Optimization and Control

Performance Monitoring

Expected Value of Control

Limitation

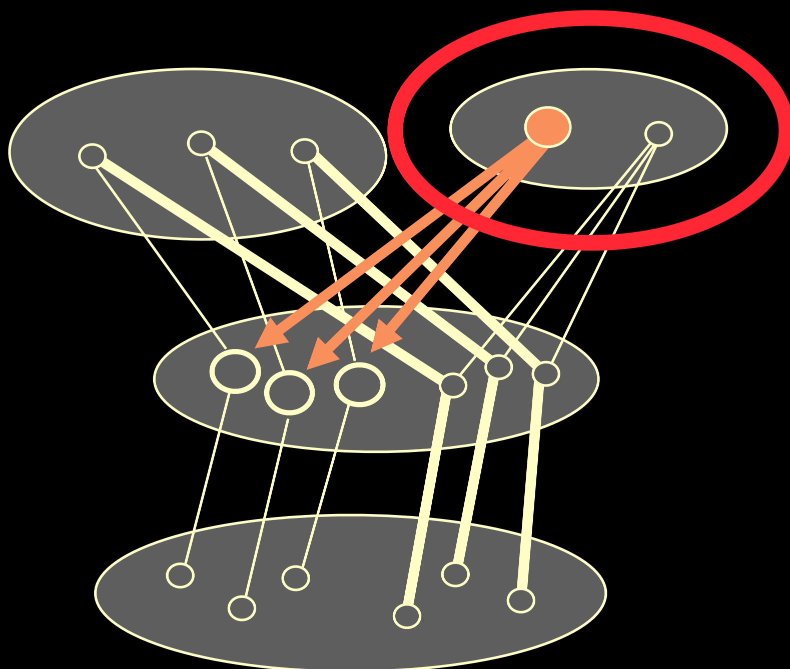

Limitation

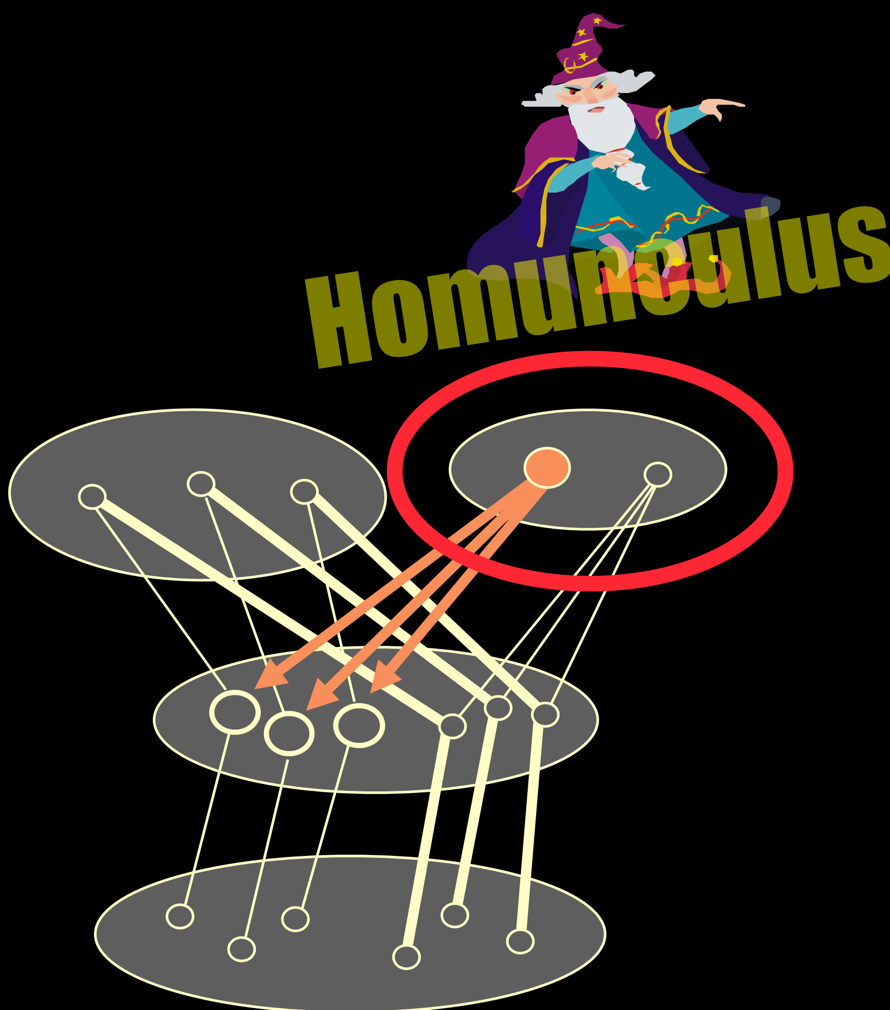

Challenges

- How are control representations maintained w/in PFC?

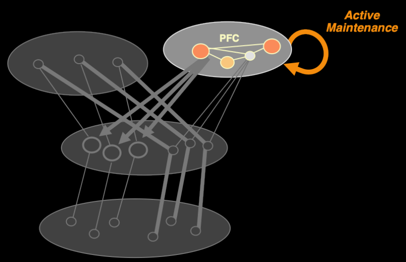

- How are control representations updated?
- How are adjustments made in the degree of control?
- How do representations develop, and what do they look like?

Challenges

- How are control representations maintained w/in PFC?

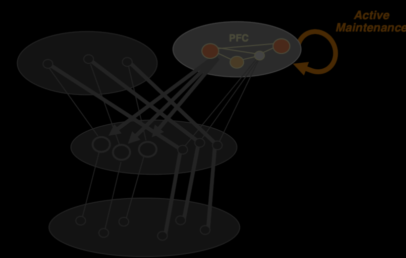

- How are control representations updated?
- How are adjustments made in the degree of control?
- How do representations develop, and what do they look like?

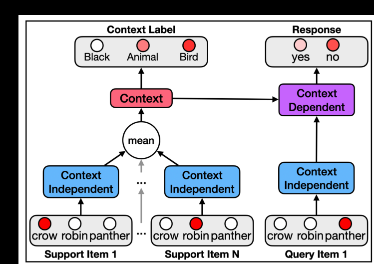

Challenges

- How are control representations maintained w/in PFC?

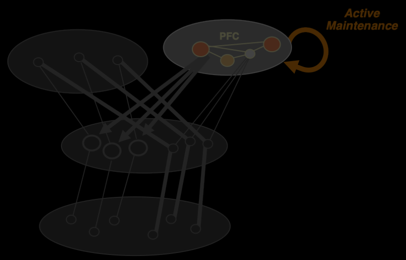

- How are control representations updated?

- How are adjustments made in the degree of control?

- How do representations develop, and what do they look like?

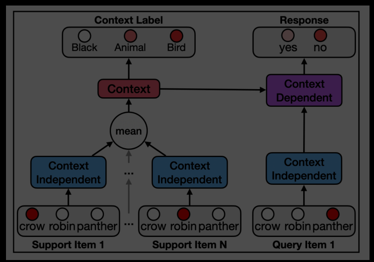

Control

Control

- Control as ***parmeter* *tuning:***

  DDM as example:

  - **Starting point:** ***expectations*** **(priors)**
  - **Drift rate: signal strength /** ***attention***
  - **Threshold:** ***speed-accuracy trade-off***

Control

- Control as *<u>parmeter</u> tuning:*

  DDM as example:

  - Starting point: *expectations* (priors)
  - Drift rate: signal strength / *attention*
  - Threshold: *speed-accuracy trade-off*

- Control as *<u>global gain</u> (NE) adjustment:*

  explore / exploit

  flexibility / stability

Control

- Control as *__parmeter__ tuning:*

  DDM as example:

  - Starting point: *expectations* (priors)
  - Drift rate: signal strength / *attention*
  - Threshold: *speed-accuracy trade-off*

- Control as *__global gain__ (NE) adjustment:*

  explore / exploit

  flexibility / stability

- Control as *__conflict__ management:*

  serialization of processing to avoid interference

Control

- Control as *__parmeter__ tuning:*

  DDM as example:

  - Starting point: *expectations* (priors)
  - Drift rate: signal strength / *attention*
  - Threshold: *speed-accuracy trade-off*

- Control as *__global gain__ (NE) adjustment:*

  explore / exploit

  flexibility / stability

- Control as *__conflict__ management:*

  serialization of processing to avoid interference

- Its all about *reward rate __optimization__*

Control

- **Control as *<u>parmeter</u> tuning:***

  DDM as example:

  - **Starting point:** ***expectations* (priors)**
  - **Drift rate: signal strength / *attention***
  - **Threshold: *speed-accuracy trade-off***

