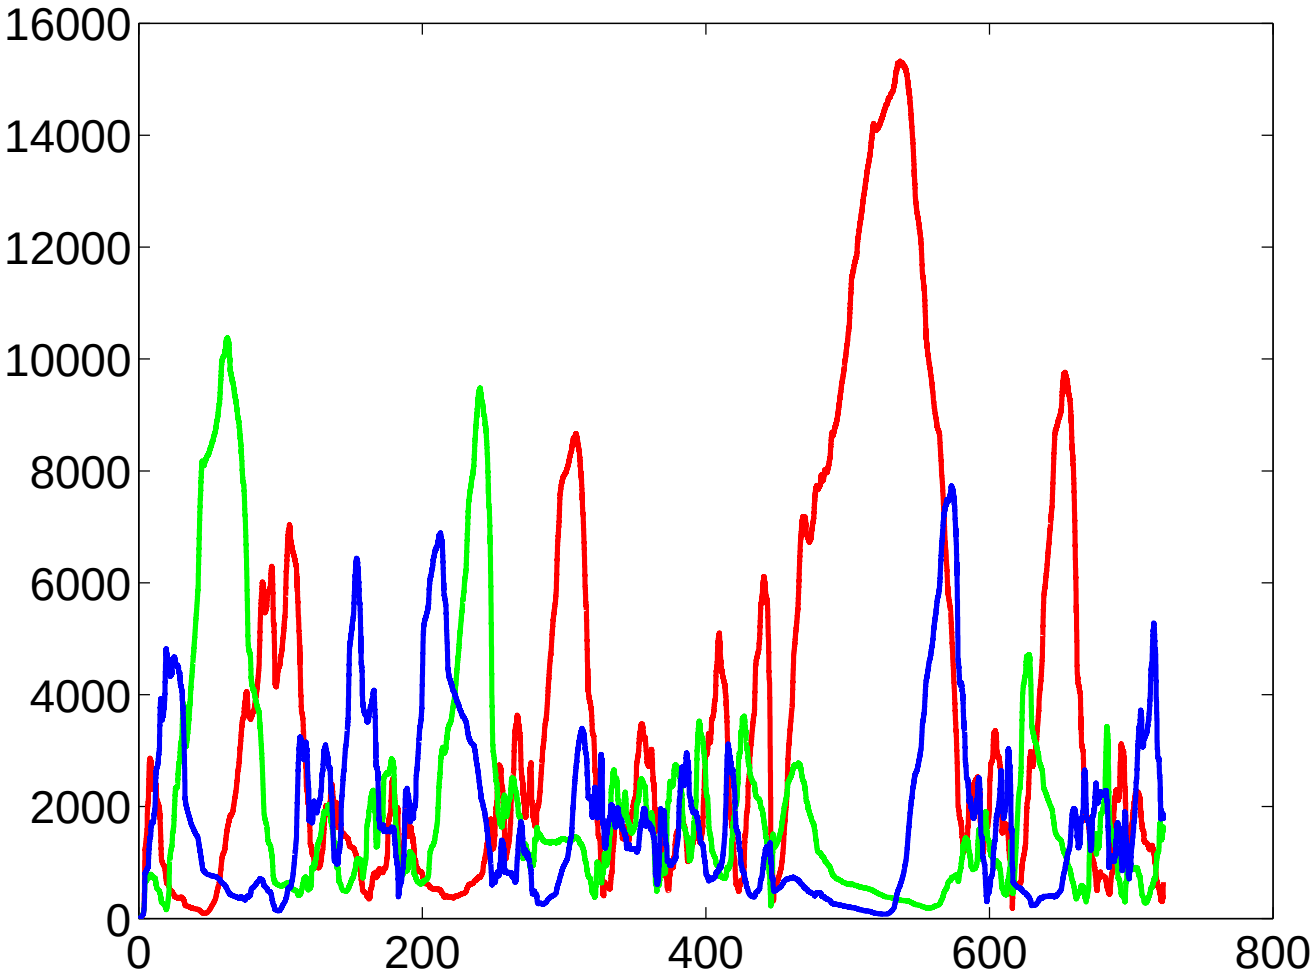

Figure: Simulation of repressilator using Gillespie algorithm.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If $\mathbf{x}$ is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If $\mathbf{x}$ is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If x is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If x is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If x is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If $\mathbf{x}$ is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If $\mathbf{x}$ is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If $\mathbf{x}$ is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

What Next?

- Simulation from the system assumes we know *structure* (stoichiometric matrix, $\mathbf{S}$) and parameters (stochastic rate parameters, $\mathbf{c}$).
- Structure *may* be known or assumed.
- Specifying parameters is more complex.
  - In chemistry *in vitro* measurements can be made.
  - In biology this is more difficult and perhaps less valid.
- Can we do learning? — **this is where we come in!**
  - If $\mathbf{x}$ is observed directly in v. high time resolution: yes.
  - In practice it is indirectly observed in lower time resolution.
- Learning in stochastic systems is difficult as marginalisation of these unknowns is required.

A Deterministic Approximation

- Approximate the stochastic system by dealing in *deterministic* concentrations.
- In *chemistry* concentrations involve large numbers, and the approximation is good.
- In *biology* this is less true.
- For Mass Action Kinetics:

$$X_1 + X_2 \rightarrow X_3$$

$$X_3 \rightarrow 0$$

leads to

$$\frac{\mathrm{d}\left[X_3\right]}{\mathrm{d}t} = k_1\left[X_1\right]\left[X_2\right] - k_2\left[X_3\right]$$

with $[X_i]$ representing concentration of species X_i.

A Deterministic Approximation

- Approximate the stochastic system by dealing in *deterministic* concentrations.
- In *chemistry* concentrations involve large numbers, and the approximation is good.
- In *biology* this is less true.
- For Mass Action Kinetics:

$$X_1 + X_2 \rightarrow X_3$$

$$X_3 \rightarrow 0$$

leads to

$$\frac{\mathrm{d}\left[X_3\right]}{\mathrm{d}t} = k_1\left[X_1\right]\left[X_2\right] - k_2\left[X_3\right]$$

with $[X_i]$ representing concentration of species X_i.

A Deterministic Approximation

- Approximate the stochastic system by dealing in *deterministic* concentrations.
- In *chemistry* concentrations involve large numbers, and the approximation is good.
- In *biology* this is less true.
- For Mass Action Kinetics:

$$X_1 + X_2 \rightarrow X_3$$

$$X_3 \rightarrow 0$$

leads to

$$\frac{\mathrm{d}\left[X_3\right]}{\mathrm{d}t} = k_1\left[X_1\right]\left[X_2\right] - k_2\left[X_3\right]$$

with $[X_i]$ representing concentration of species X_i.

A Deterministic Approximation

- Approximate the stochastic system by dealing in *deterministic* concentrations.
- In *chemistry* concentrations involve large numbers, and the approximation is good.
- In *biology* this is less true.
- For Mass Action Kinetics:

$$X_1 + X_2 \rightarrow X_3$$

$$X_3 \rightarrow 0$$

leads to

$$\frac{\mathrm{d}\left[X_3\right]}{\mathrm{d}t} = k_1\left[X_1\right]\left[X_2\right] - k_2\left[X_3\right]$$

with $[X_i]$ representing concentration of species X_i.

Repressilator Simulation

Translation:

$$\text{mlac} \rightarrow \text{Lac} + \text{mlac}$$

$$\frac{\mathrm{d}\left[\text{Lac}\right]}{\mathrm{d}t} = -k_3\left[\text{Lac}\right] - k_4\left[\text{Lac}\right]\left[\text{mtetR}\right] + k_5\left[\text{mlac}\right] + k_6\left[\text{Lac}\cdot\text{tetR}\right]$$

First order reaction of mRNA from *lac* gene to protein plus mRNA from *lac* gene.

Repressilator Simulation

mRNA decay:

$$\text{mlac} \rightarrow 0$$

$$\frac{\mathrm{d}\,[\text{mlac}]}{\mathrm{d}t} = k_1\,[\text{RNAP}]\,[\text{lacI1}] - k_2\,[\text{mlac}]$$

First order reaction of mRNA from *lac* gene.

Repressilator Simulation

Protein decay:

$$\text{Lac} \rightarrow 0$$

$$\frac{\mathrm{d}\left[\text{Lac}\right]}{\mathrm{d}t} = -k_3\left[\text{Lac}\right] - k_4\left[\text{Lac}\right]\left[\text{mtetR}\right] + k_5\left[\text{mlac}\right] + k_6\left[\text{Lac}\cdot\text{tetR}\right]$$

First order reaction of Lac protein.

Repressilator Simulation

Transcription:

$$\text{lac} + \text{RNAP} \rightarrow \text{mlac} + \text{RNAP} + \text{lac}$$

$$\frac{\mathrm{d}\,[\text{mlac}]}{\mathrm{d}t} = k_1\,[\text{RNAP}]\,[\text{lacI}] - k_2\,[\text{mlac}]$$

Second order reaction of *lac* gene and RNA polymerase to *lac* mRNA, *lac* gene and RNA polymerase.

Repressilator Simulation

Protein (TF) bound to promoter:

$$\mathrm{CI} + \mathrm{lac} \rightarrow \mathrm{CI} \cdot \mathrm{lac}$$

$$\frac{\mathrm{d}\left[\mathrm{CI} \cdot \mathrm{lac}\right]}{\mathrm{d}t} = k_8 \left[\mathrm{CI}\right]\left[\mathrm{lac}\right] - k_{10}\left[\mathrm{CI} \cdot \mathrm{lac}\right]$$

$$\frac{\mathrm{d}\left[\mathrm{CI}\right]}{\mathrm{d}t} = -k_7\left[\mathrm{CI}\right] - k_8\left[\mathrm{CI}\right]\left[\mathrm{lac}\right] + k_9\left[\mathrm{mcI}\right] + k_{10}\left[\mathrm{CI} \cdot \mathrm{lac}\right]$$

$$\frac{\mathrm{d}\left[\mathrm{lac}\right]}{\mathrm{d}t} = -k_8\left[\mathrm{CI}\right]\left[\mathrm{lac}\right] + k_{10}\left[\mathrm{CI} \cdot \mathrm{lac}\right]$$

Second order reaction, TF protein (CI) from another gene binds to *lac* promoter (represented by the gene). This prevents transcription.

Repressilator Simulation

Protein unbinds from promoter:

$$\mathrm{CI} \cdot \mathrm{lac} \rightarrow \mathrm{CI} + \mathrm{lac}$$

$$\frac{\mathrm{d}\left[\mathrm{CI} \cdot \mathrm{lac}\right]}{\mathrm{d}t} = k_8 \left[\mathrm{CI}\right] \left[\mathrm{lac}\right] - k_{10} \left[\mathrm{CI} \cdot \mathrm{lac}\right]$$

$$\frac{\mathrm{d}\left[\mathrm{CI}\right]}{\mathrm{d}t} = -k_7 \left[\mathrm{CI}\right] - k_8 \left[\mathrm{CI}\right] \left[\mathrm{lac}\right] + k_9 \left[\mathrm{mcI}\right] + k_{10} \left[\mathrm{CI} \cdot \mathrm{lac}\right]$$

$$\frac{\mathrm{d}\left[\mathrm{lac}\right]}{\mathrm{d}t} = -k_8 \left[\mathrm{CI}\right] \left[\mathrm{lac}\right] + k_{10} \left[\mathrm{CI} \cdot \mathrm{lac}\right]$$

First order reaction, TF protein and *lac* promoter region unbind, allowing transcription to take place.

Simulated Repressilator

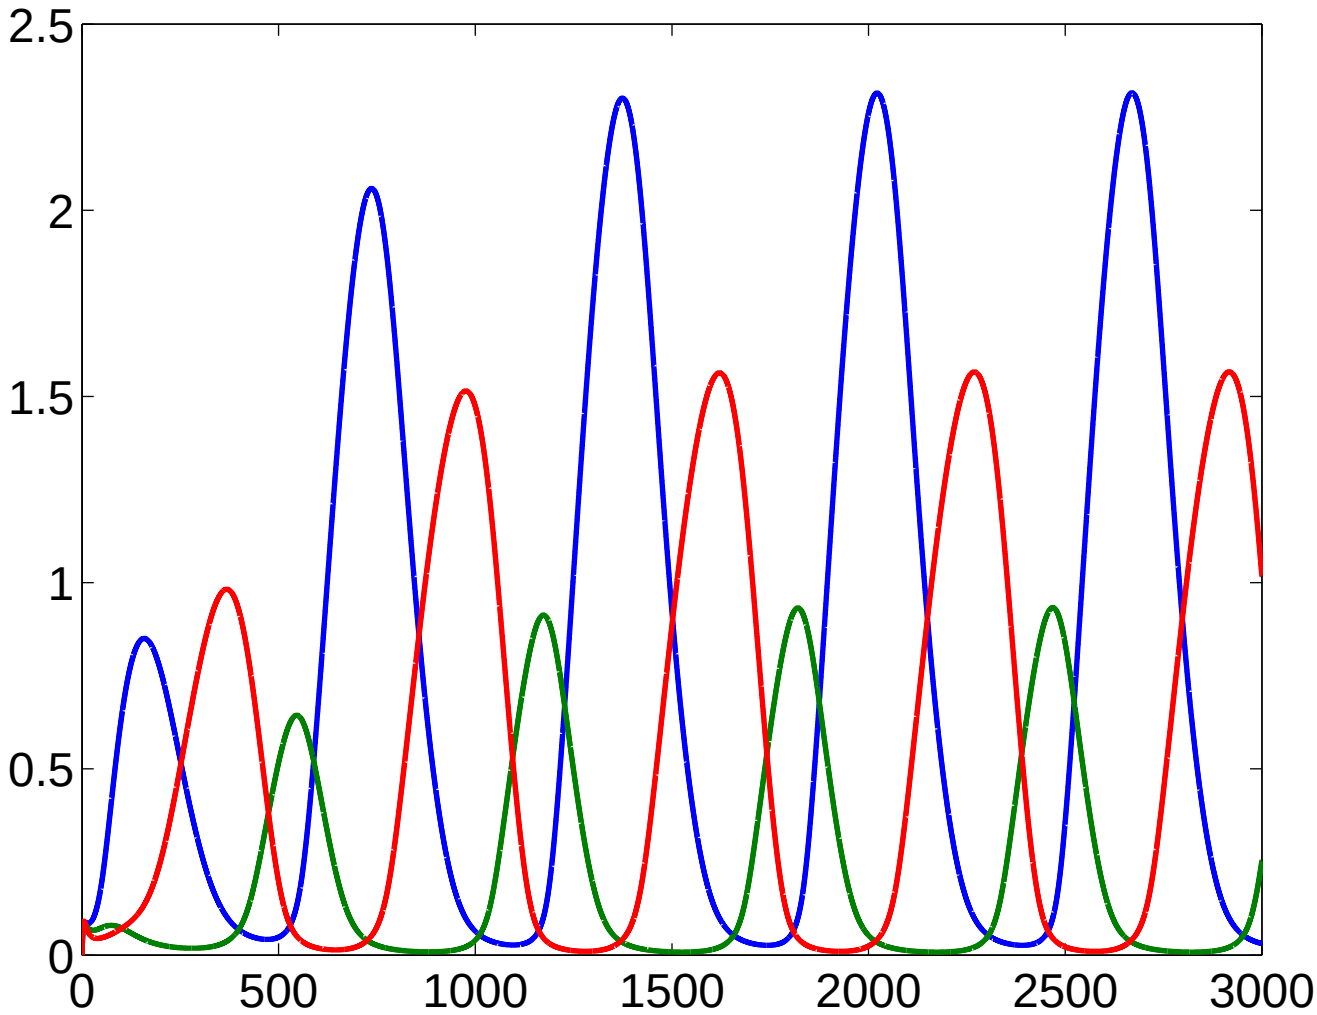

Figure: Simulation of repressilator based on ODEs from COPASI Hoops et al. (2006).

Fitting ODE Models

- Find parameters that allow model to fit a given data set.
- For given parameters and inital conditions solve the system and compare to data.
- Minimise the least squares match to the data with respect to parameters and initial conditions.
- Multimodal optimisation: tools available for fitting (COPASI Hoops et al. (2006)).
- Problems remain:
  1. **How do we deal with a missing chemical species (e.g. TF concentration)?**
     We'll look at this next and in Part II.
  2. **What to do if certain parameters aren't well identified?**
     The system outputs may be insensitive to some parameters.
  3. **If several hypothesised models exist, which should we choose?**
     We'll look briefly at this at the end if there's time.

Outline

ODE Model of Activation

- Linear Activation Model (Barenco et al., 2006, Genome Biology)

$$\frac{\mathrm{d}x_j(t)}{\mathrm{d}t} = B_j + S_j f(t) - D_j x_j(t)$$

- Slight change in notation:
  - $x_j(t)$ – concentration of gene j's mRNA
  - $f(t)$ – concentration of active transcription factor
  - Model parameters: baseline B_j, sensitivity S_j and decay D_j
  - Application: identifying co-regulated genes (targets)
  - Problem: how do we fit the model when $f(t)$ is not observed?

ODE Model of Activation

- Linear Activation Model (Barenco et al., 2006, Genome Biology)

$$\frac{\mathrm{d}x_j(t)}{\mathrm{d}t} = B_j + S_j f(t) - D_j x_j(t)$$

- Slight change in notation:
  - $x_j(t)$ – concentration of gene j's mRNA
  - $f(t)$ – concentration of active transcription factor
  - Model parameters: baseline B_j, sensitivity S_j and decay D_j
  - Application: identifying co-regulated genes (targets)
  - Problem: how do we fit the model when $f(t)$ is not observed?

ODE Model of Activation

- Linear Activation Model (Barenco et al., 2006, Genome Biology)

$$\frac{\mathrm{d}x_j(t)}{\mathrm{d}t} = B_j + S_j f(t) - D_j x_j(t)$$

- Slight change in notation:
  - $x_j(t)$ – concentration of gene j's mRNA
  - $f(t)$ – concentration of active transcription factor
  - Model parameters: baseline B_j, sensitivity S_j and decay D_j
  - Application: identifying co-regulated genes (targets)
  - Problem: how do we fit the model when $f(t)$ is not observed?

ODE Model of Activation

- Linear Activation Model (Barenco et al., 2006, Genome Biology)

$$\frac{\mathrm{d}x_j(t)}{\mathrm{d}t} = B_j + S_j f(t) - D_j x_j(t)$$

- Slight change in notation:
  - $x_j(t)$ – concentration of gene j's mRNA
  - $f(t)$ – concentration of active transcription factor
  - Model parameters: baseline B_j, sensitivity S_j and decay D_j
  - Application: identifying co-regulated genes (targets)
  - Problem: how do we fit the model when $f(t)$ is not observed?

ODE Model of Activation

- Linear Activation Model (Barenco et al., 2006, Genome Biology)

$$\frac{\mathrm{d}x_j(t)}{\mathrm{d}t} = B_j + S_j f(t) - D_j x_j(t)$$

- Slight change in notation:
  - $x_j(t)$ – concentration of gene j's mRNA
  - $f(t)$ – concentration of active transcription factor
  - Model parameters: baseline B_j, sensitivity S_j and decay D_j
  - Application: identifying co-regulated genes (targets)
  - Problem: how do we fit the model when $f(t)$ is not observed?

ODE Model of Activation

- Linear Activation Model (Barenco et al., 2006, Genome Biology)

$$\frac{\mathrm{d}x_j(t)}{\mathrm{d}t} = B_j + S_j f(t) - D_j x_j(t)$$

- Slight change in notation:
  - $x_j(t)$ – concentration of gene j's mRNA
  - $f(t)$ – concentration of active transcription factor
  - Model parameters: baseline B_j, sensitivity S_j and decay D_j
  - Application: identifying co-regulated genes (targets)
  - Problem: how do we fit the model when $f(t)$ is not observed?

Why use a model-based approach?

- Model based approach to co-regulated targets ...
  - clustering is often used but,
  - co-regulated genes can differ greatly in their expression profiles

Why use a model-based approach?

- Model based approach to co-regulated targets ...
  - clustering is often used but,
  - co-regulated genes can differ greatly in their expression profiles

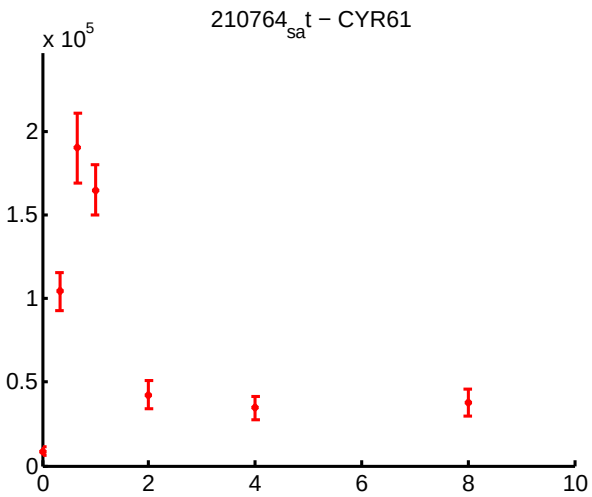

Why use a model-based approach?

- Model based approach to co-regulated targets ...
  - clustering is often used but,
  - co-regulated genes can differ greatly in their expression profiles

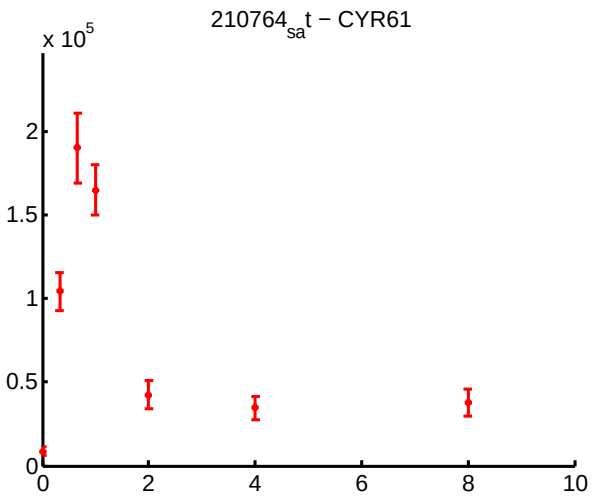

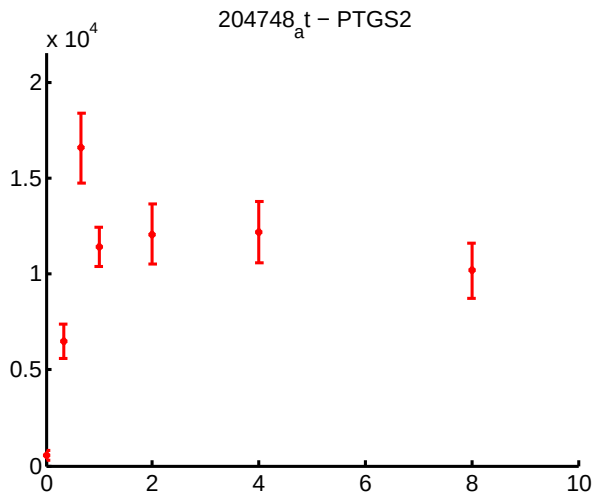

Why use a model-based approach?