- **Control as *<u>global gain</u> (NE) adjustment:***

  explore / exploit

  flexibility / stability

- **Control as *<u>conflict</u> management:***

  serialization of processing to avoid interference

- **Its all about *reward rate <u>optimization</u>***

  requires *monitoring* for reward, i.e., *performance*

Control

- Control as *<u>parmeter</u> tuning:*

  DDM as example:

  - Starting point: *expectations* (priors)
  - Drift rate: signal strength / *attention*
  - Threshold: *speed-accuracy trade-off*

- Control as *<u>global gain</u> (NE) adjustment:*

  explore / exploit

  flexibility / stability

- Control as *<u>conflict</u> management:*

  serialization of processing to avoid interference

- Its all about *reward rate <u>optimization</u>*

  requires *monitoring* for reward, i.e., *performance*

  we've considered reward, but what about *conflict?*

Less Need for Control…

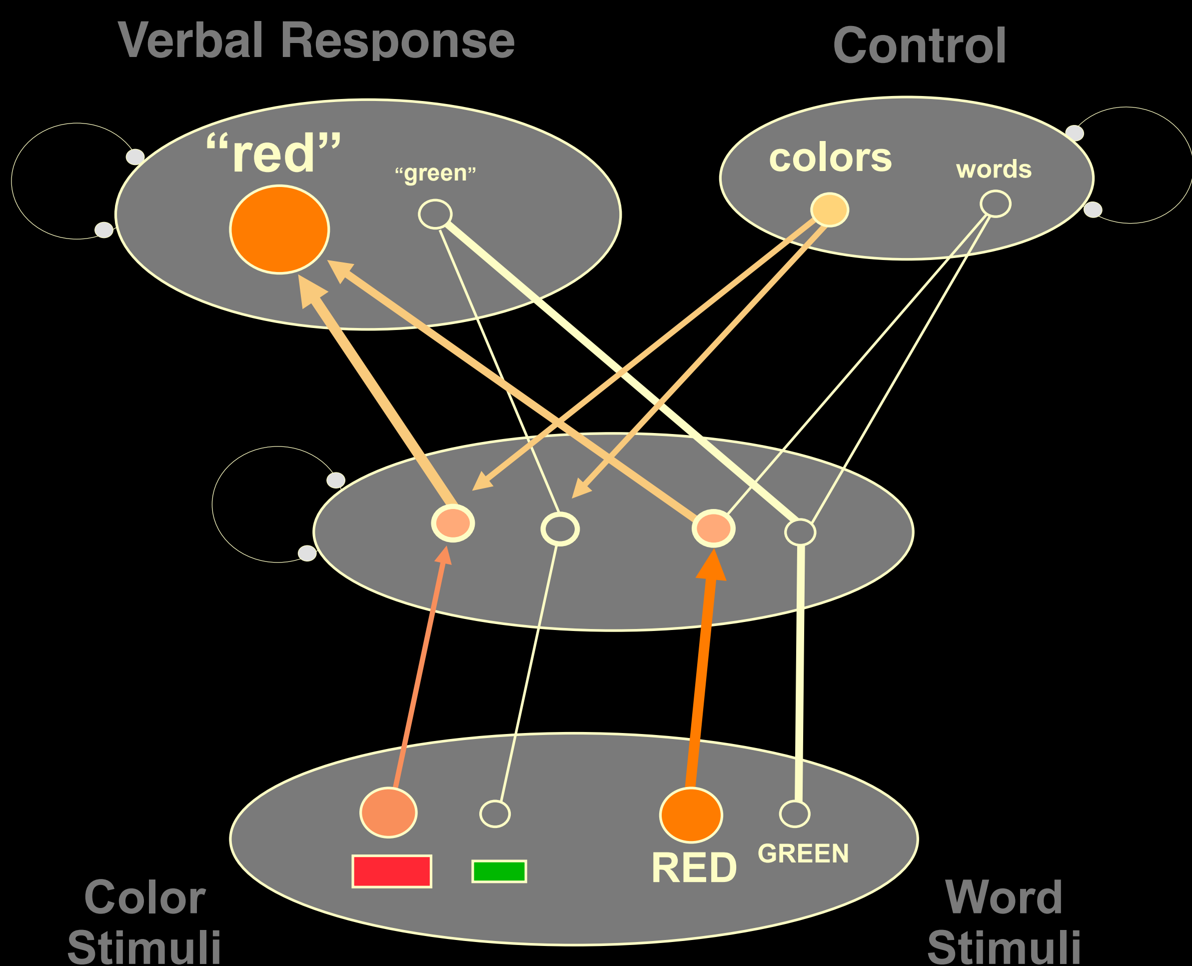

Stimulus: RED

Less Need for Control…

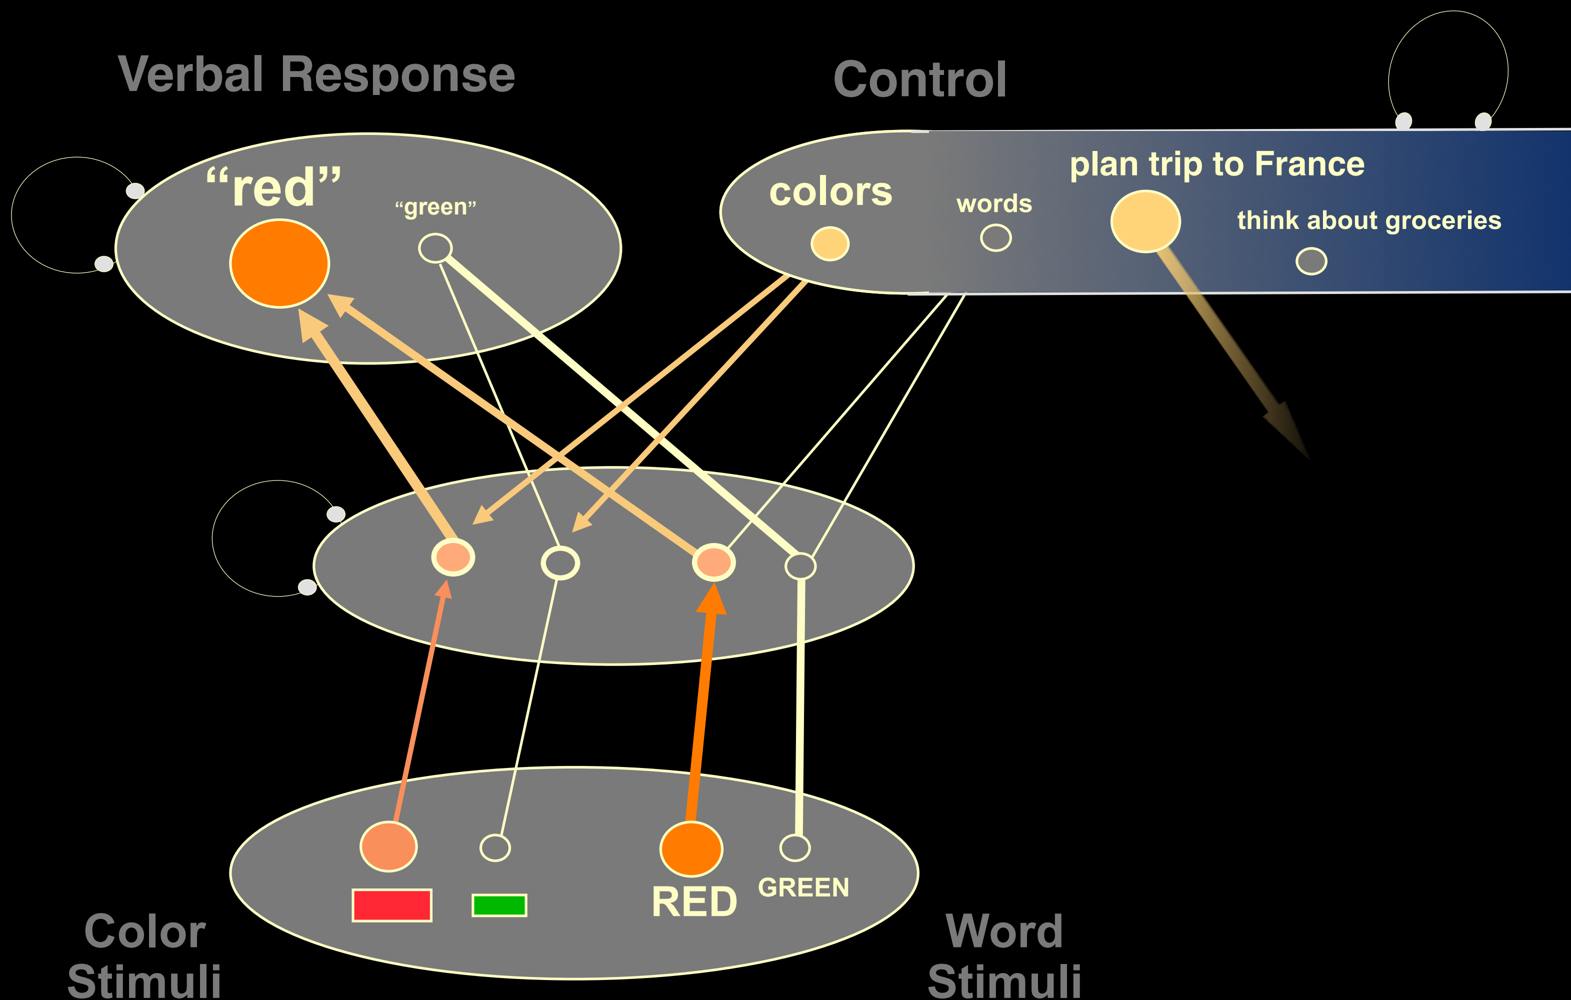

Stimulus: RED

More Need for Control

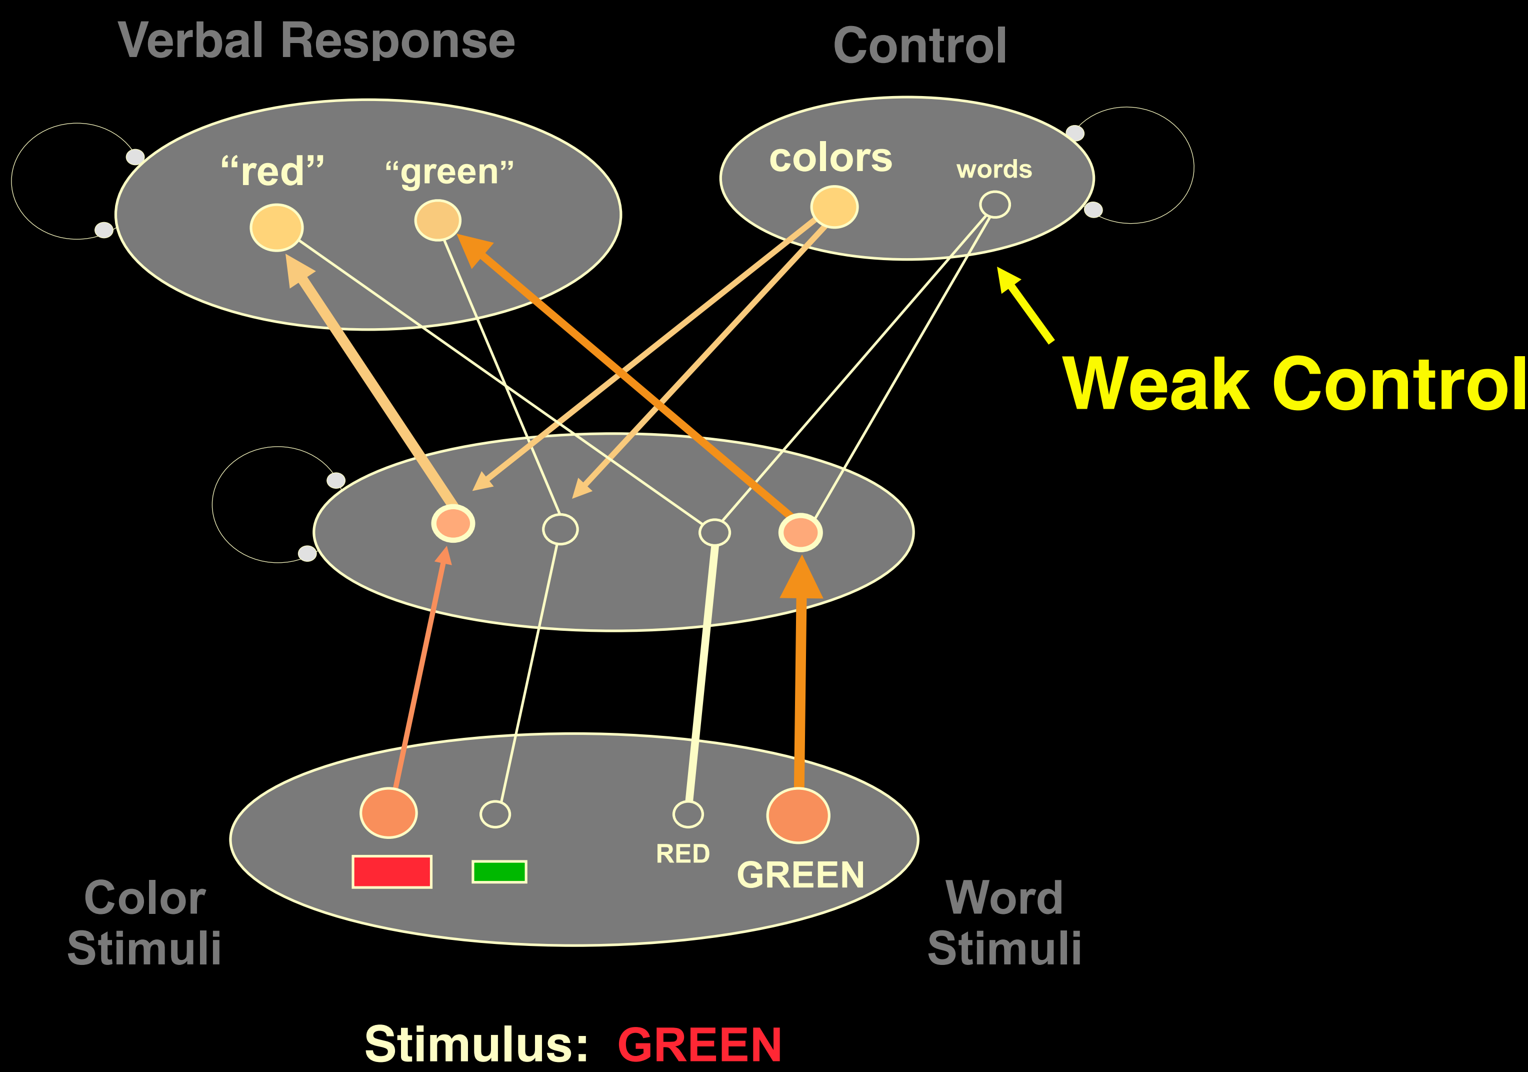

More Need for Control

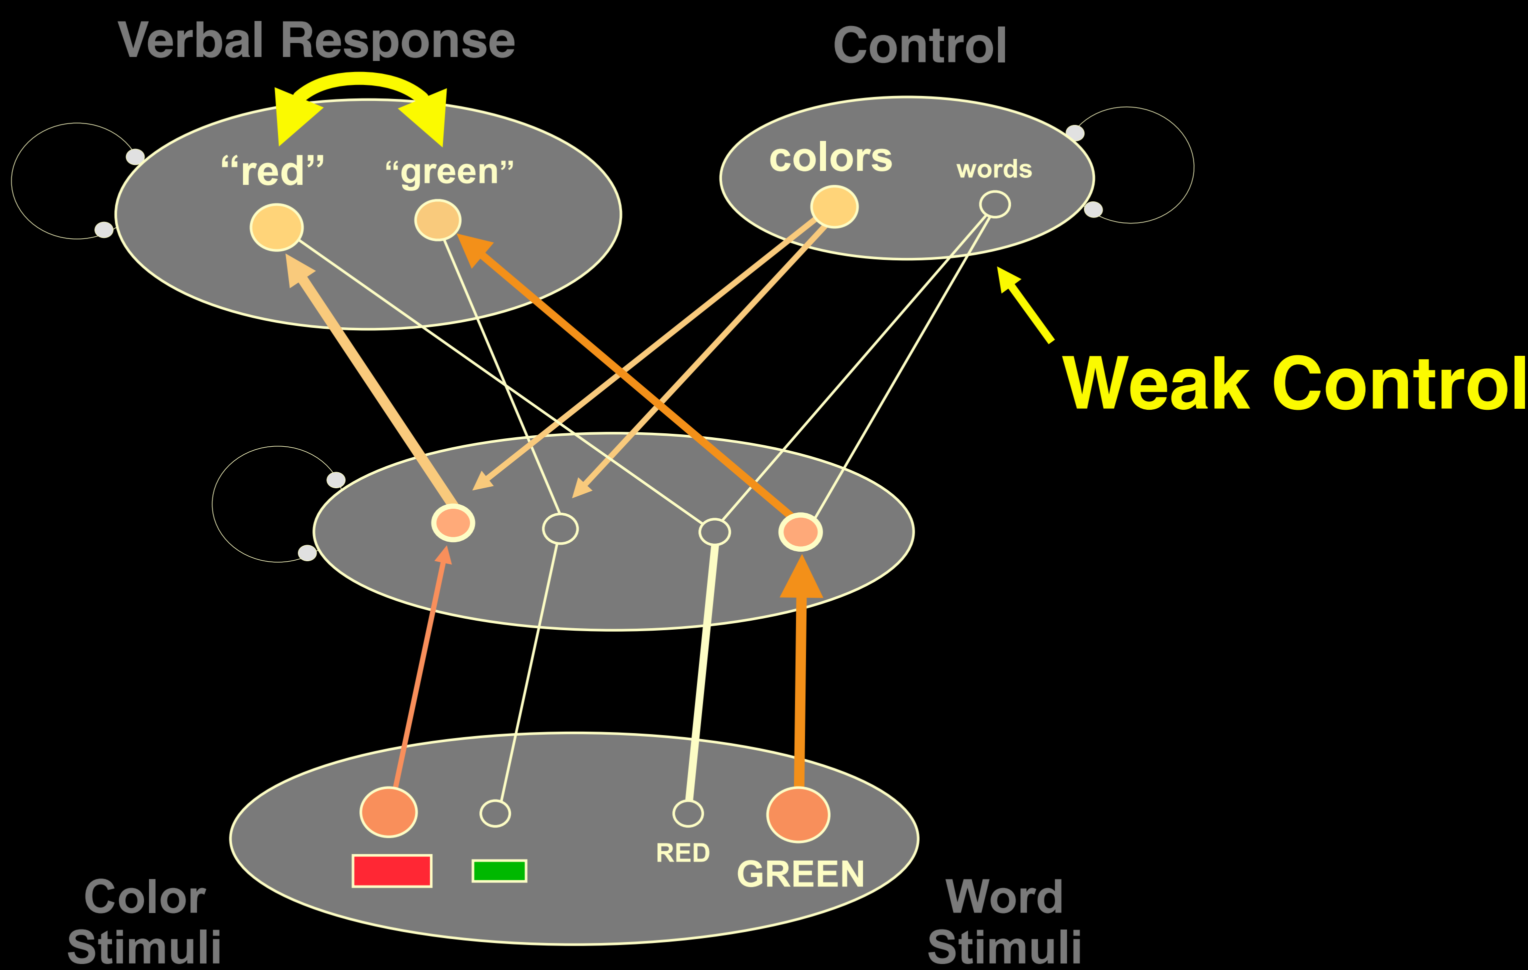

Stimulus: GREEN

More Need for Control

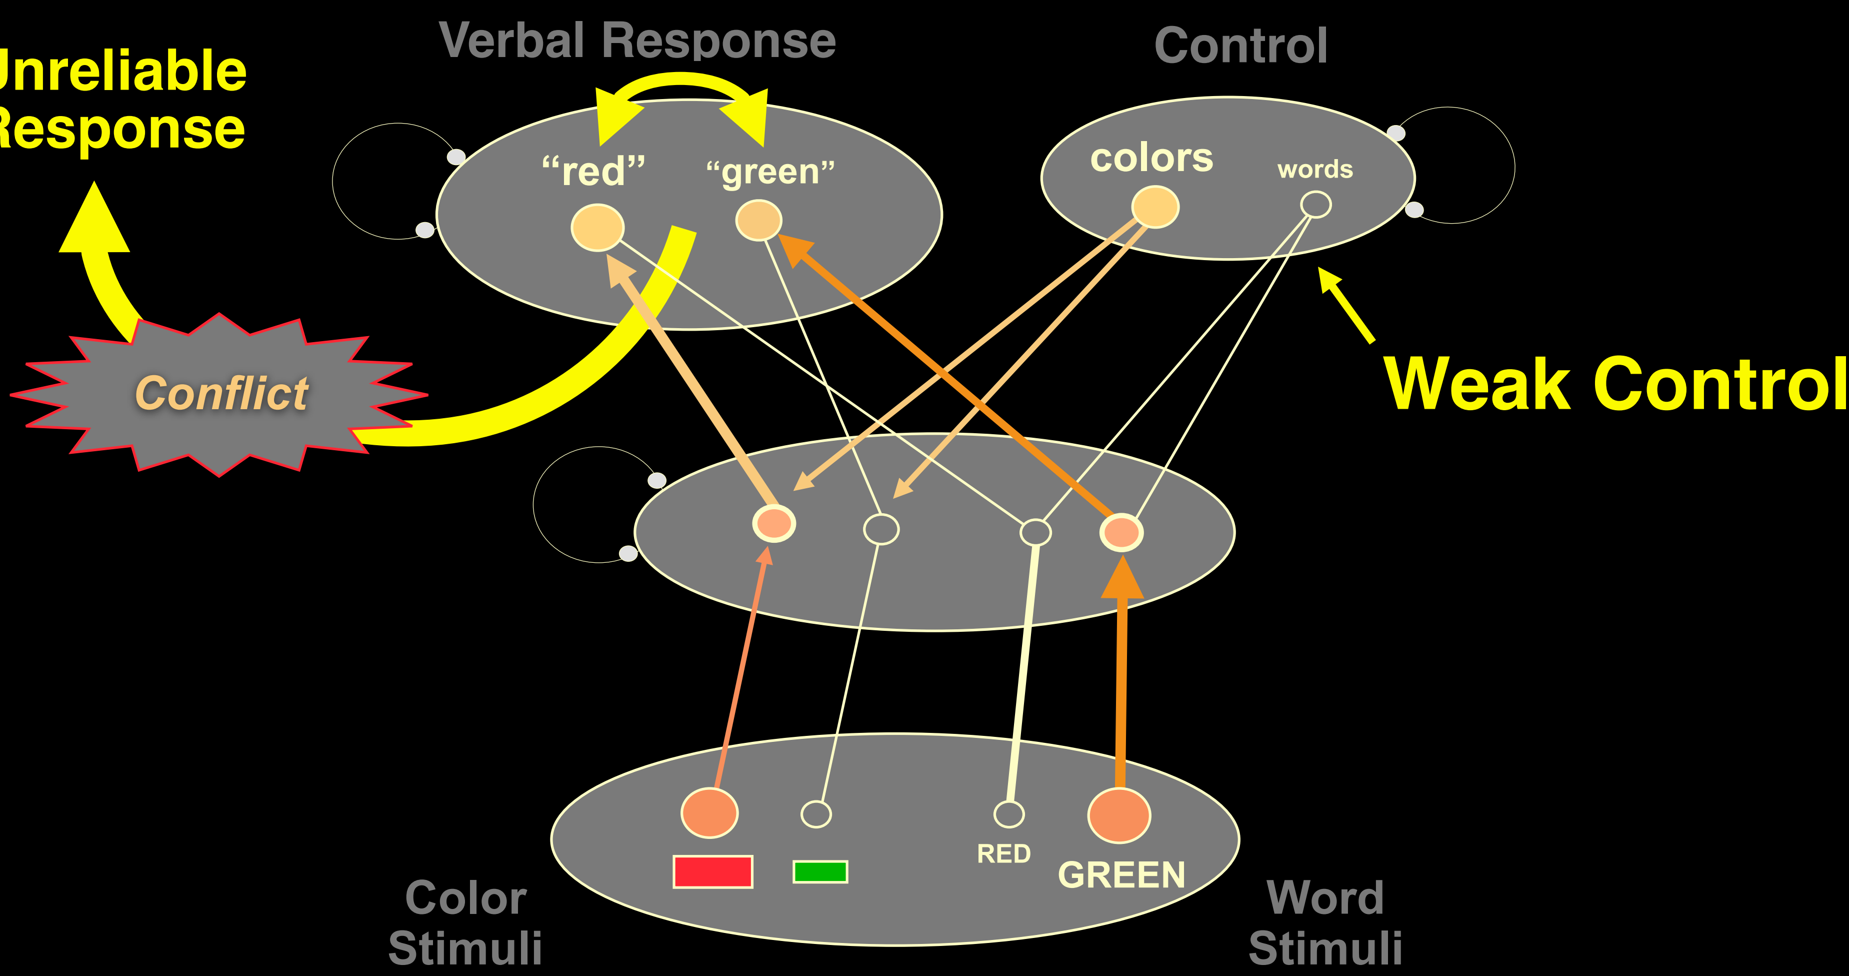

Stimulus: GREEN

Conflict Monitoring and Modulation of Control

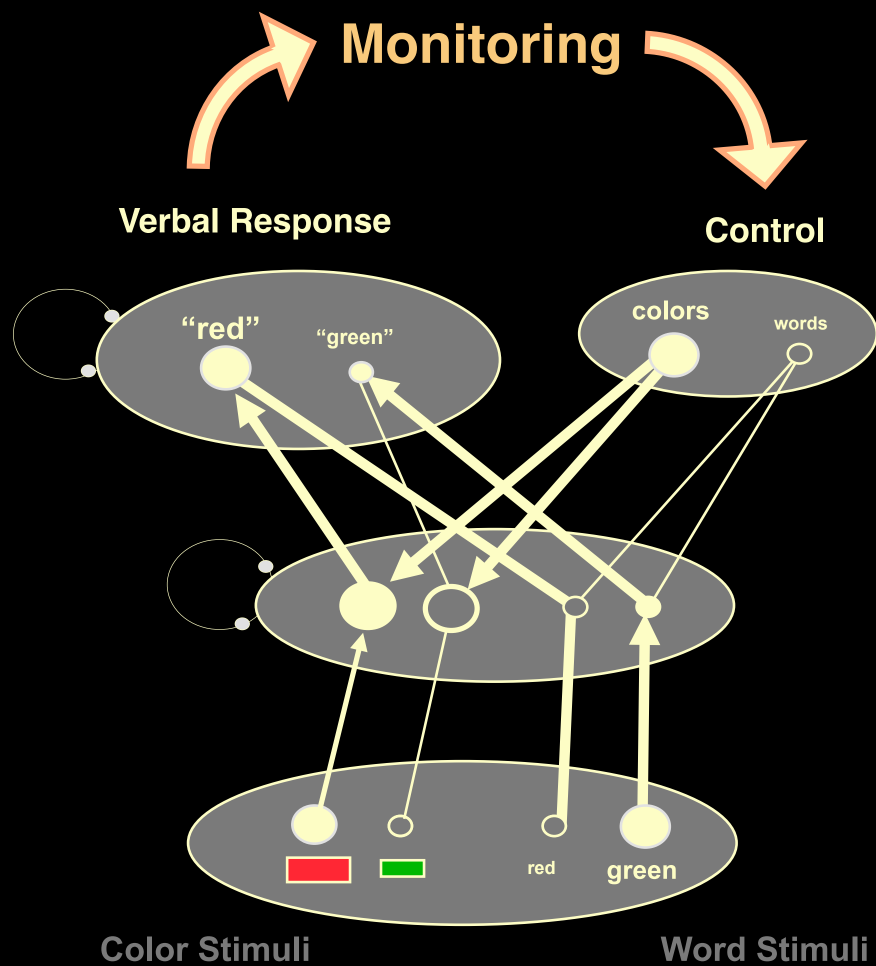

Stimulus: RED

Evidence for Performance Monitoring and Dynamic Adjustments in Control

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992))