- Model based approach to co-regulated targets ...
  - clustering is often used but,
  - co-regulated genes can differ greatly in their expression profiles

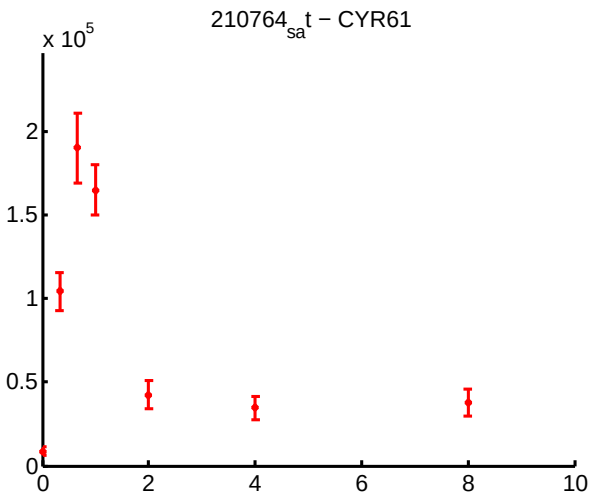

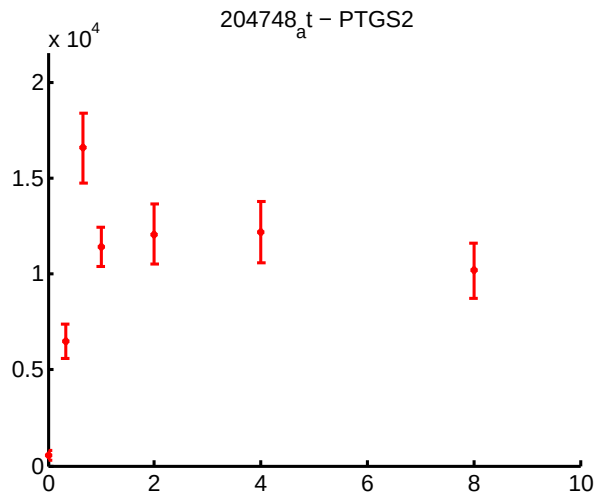

- Clustering cannot be relied on to identify co-regulated genes

Why use a model-based approach?

- Model based approach to co-regulated targets ...
  - clustering is often used but,
  - co-regulated genes can differ greatly in their expression profiles

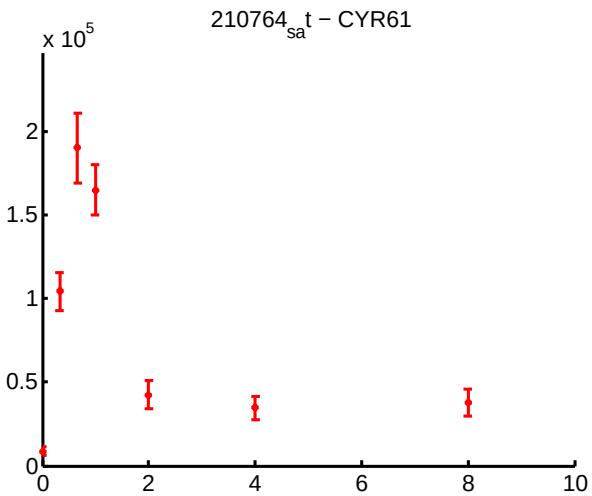

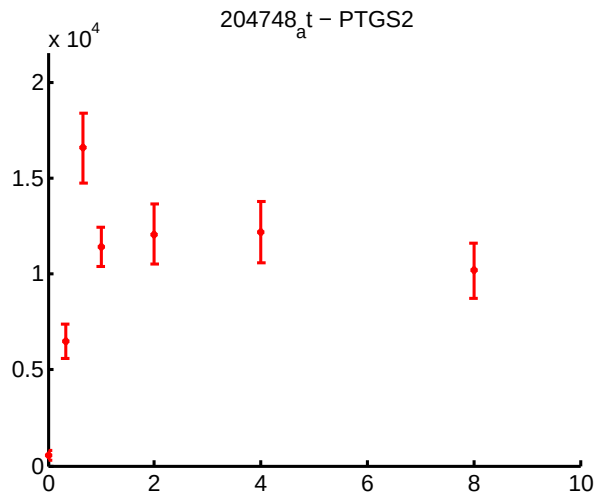

- Clustering cannot be relied on to identify co-regulated genes
- A model-based approach is required

Cell Damage

- Radiation damages molecules in the cell.
- Most of this damage is quickly repaired — single strand breaks, backbone break.
- Double strand breaks are more serious — a complete disconnect along the chromosome.
- Cell cycle stages:
  - G_1: Cell is not dividing.
  - G_2: Cell is preparing for meitosis, chromosomes have divided.
  - S: Cell is undergoing meitosis (DNA synthesis).
- Main problem is in G1. In G2 there are two copies of the chromosome. In G1 only one copy.

p53 "Guardian of the Cell"

- Responsible for Repairing DNA damage
- Activates DNA Repair proteins
- Pauses the Cell Cycle (prevents replication of damage DNA)
- Initiates *apoptosis* (cell death) in the case where damage can't be repaired.
- Large scale feeback loop with NF-κB.

p53 DNA Damage Repair

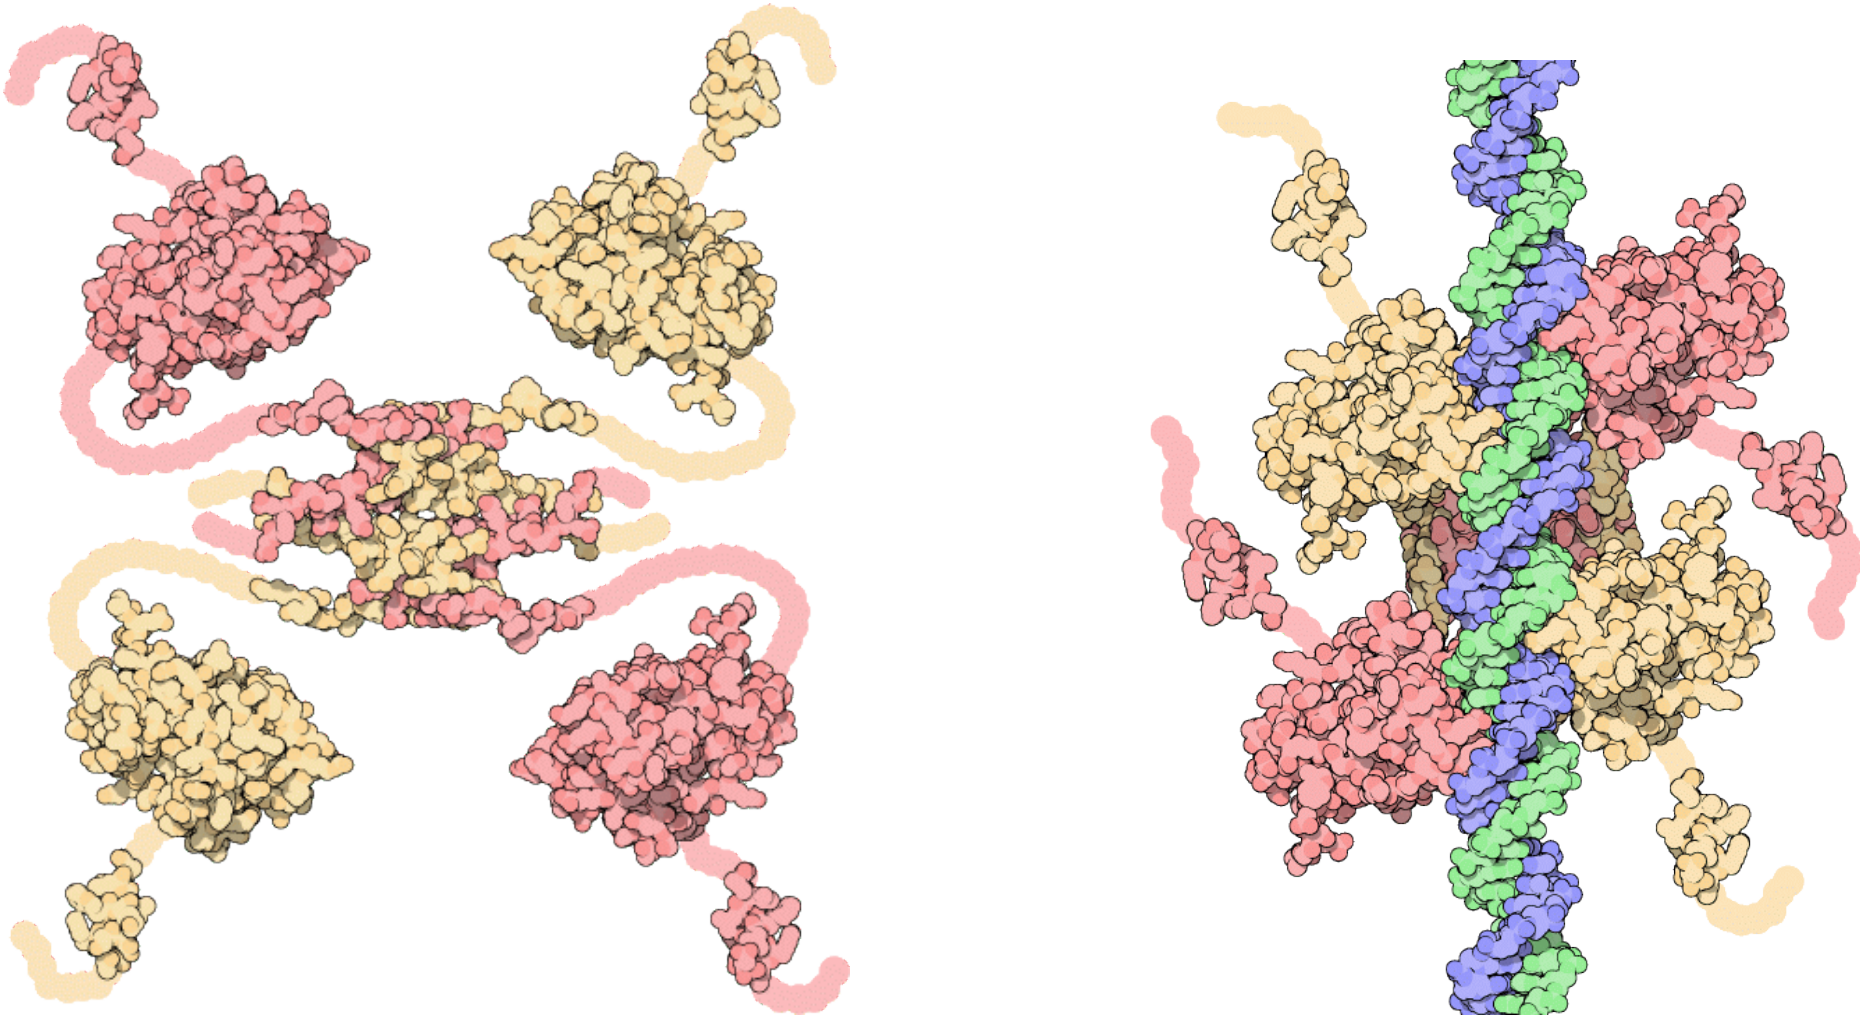

Figure: p53. *Left* unbound, *Right* bound to DNA. Images by David S. Goodsell from `http://www.rcsb.org/` (see the "Molecule of the Month" feature).

p53

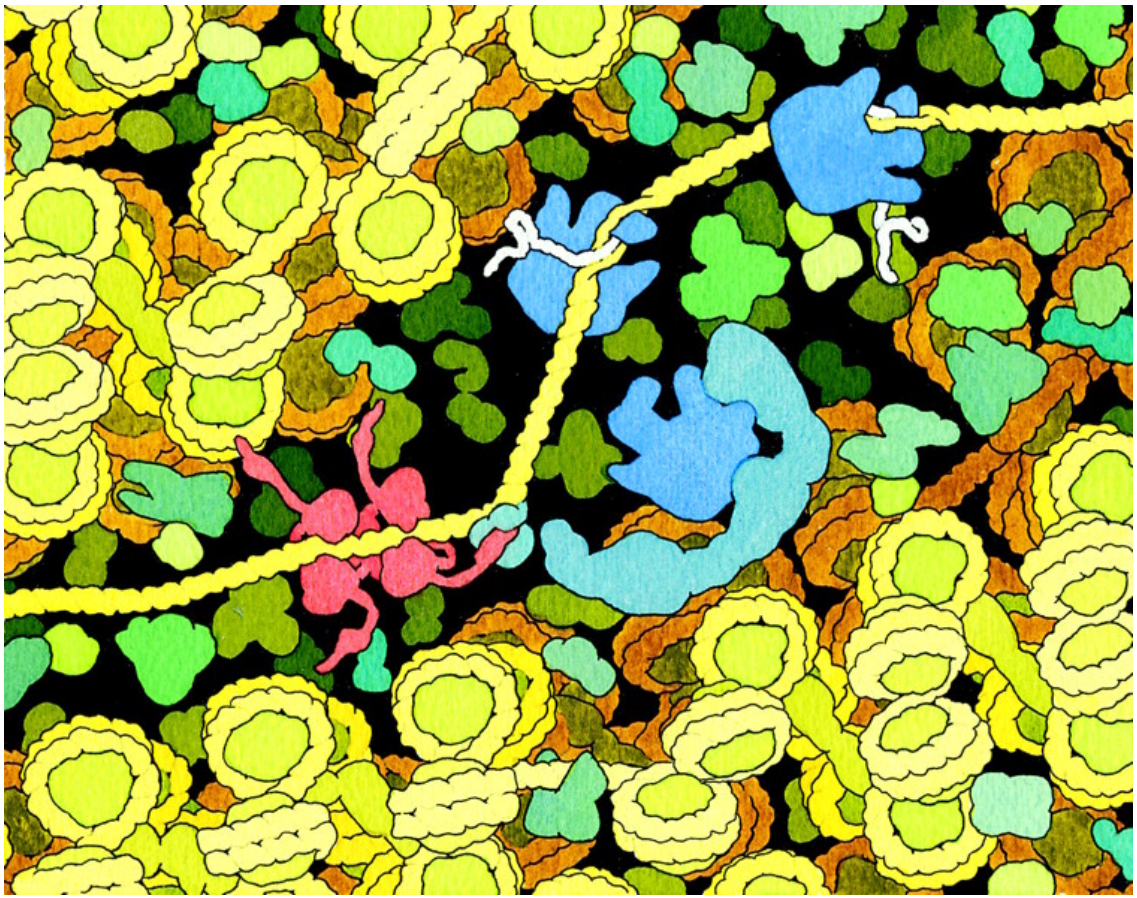

Figure: Repair of DNA damage by p53. Image fromGoodsell (1999).

Some p53 Targets

DDB2 DNA Damage Specific DNA Binding Protein 2. (also governed by C/ EBP-beta, E2F1, E2F3,...).

p21 Cycline-dependent kinase inhibitor 1A (CDKN1A). A regulator of cell cycle progression. (also goverened by SREBP-1a, Sp1, Sp3,...).

hPA26/SESN1 sestrin 1 Cell Cycle arrest.

BIK BCL2-interacting killer. Induces cell death (apoptosis)

TNFRSF10b tumor necrosis factor receptor superfamily, member 10b. A transducer of apoptosis signals.

Modelling Assumption

- Assume p53 affects targets as a single input module network motif (SIM).

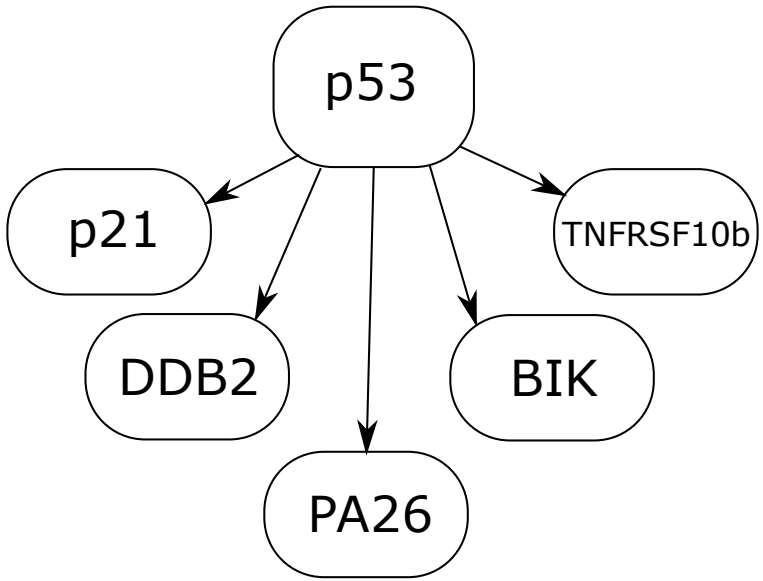

Figure: p53 SIM network motif as modelled by Barenco et al. 2006.

Response of p53 to Ionizing Radiation

- Experiment by Barenco et al. 2006.
- Human leukemia cell line (MOLT4) containing functional p53 and harvested protein and RNA at regular intervals after irradiation.
- The time course was performed in triplicate, and mRNA concentrations measured using Affymetrix U133A microarrays.

Mathematical Model

- Reorder differential equations

$$\frac{\mathrm{d}x_j(t)}{\mathrm{d}t} + D_j x_j(t) = B_j + S_j f(t)$$

- We have observation of $x_j(t)$.
- An estimate of $\frac{\mathrm{d}x_j(t)}{\mathrm{d}t}$ is obtained through fitting polynomials.
- Jointly estimate $f(t)$ at observations of time points along with $\{B_j, D_j, S_j\}_{j=1}^{g}$.
- Use MCMC sampling or maximum likelihood for parameters.

Response of p53

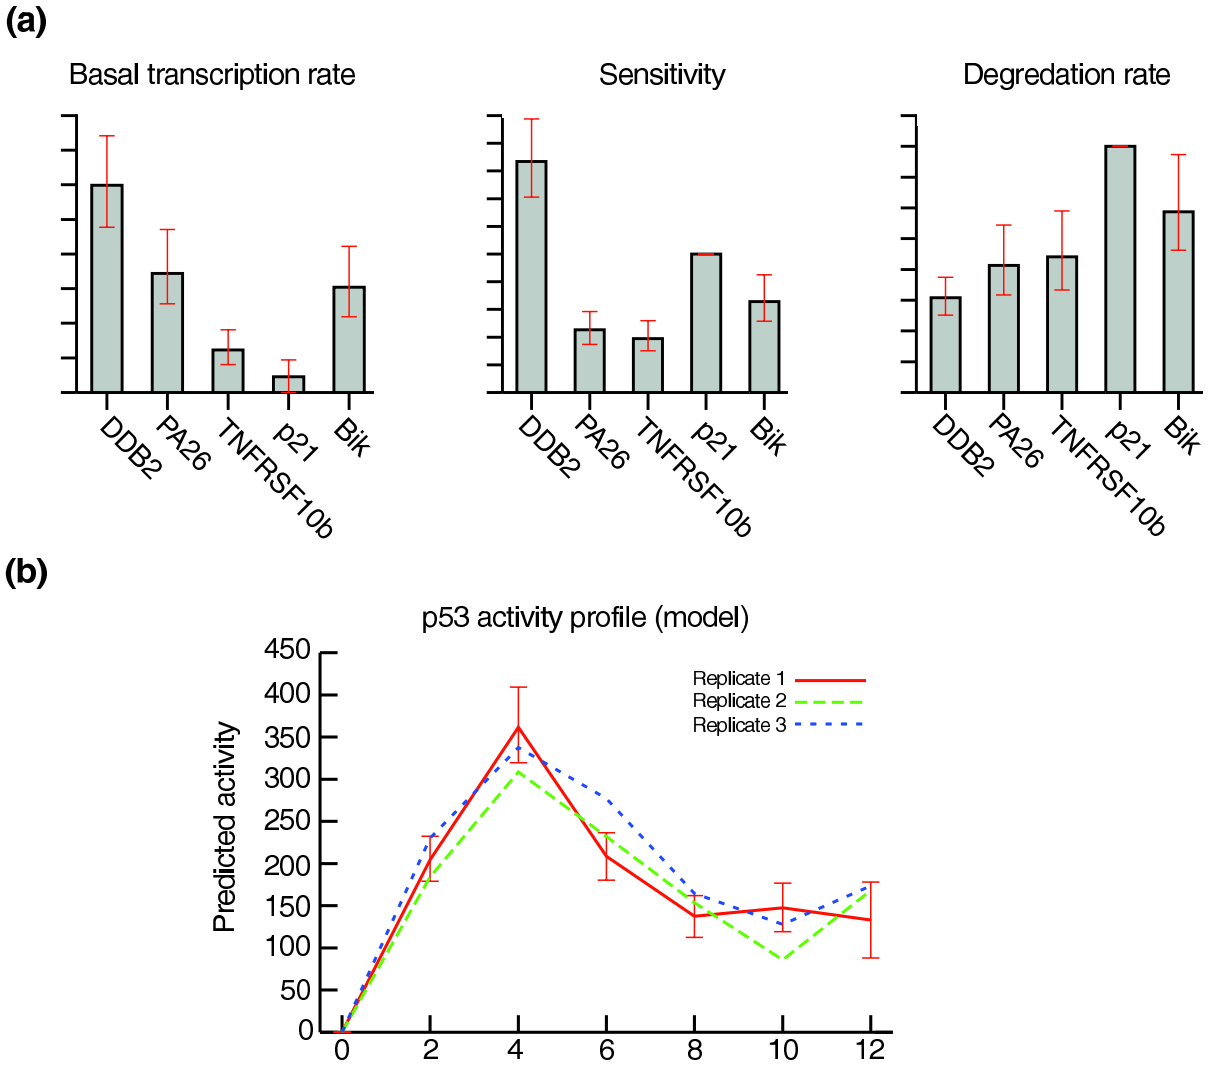

Figure: Results from Barenco et al. (2006). Top is parameter estimates. Bottom is inferred profile.

Respose to p53 ...