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)
- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes"**
    Their operation can only be *inferred* from behavioral data

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)
- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes"**
    Their operation can only be *inferred* from behavioral data
  - **Measuring brain function may provide a more direct test...**

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)
- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes" Their operation can only be *inferred* from behavioral data**
  - **Measuring brain function may provide a more direct test…**
- **Neurophysiological Evidence:**
  - **Error Related Negativity (ERN)** *(Falkenstein et al. (1993); Gehring et al. (1993)*

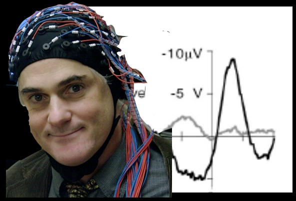

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)
- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes"**
    Their operation can only be *inferred* from behavioral data
  - **Measuring brain function may provide a more direct test...**
- **Neurophysiological Evidence:**
  - **Error Related Negativity (ERN)** *(Falkenstein et al. (1993); Gehring et al. (1993)*
  - **Anterior Cingulate Cortex (ACC) response to task difficulty**
    (Evans et al., 1992; Owens & Duncan, 2001)

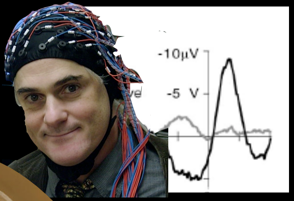

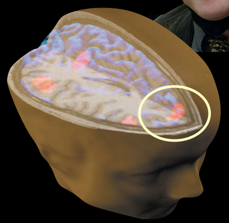

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)

- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes" Their operation can only be *inferred* from behavioral data**