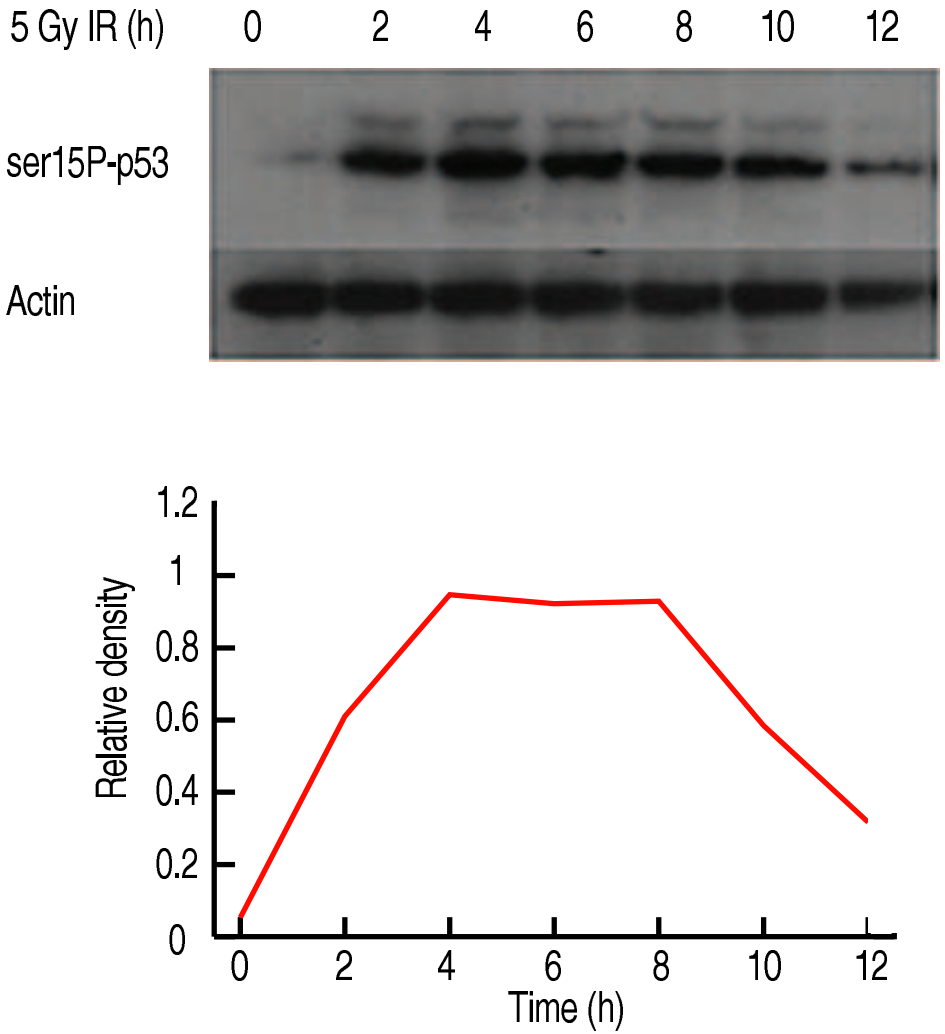

Figure: Results from Barenco et al. (2006). Activity profile of p53 was measured by Western blot to determine the levels of ser-15 phosphorylated p53 (ser15P-p53)

Models of non-linear regulation

- Non-linear Activation: Michaelis-Menten Kinetics

$$\frac{\mathrm{d}x_i(t)}{\mathrm{d}t} = B_i + \frac{S_i f(t)}{\gamma_i + f(t)} - D_i x_i(t)$$

used by Rogers and Girolami (2006)

- Non-linear Repression

$$\frac{\mathrm{d}x_i(t)}{\mathrm{d}t} = B_i + \frac{S_i}{\gamma_i + f(t)} - D_i x_i(t)$$

used by Khanin et al., 2006, PNAS 103

Models of non-linear regulation

- Non-linear Activation: Michaelis-Menten Kinetics

$$\frac{\mathrm{d}x_i(t)}{\mathrm{d}t} = B_i + \frac{S_i f(t)}{\gamma_i + f(t)} - D_i x_i(t)$$

used by Rogers and Girolami (2006)

- Non-linear Repression

$$\frac{\mathrm{d}x_i(t)}{\mathrm{d}t} = B_i + \frac{S_i}{\gamma_i + f(t)} - D_i x_i(t)$$

used by Khanin et al., 2006, PNAS 103

Models of non-linear regulation

- Non-linear Activation: Michaelis-Menten Kinetics

$$\frac{\mathrm{d}x_i(t)}{\mathrm{d}t} = B_i + \frac{S_i f(t)}{\gamma_i + f(t)} - D_i x_i(t)$$

  used by Rogers and Girolami (2006)

- Non-linear Repression

$$\frac{\mathrm{d}x_i(t)}{\mathrm{d}t} = B_i + \frac{S_i}{\gamma_i + f(t)} - D_i x_i(t)$$

  used by Khanin et al., 2006, PNAS 103

SOS Response

- Post replication DNA system: allows DNA replication to bypass errors in the DNA.
- DNA damage may occur as a result of activity of antibiotics.
- LexA is bound to the genome preventing transcription of the SOS genes.
- RecA protein is stimulated by single stranded DNA, inactivates the LexA repessor.
- This allows several of the LexA targets to transcribe.
- The SOS pathway may be essential in antibiotic resistance Cirz et al. (2005).
- Aim is to target these proteins to produce drugs to increase efficacy of antibiotics Lee et al. (2005).

LexA Experimental Description

- Data from Courcelle et al. (2001)
- UV irradiation of *E. coli.* in both wild-type cells and lexA1 mutants, which are unable to induce genes under LexA control.
- Response measured with two color hybridization to cDNA arrays.

Their Model

Given measurements of gene expression at N time points $(t_0, t_1, \ldots, t_{N-1})$, the temporal profile of a gene i, $x_i(t)$, that solves the ODE in Eq. 1 can be approximated by

$$x_i(t) = x_i^0 e^{-\delta_i t} + \frac{B_i}{D_i} + S_i e^{-\delta_i t} \frac{1}{D_i} \sum_{j=0}^{N-2} \left(e^{D_i t_j + 1} - e^{D_i t_j}\right) \frac{1}{\gamma_i + \bar{f}_j}$$

where $\bar{f}_j = \frac{(f(t_j) + f(t_j+1))}{2}$ on each subinterval $(t_j, t_j + 1), j = 0, \ldots, N-2$. This is under the simplifying assumption that $f(t)$ is a piece-wise constant function on each subinterval $(t_j, t_j + 1)$. **One can come up with linear (or higher order) $f(t)$ approximations on each subinterval. This will introduce additional parameters, which will be impossible to infer with any certainty given limited amount of data.**

Khanin et al. (2006)

Their Results

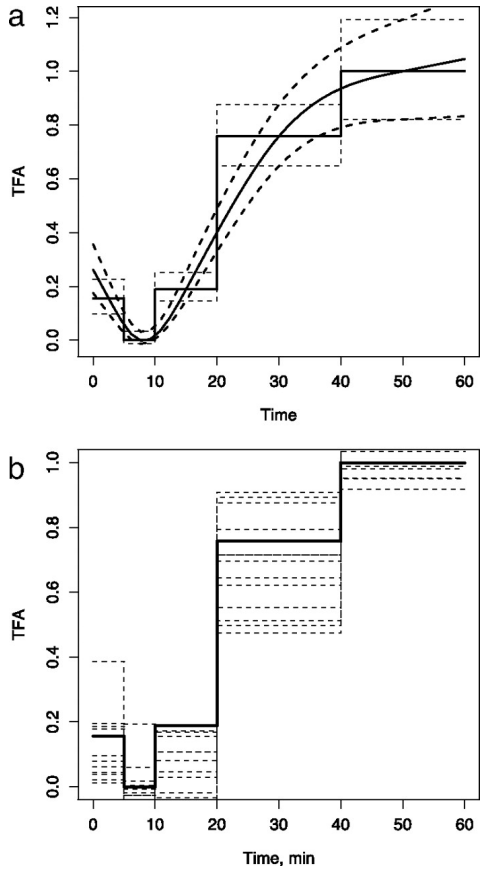

Figure: Fig. 2 from Khanin et al. (2006): Reconstructed activity level of master repressor LexA, following a UV dose of 40 J/m2.

Their Results

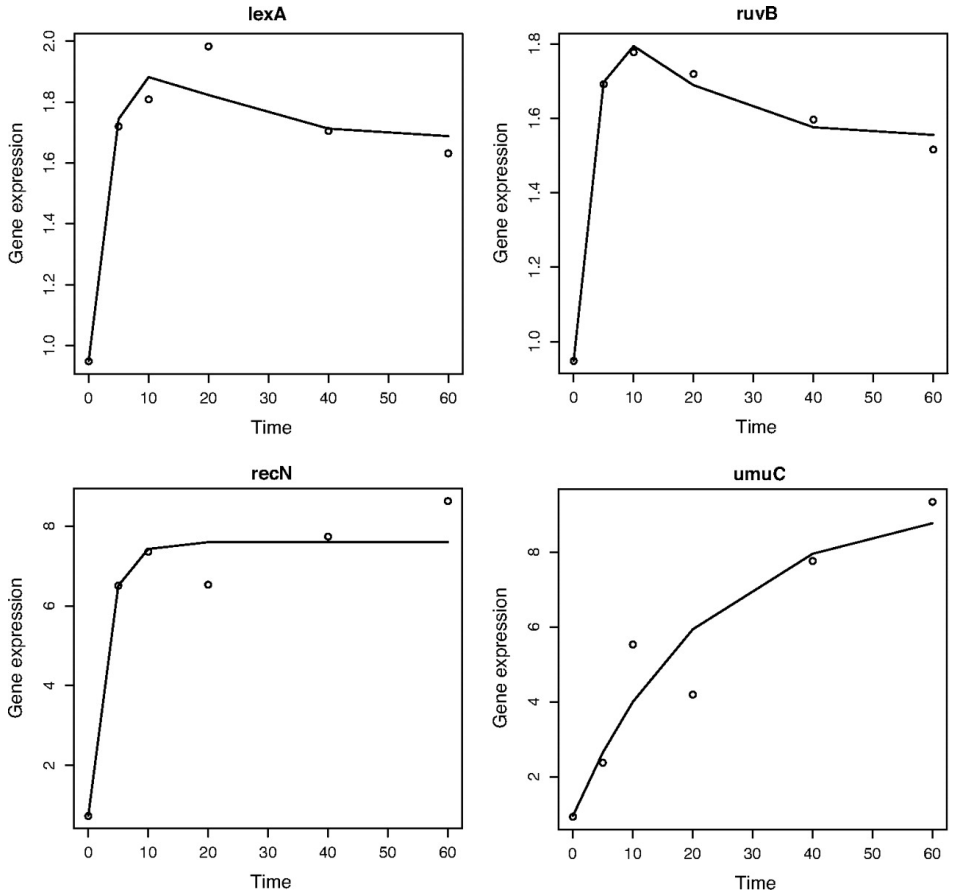

Figure: Fig. 3 from Khanin et al. (2006): Reconstructed profiles for four genes in the LexA SIM.

Actin and Ribosomes

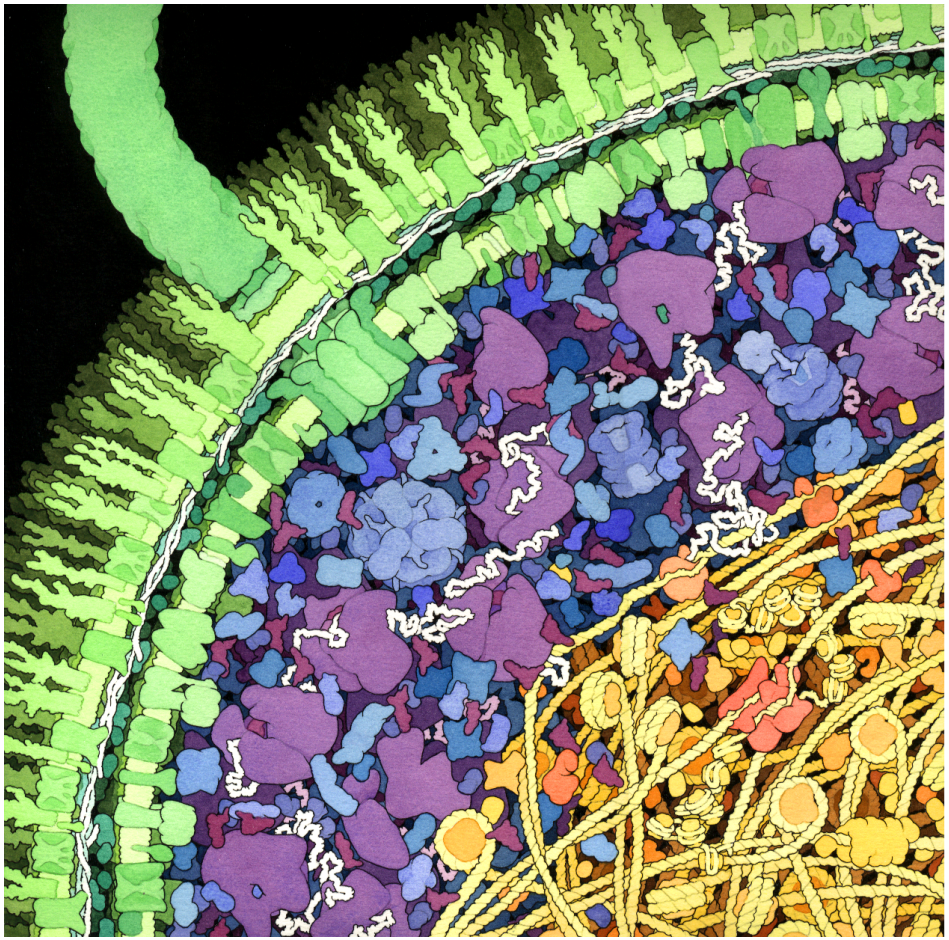

Figure: *E. coli* cell. Illustration courtesy of David S. Goodsell `http://mgl.scripps.edu/people/goodsell/illustration/public`. Confined structure leads to attempts to characterise diffusion in confined spaces, *e.g.* Schuss et al. (2007)

Outline

ERK Signalling Pathway

- Epidermal Growth Factor 40,000-100,000 EGFR per cell.
- Over expressed in tumours — some breast cancer cells 2×10^6 receptors per cell Herbst (2004).
- Over expression leads to an intense signal generation and activation of down stream signalling pathways.

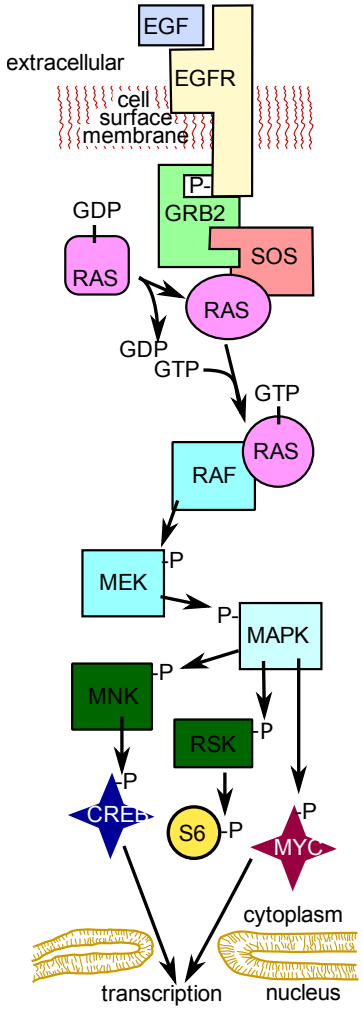

Figure: MAPK Pathway

Multiple Mechanistic Models

Vyshemirsky and Girolami (2008).

- Multiple mechanistic models describing a pathway.

Multiple Mechanistic Models

Vyshemirsky and Girolami (2008).

- Multiple mechanistic models describing a pathway.

Model 1

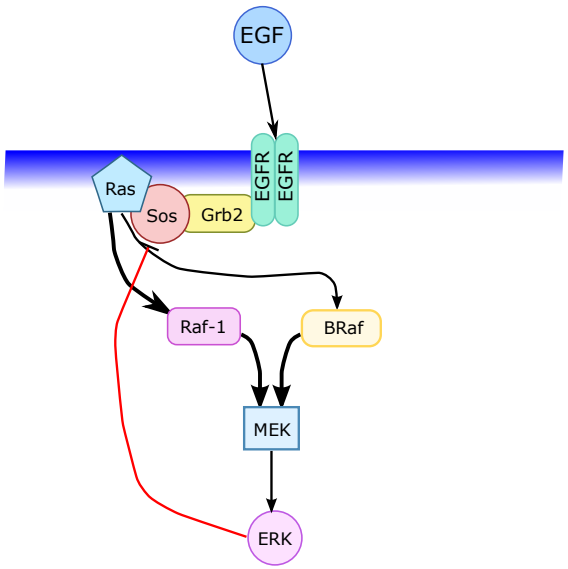

Multiple Mechanistic Models

Vyshemirsky and Girolami (2008).

- Multiple mechanistic models describing a pathway.

Model 1

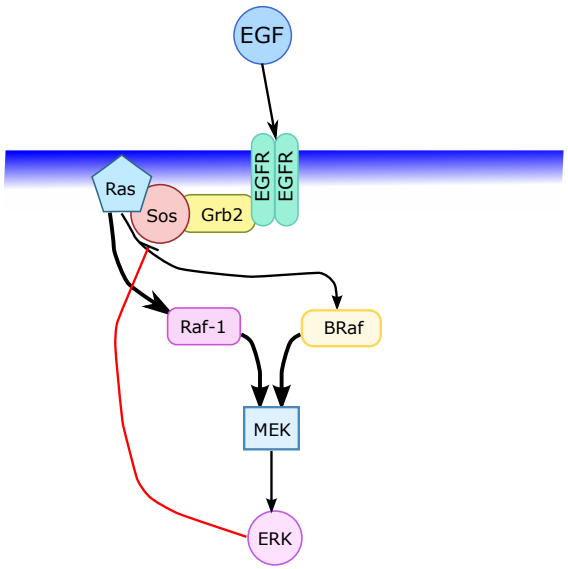

Model 2

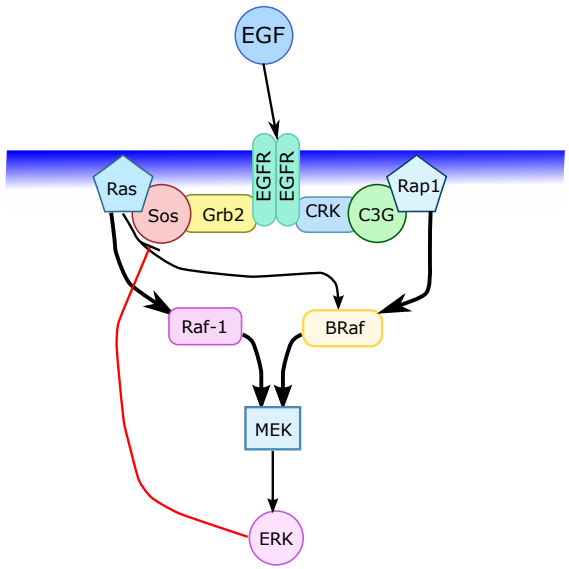

Multiple Mechanistic Models

Vyshemirsky and Girolami (2008).

- Multiple mechanistic models describing a pathway.

Model 1

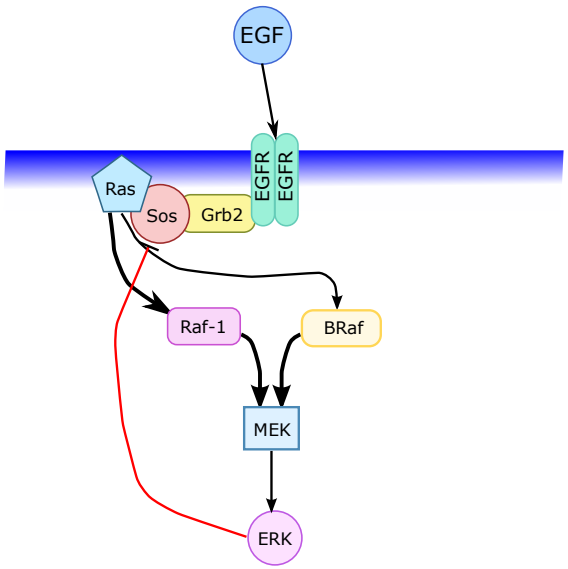

Model 2

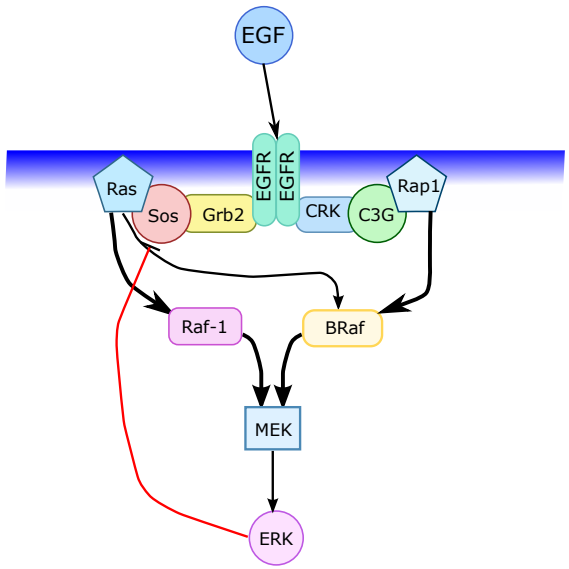

Model 3

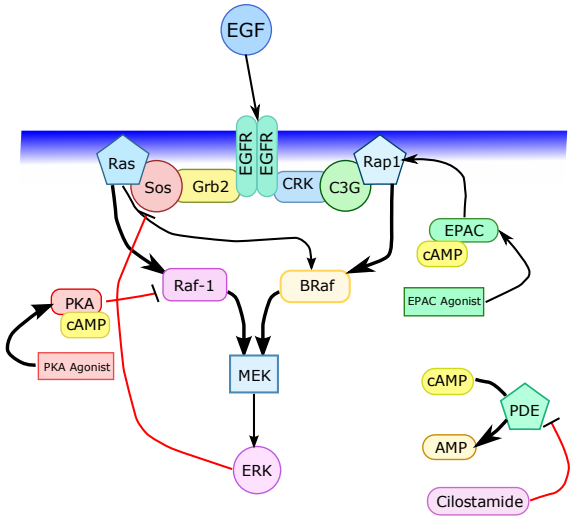

Differential Equations

Models are formally defined using systems of ordinary differential equations:

$$\frac{\mathrm{d}\,[\mathrm{EGF}]}{\mathrm{d}t} = -k_1\,[\mathrm{EGF}]\,[\mathrm{EGFR}]$$

$$\frac{\mathrm{d}\,[\mathrm{Rap1}_a]}{\mathrm{d}t} = \frac{K_{cat12}\,[\mathrm{Rap1}_i]}{K_{m12} + [\mathrm{Rap1}_i]}\,[\mathrm{EPAC}] - \frac{V_{13}\,[\mathrm{Rap1}_a]}{K_{13} + [\mathrm{Rap1}_a]}$$

$$\frac{\mathrm{d}\,[\mathrm{MEK}]}{\mathrm{d}t} = -\frac{K_{cat21}\,[\mathrm{MEK}]\,[\mathrm{Raf}] - 1}{K_{m21} + [\mathrm{MEK}]} - \frac{K_{cat22}\,[\mathrm{MEK}]}{K_{m22} + [\mathrm{MEK}]}\,[\mathrm{BRaf}]$$

Model 1	**Model 2**
50 kinetic parameters	55 kinetic parameters

Bayes Factors

- Which hypothesised structure is best supported by the data?
- Use Bayes factors: $\frac{P(M_1|D)}{P(M_2|D)}$, ratio of model marginal likelihoods.
- Difficulty is computing $P(M_1|D)$.
- Turn to the *thermodynamic integral* for results.