  - **Measuring brain function may provide a more direct test...**

- **Neurophysiological Evidence:**
  - **Error Related Negativity (ERN)** *(Falkenstein et al. (1993); Gehring et al. (1993)*
  - **Anterior Cingulate Cortex (ACC) response to task difficulty**
    (Evans et al., 1992; Owens & Duncan, 2001)
  - **ACC is an error monitor!**

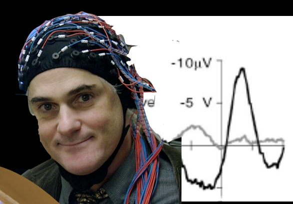

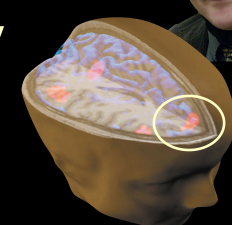

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)
- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes" Their operation can only be *inferred* from behavioral data**
  - **Measuring brain function may provide a more direct test...**
- **Neurophysiological Evidence:**
  - **Error Related Negativity (ERN)** *(Falkenstein et al. (1993); Gehring et al. (1993)*
  - **Anterior Cingulate Cortex (ACC) response to task difficulty**
    (Evans et al., 1992; Owens & Duncan, 2001)
  - **ACC is an error monitor!**

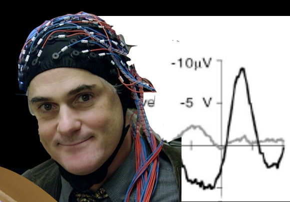

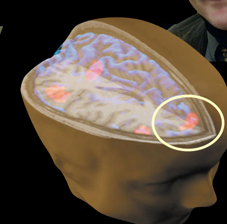

- ***Problems:***
  - **ACC activates when task is difficult but no errors are made** *(e.g., Stroop)*

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)
- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes"**
    Their operation can only be *inferred* from behavioral data
  - **Measuring brain function may provide a more direct test…**
- **Neurophysiological Evidence:**
  - **Error Related Negativity (ERN)** *(Falkenstein et al. (1993); Gehring et al. (1993)*