Thermodynamic Integral

Gelman and Meng (1998)

$$p\left(\boldsymbol{\theta}|\mathbf{x}, M, \alpha\right) = \frac{p\left(\mathbf{x}|\theta, M\right)^{\alpha} p\left(\theta|M\right)}{Z_{\alpha}}$$

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} \log Z_{\alpha} = \frac{1}{Z_{\alpha}} \frac{\mathrm{d}}{\mathrm{d}\, T} Z_{\alpha} = \left\langle \log p\left(\mathbf{x}|\boldsymbol{\theta}\right) \right\rangle_{p(\boldsymbol{\theta}|\mathbf{x}, M, \alpha)}$$

giving

$$\log p\left(\mathbf{x}|M\right) = \int_{0}^{1} \left\langle \log p\left(\mathbf{x}|\boldsymbol{\theta}\right) \right\rangle_{p(\boldsymbol{\theta}|\mathbf{x}, M, \alpha)} \mathrm{d}\alpha$$

Need samples from different temperatures.

Posterior for Different α

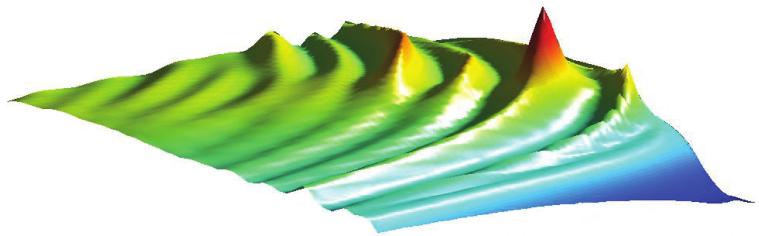

Figure: Annnealing of likelihood. Top is prior bottom is posterior (here $\alpha = 1$)

Posterior for Different α

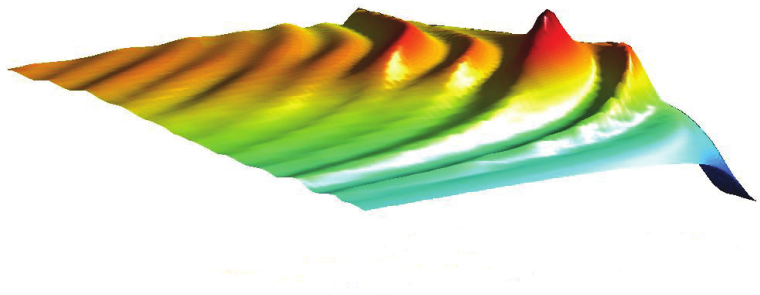

Figure: Annnealing of likelihood. Top is prior bottom is posterior (here $\alpha = 0.55$)

Posterior for Different α

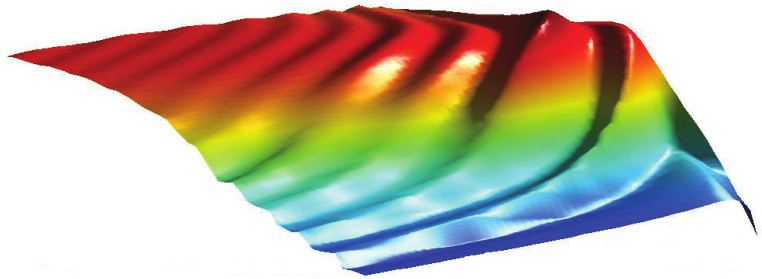

Figure: Annnealing of likelihood. Top is prior bottom is posterior (here $\alpha = 0.28$)

Posterior for Different α

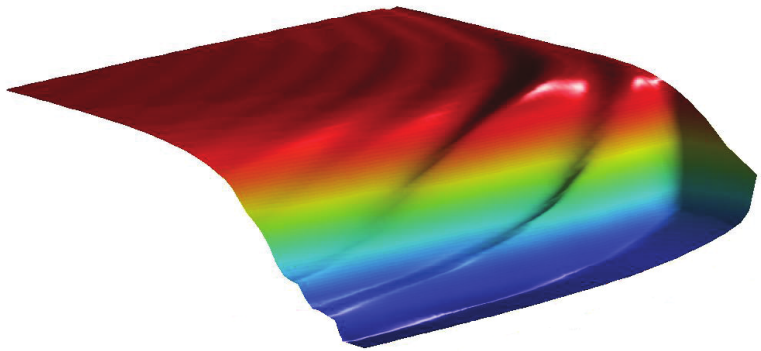

Figure: Annnealing of likelihood. Top is prior bottom is posterior (here $\alpha = 0.13$)

Posterior for Different α

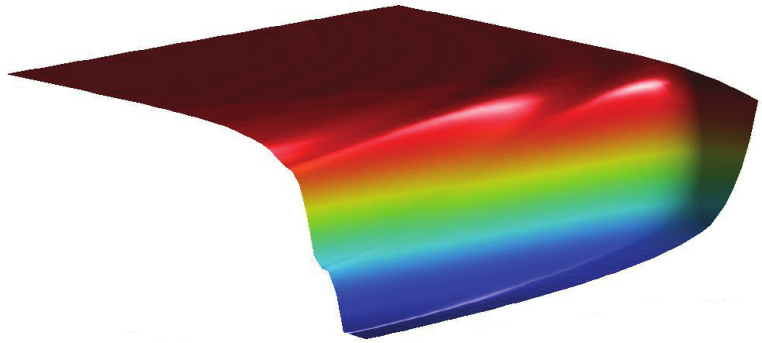

Figure: Annnealing of likelihood. Top is prior bottom is posterior (here $\alpha = 0.05$)

Posterior for Different α

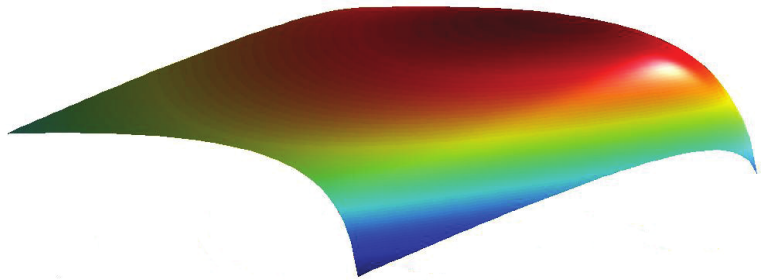

Figure: Annnealing of likelihood. Top is prior bottom is posterior (here $\alpha = 0$)

Population Monte Carlo

- Further problems from highly multimodal posteriors — use population Monte Carlo methods.

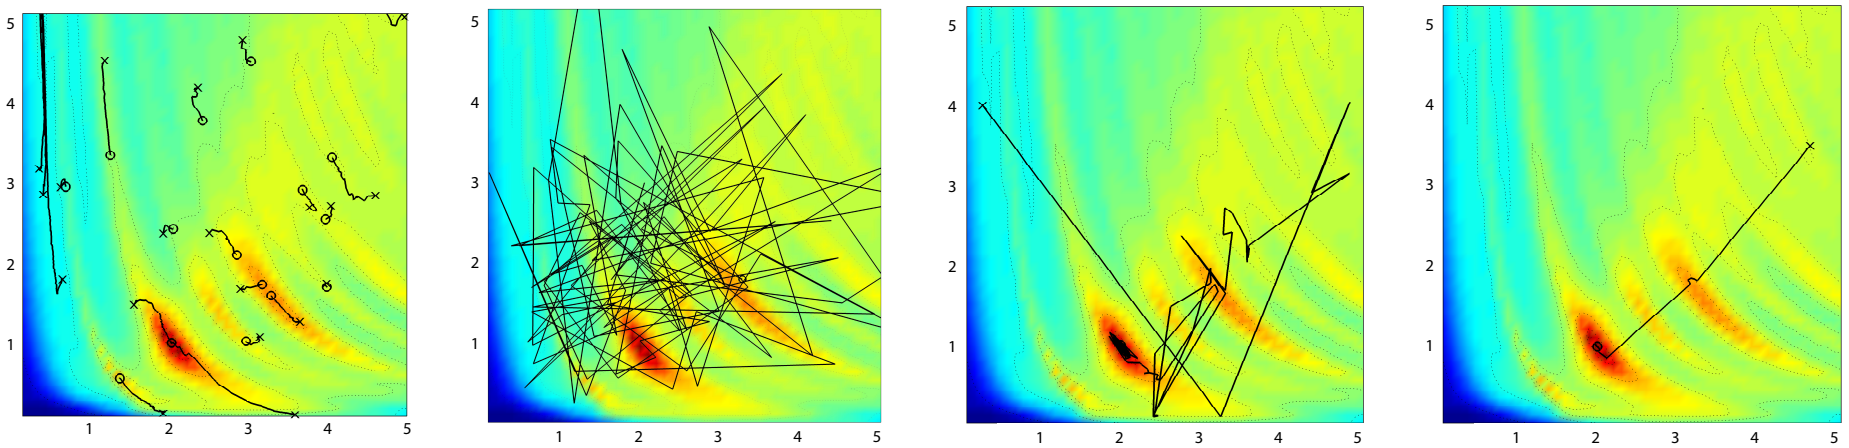

Figure: *Far Left*: standard Monte Carlo gets stuck in different modes. *Middle left*: exploration of space for low α. *Middle right:* intermediate α allows movement between modes. *Far left:* information is exchanged between samples to allow full exploration of posterior.

Bayes Factors for ERK signalling: Result

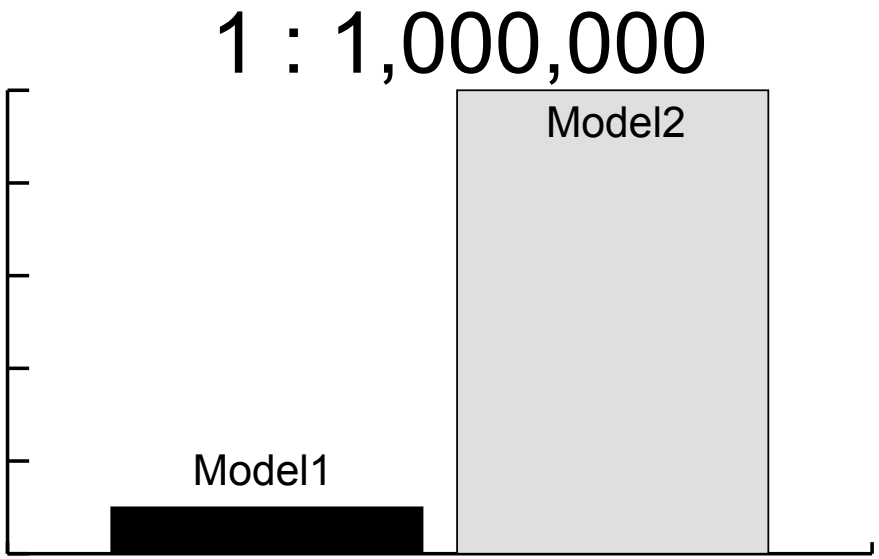

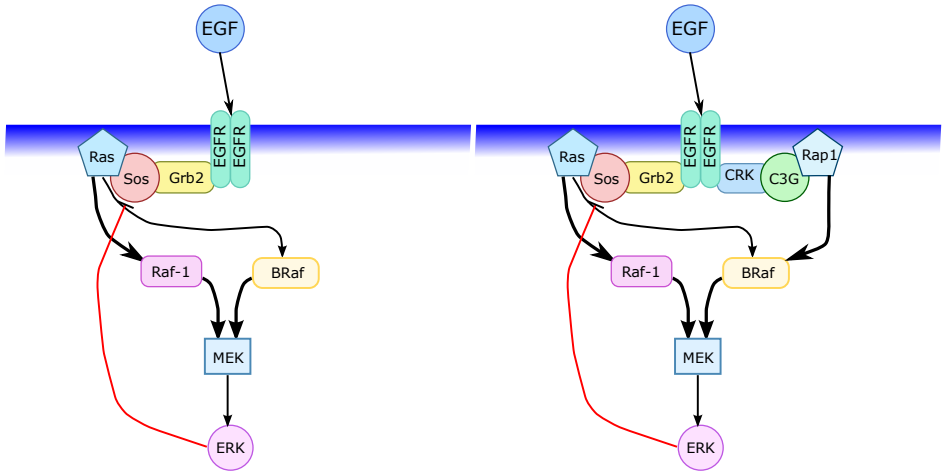

Hypothesis Implications

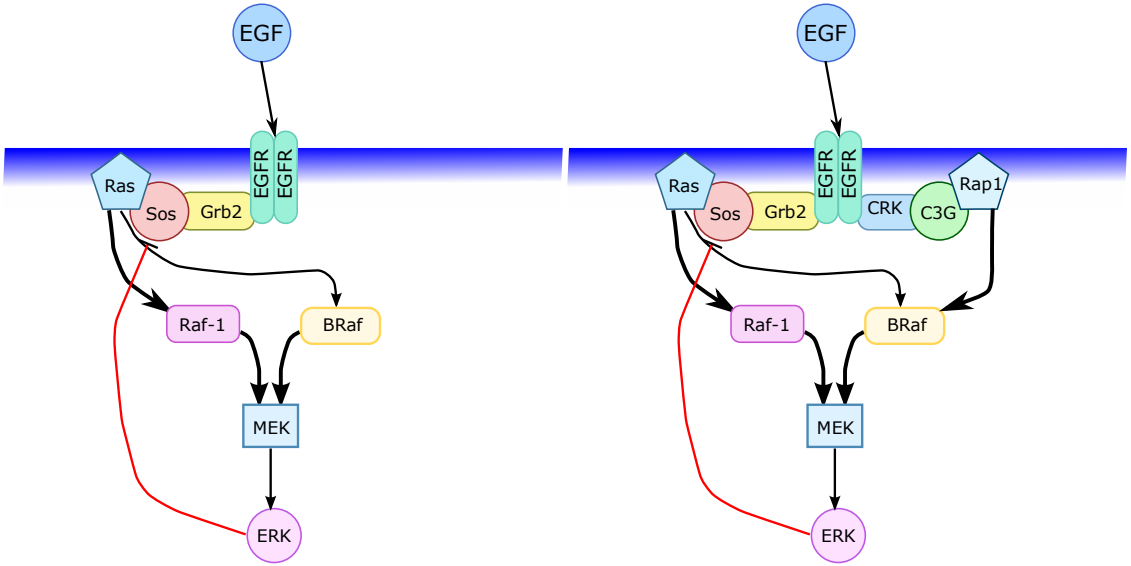

- Double branched model has much better support from the experimental evidence: leads to a robust system.
- BRaf was found to be more active than Raf-1. This is confirmed by a number of publications in biochemical journals.
- siRNA Knock-Down experiments have confirmed dual-branch hypothesis (Walter Kolch).

Hypothesis Implications

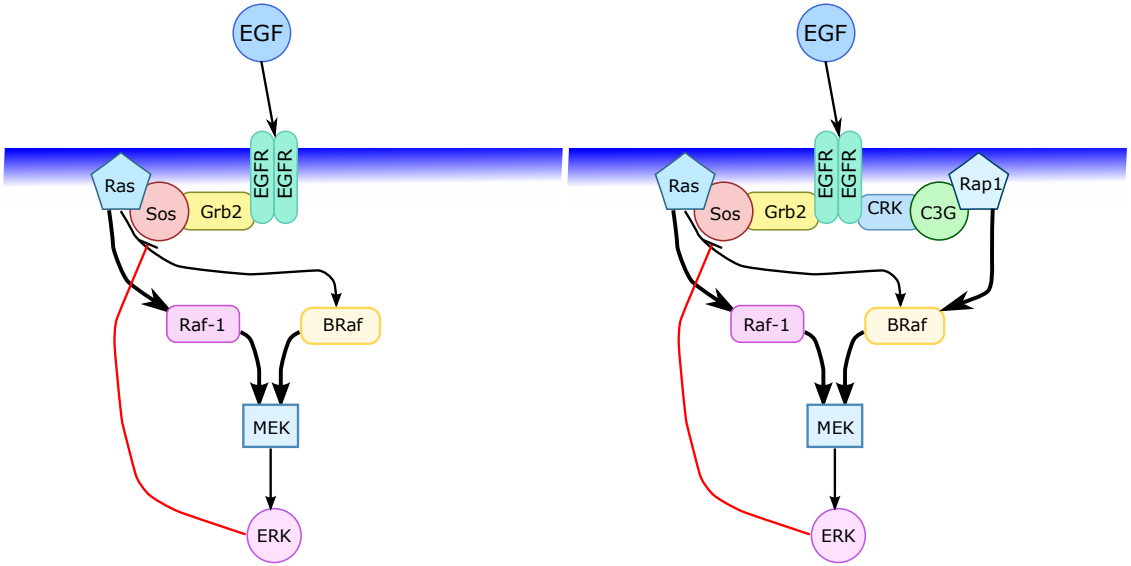

- Double branched model has much better support from the experimental evidence: leads to a robust system.
- BRaf was found to be more active than Raf-1. This is confirmed by a number of publications in biochemical journals.
- siRNA Knock-Down experiments have confirmed dual-branch hypothesis (Walter Kolch).

Hypothesis Implications

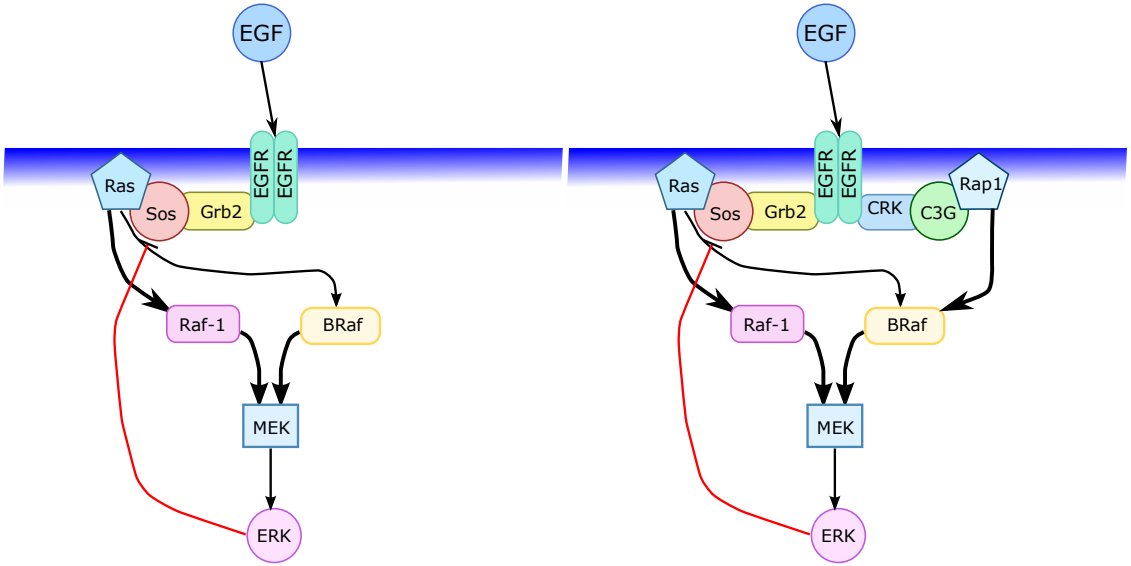

- Double branched model has much better support from the experimental evidence: leads to a robust system.
- BRaf was found to be more active than Raf-1. This is confirmed by a number of publications in biochemical journals.
- siRNA Knock-Down experiments have confirmed dual-branch hypothesis (Walter Kolch).

Outline

Summary and Conclusions

- Systems biology presents us with models and data.
- Challenge for machine learning: introduce our inference techniques to this domain.
- Lots of work on methodological developments necessary still.
- **Next part**: an approach to dealing with differential equations with missing chemical species.
  - Gaussian processes allow integration of Bayesian probabilistic inference with differential equations.

References I

M. Barenco, D. Tomescu, D. Brewer, R. Callard, J. Stark, and M. Hubank. Ranked prediction of p53 targets using hidden variable dynamic modeling. *Genome Biology*, 7(3):R25, 2006. [PDF].

R. T. Cirz, J. K. Chin, D. R. Andes, V. de Crécy-Lagard, W. A. Craig, and F. E. Romesberg. Inhibition of mutation and combating the evolution of antibiotic resistance. *PLoS Biology*, 3(6), 2005.

J. Courcelle, A. Khodursky, B. Peter, P. O. Brown, , and P. C. Hanawalt. Comparative gene expression profiles following UV exposure in wild-type and SOS-deficient *Escherichia coli*. *Genetics*, 158:41–64, 2001.

M. B. Elowitz and S. Leibler. Synthetic oscillatory network of transcriptional regulators. *Nature*, 403(6767):335–338, 2000. [DOI].

R. P. Feynman. There's plenty of room at the bottom: An invitation to enter a new field of physics. Talk at Annual meeting of the American Physical Society, 1959. Available from `http://www.zyvex.com/nanotech/feynman.html`.

A. Garfinkel. Reductionism. In R. Boyd, P. Gasper, and J. D. Trout, editors, *The Philosophy of Science*, pages 443–459. MIT Press, 1991. [Google Books] .

A. Gelman and X.-L. Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163–185, 1998.

M. A. Gibson and J. Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *J. Phys Chem. A*, 104:1876–1889, 2000.

D. S. Goodsell. The molecular perspective: p53 tumor suppressor. *The Oncologist, Vol. 4, No. 2, 138-139, April 1999*, 4(2): 138–139, 1999.

R. S. Herbst. Review of epidermal growth factor receptor biology. *International Journal of Radiation Oncology*, 59(2):S21–S26, 2004. [DOI].

S. Hoops, S. Sahle, C. Lee, J. Pahle, N. Simus, M. Singhal, L. Xu, P. Mendes, and U. Kummer. COPASI: a COmplex PAthway SImulator. *Bioinformatics*, 22(24):3067–3074, 2006.

R. Khanin, V. Viciotti, and E. Wit. Reconstructing repressor protein levels from expression of gene targets in *E. Coli*. *Proc. Natl. Acad. Sci. USA*, 103(49):18592–18596, 2006. [PDF]. [DOI].

References II

A. M. Lee, C. T. Ross, B.-B. Zeng, , and S. F. Singleton. A molecular target for suppression of the evolution of antibiotic resistance: Inhibition of the escherichia coli reca protein by n6-(1-naphthyl)-adp. *J. Med. Chem.*, 48(17), 2005.

S. Rogers and M. Girolami. Model based identification of transcription factor regulatory activity via Markov chain Monte Carlo. Presentation at MASAMB '06, 2006.

Z. Schuss, A. Singer, and D. Holcman. The narrow escape problem for diffusion in cellular microdomains. *Proc. Natl. Acad. Sci. USA*, 104(41):16098–16103, 2007. [DOI].

V. Vyshemirsky and M. A. Girolami. Bayesian ranking of biochemical system models. *Bioinformatics*, 24(6):833–839, 2008. [PDF]. [DOI].

D. J. Wilkinson. *Stochastic Modelling for Systems Biology*. Chapman and Hall/CRC, 2006. [Google Books] .