  - **Anterior Cingulate Cortex (ACC) response to task difficulty**
    (Evans et al., 1992; Owens & Duncan, 2001)
  - **ACC is an error monitor!**

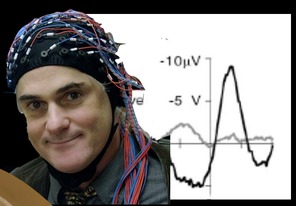

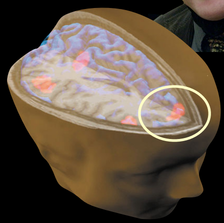

- ***Problems:***
  - **ACC activates when task is difficult but no errors are made** *(e.g., Stroop)*
  - **Who has the correct response? Why isn't s/he responding?**

Evidence for Performance Monitoring and Dynamic Adjustments in Control

- **Behavioral Evidence:**
  - **Subjects are slower and more accurate after errors**
    (Rabbitt, 1966; Laming, 1968)
  - **Subjects focus attention when interference increases**
    (Logan & Zbrodoff (1979); Gratton et al. (1992); Tzelgov et al. (1992)
- ***Problem:***
  - **Monitoring and adjustments of control are "internal processes" Their operation can only be *inferred* from behavioral data**
  - **Measuring brain function may provide a more direct test...**
- **Neurophysiological Evidence:**
  - **Error Related Negativity (ERN)** *(Falkenstein et al. (1993); Gehring et al. (1993)*
  - **Anterior Cingulate Cortex (ACC) response to task difficulty**
    (Evans et al., 1992; Owens & Duncan, 2001)
  - **ACC is an error monitor!**

- ***Problems:***
  - **ACC activates when task is difficult but no errors are made** *(e.g., Stroop)*
  - **Who has the correct response? Why isn't s/he responding?**

Hypothesis

Hypothesis

- ACC monitors processing conflict
  - ACC is activated by <u>competition</u> among task-relevant representations

Hypothesis

- ACC monitors processing conflict
  - ACC is activated by <u>competition</u> among task-relevant representations
  - Conflict predisposes to <u>errors</u>

    accounts for association between ACC activity, ERN and errors

Hypothesis

- **ACC monitors processing conflict**
  - **ACC is activated by <u>competition</u> among task-relevant representations**
  - **Conflict predisposes to <u>errors</u>**

    accounts for association between ACC activity, ERN and errors
  - **Computation:**

    product of activity of competing units

Hypothesis

- **ACC monitors processing conflict**
  - **ACC is activated by competition among task-relevant representations**
  - **Conflict predisposes to errors**

    accounts for association between ACC activity, ERN and errors
  - **Computation:**

    product of activity of competing units
- **Example…**

High Conflict

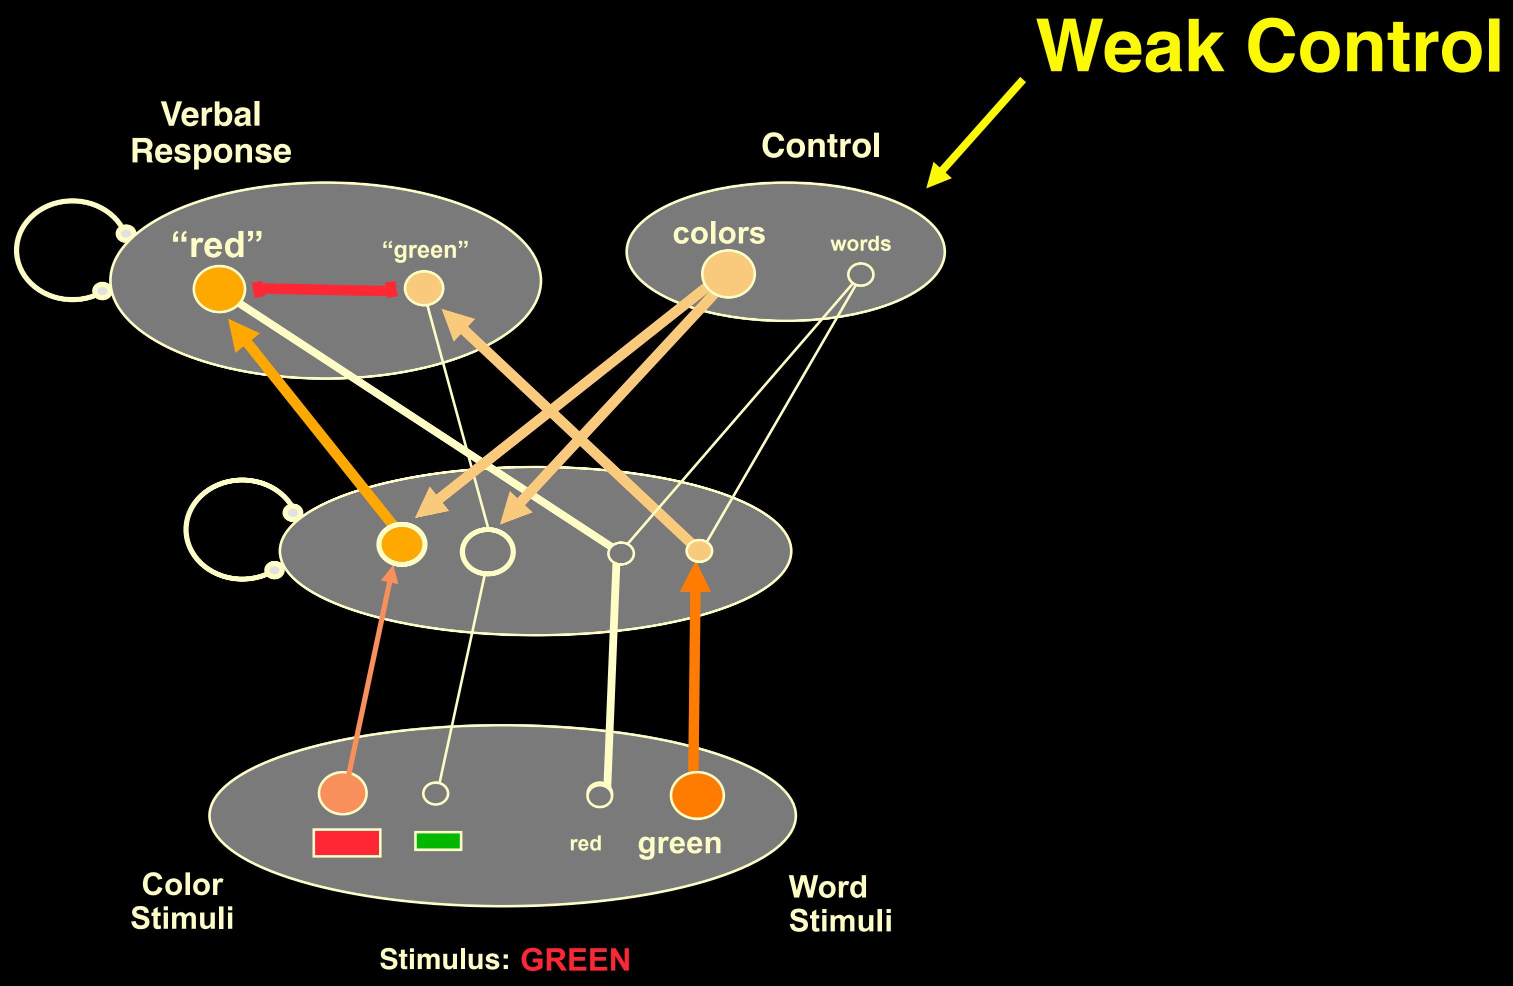

High Conflict

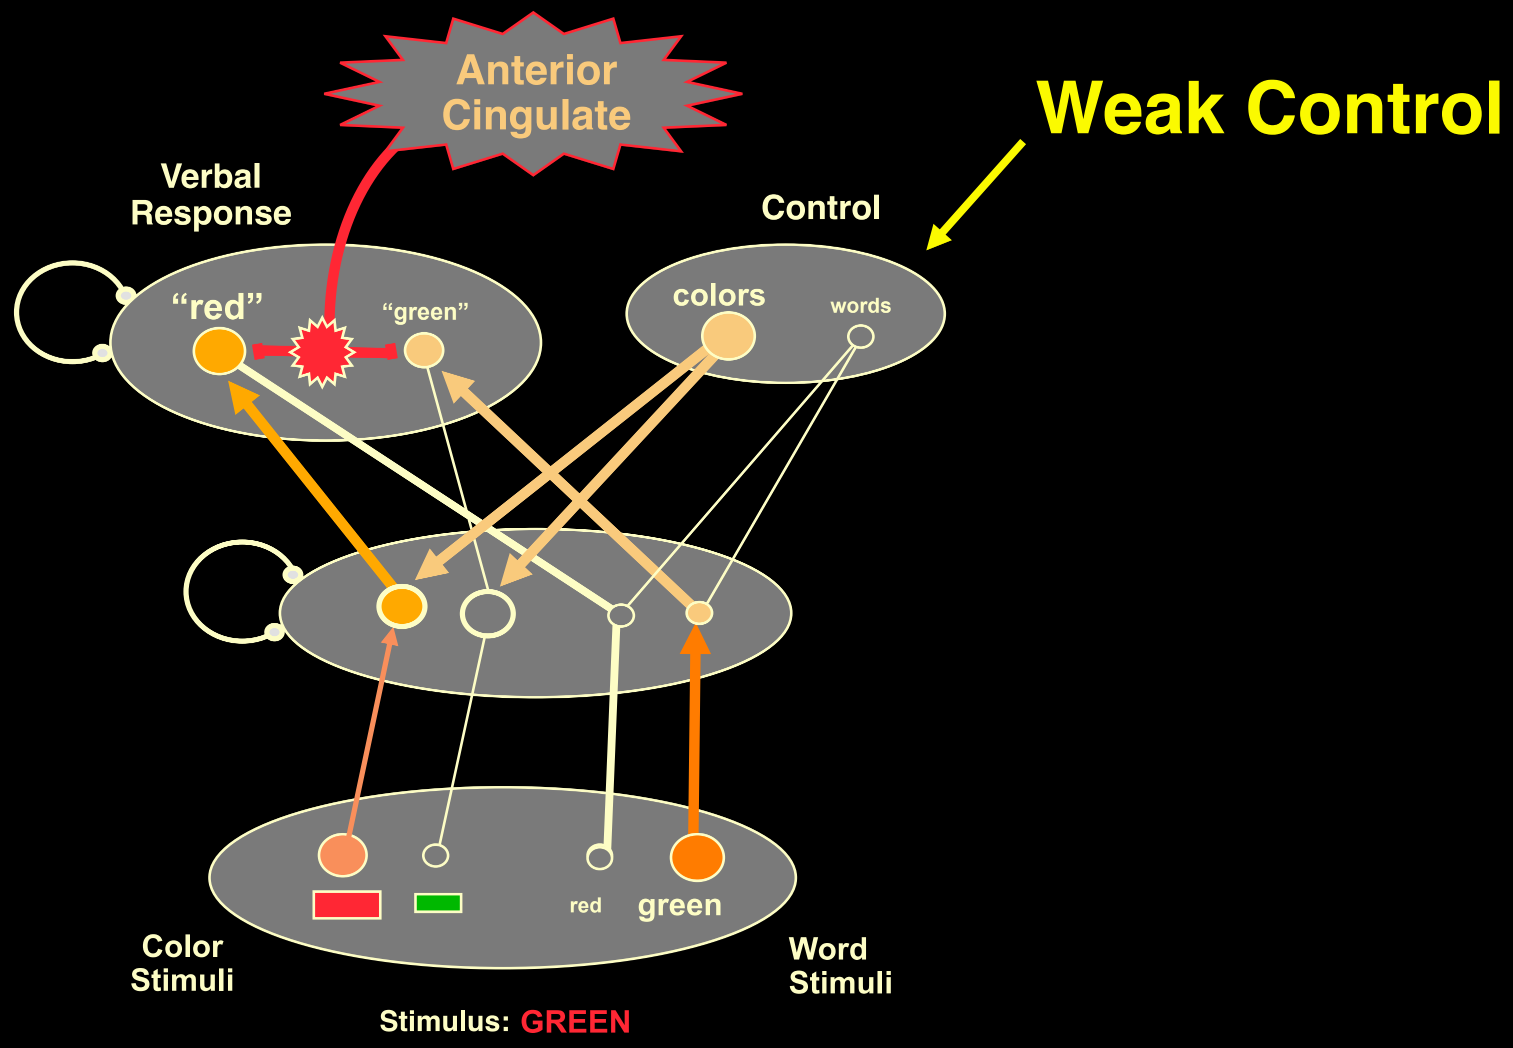

Augment Control

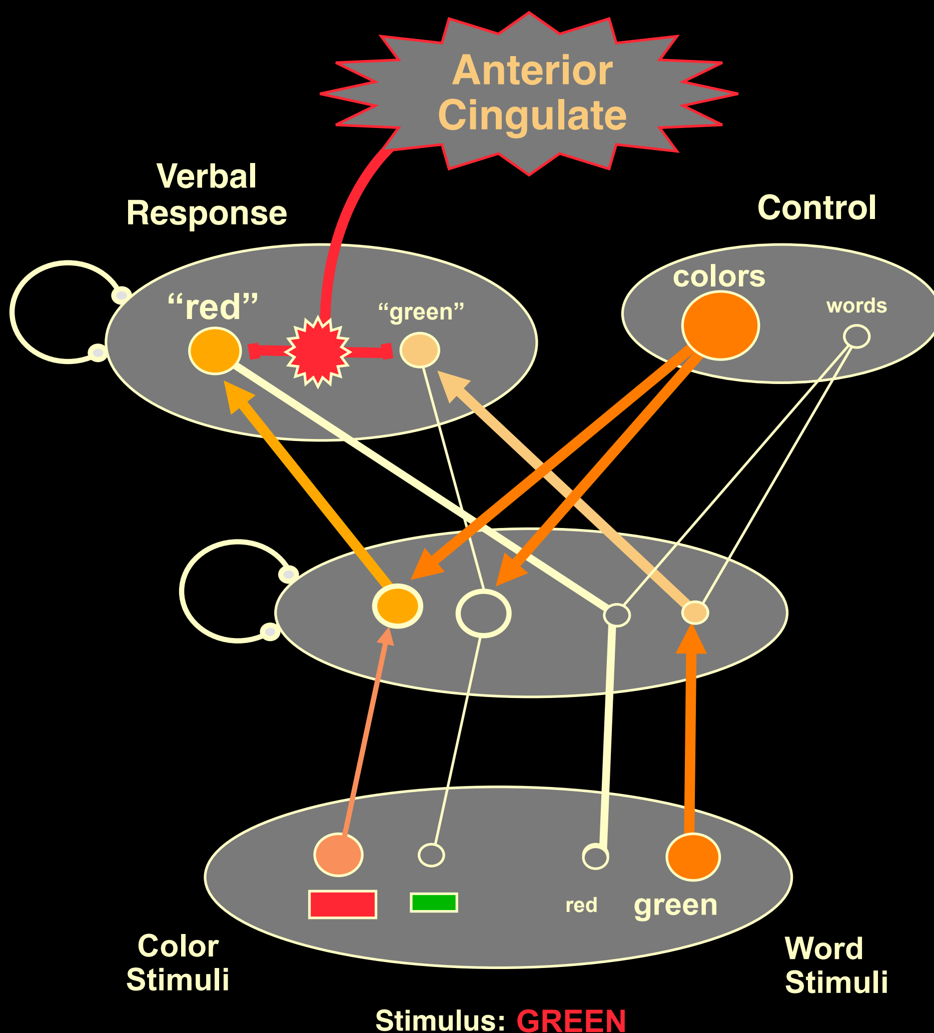

Augment Control

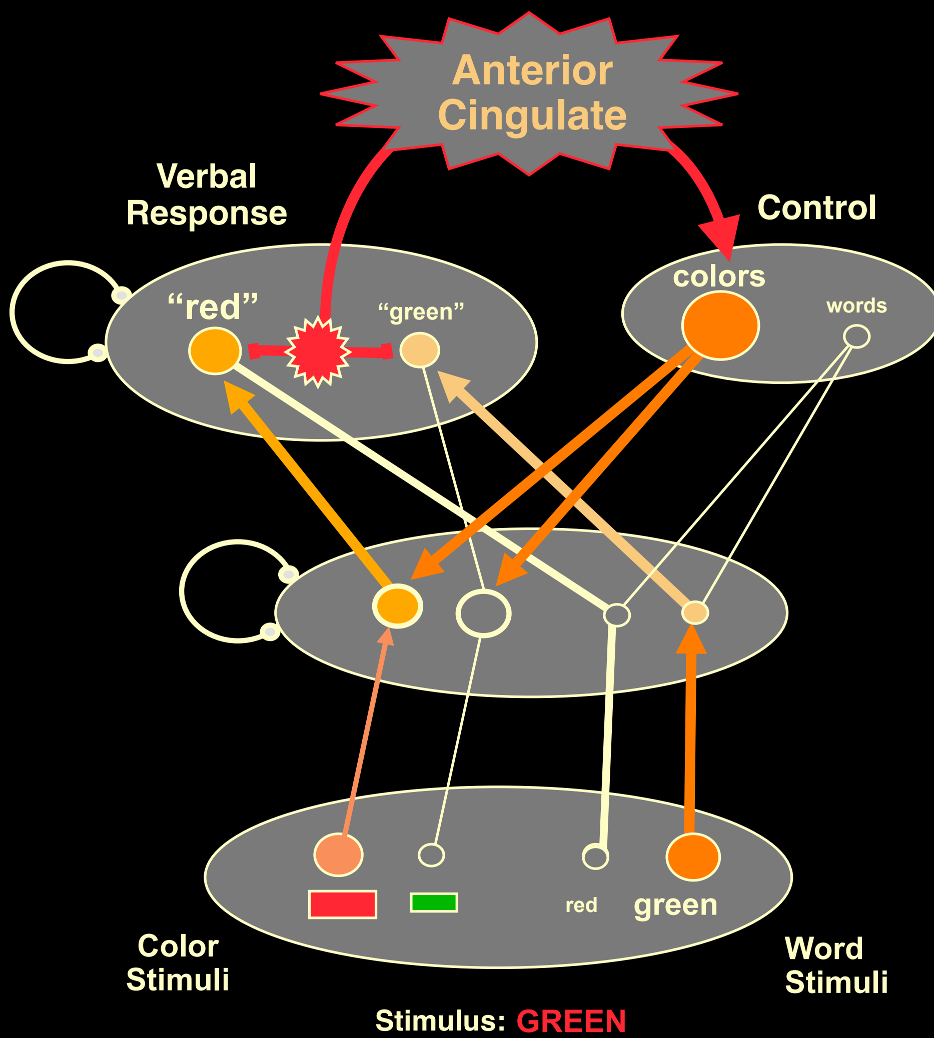

Lowers Conflict

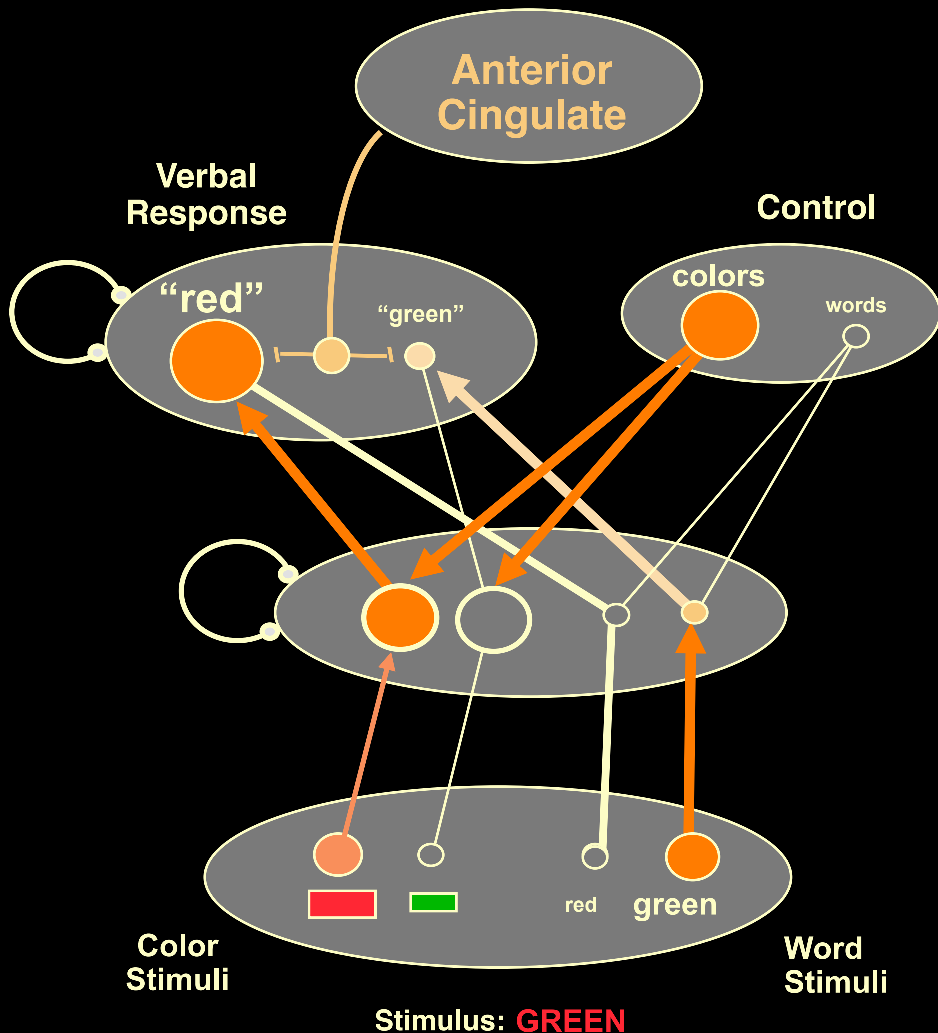

Conflict Monitoring and ACC

(Botvinick et al., 2001)

- **Hypothesis integrates a wide body of findings:**
  - *Stroop task*
    Stroop model *(Cohen et al. 1990)*
    PET: ACC activity *(Pardo et al., 1990)*

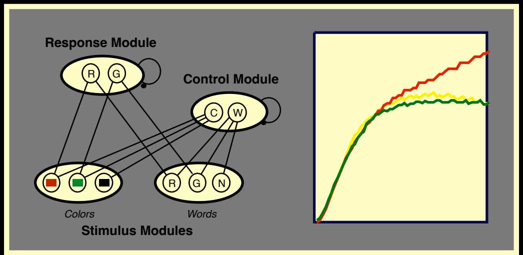

Conflict Monitoring and ACC

(Botvinick et al., 2001)

- **Hypothesis integrates a wide body of findings:**

  - ***Stroop task***
    Stroop model *(Cohen et al. 1990)*
    PET: ACC activity *(Pardo et al., 1990)*

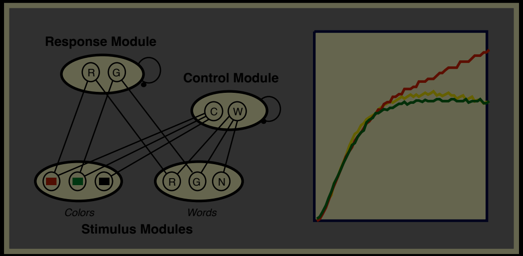

  - ***Eriksen flanker paradigm***
    Eriksen model *(Servan-Schreiber et al., 1992)*
    Electrophysiology: ERN *(Gehring et al., 1993)*

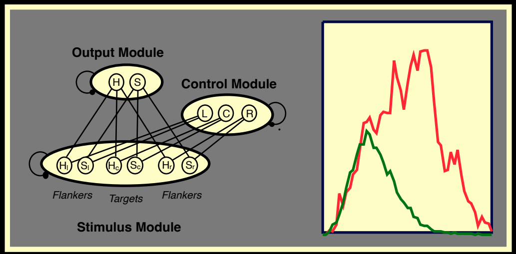

Conflict Monitoring and ACC

(Botvinick et al., 2001)

- **Hypothesis integrates a wide body of findings:**

  - ***Stroop task***
    Stroop model *(Cohen et al. 1990)*
    PET: ACC activity *(Pardo et al., 1990)*

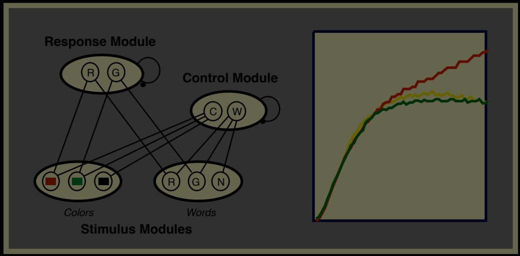

  - ***Eriksen flanker paradigm***
    Eriksen model *(Servan-Schreiber et al., 1992)*
    Electrophysiology: ERN *(Gehring et al., 1993)*

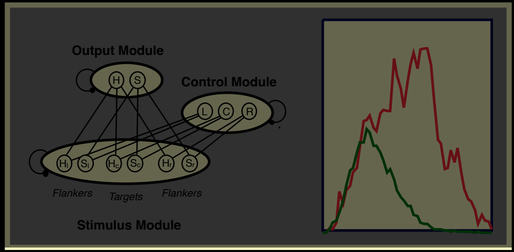

  - ***Stem completion*** **(and verbal fluency)**
    Interactive Activation Model *(McClelland and Rumelhart, 1981)*
    fMRI: ACC activity *(Thompson-Schill et al., 1998)*

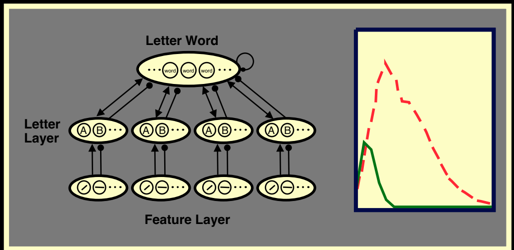

Conflict Monitoring and Control

(Botvinick et al., 2001)

Forced Choice RT Task
Sequential adjustment effects
(Laming., 1968)

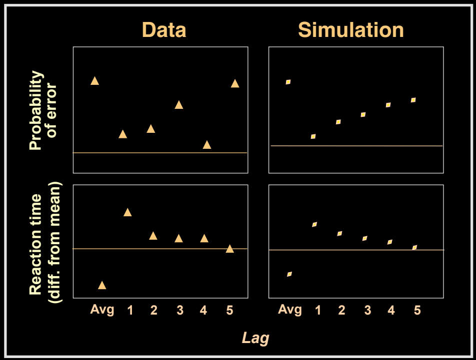

Eriksen Task
Sequential adjustment effects
(Gratton et al., 1992)

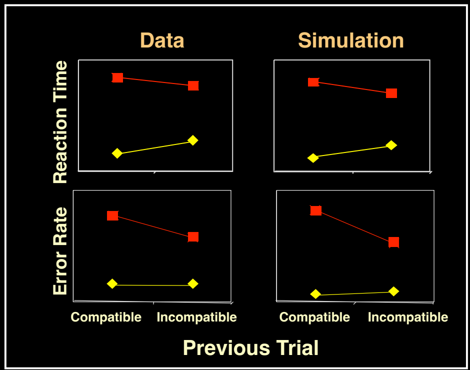

Stroop Task
Frequency effects
(Tzelgov, 1992)

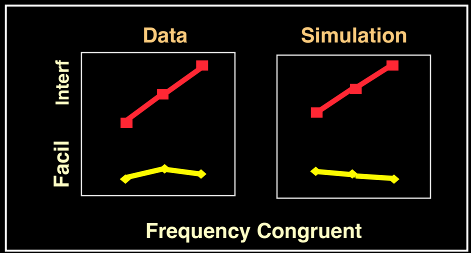

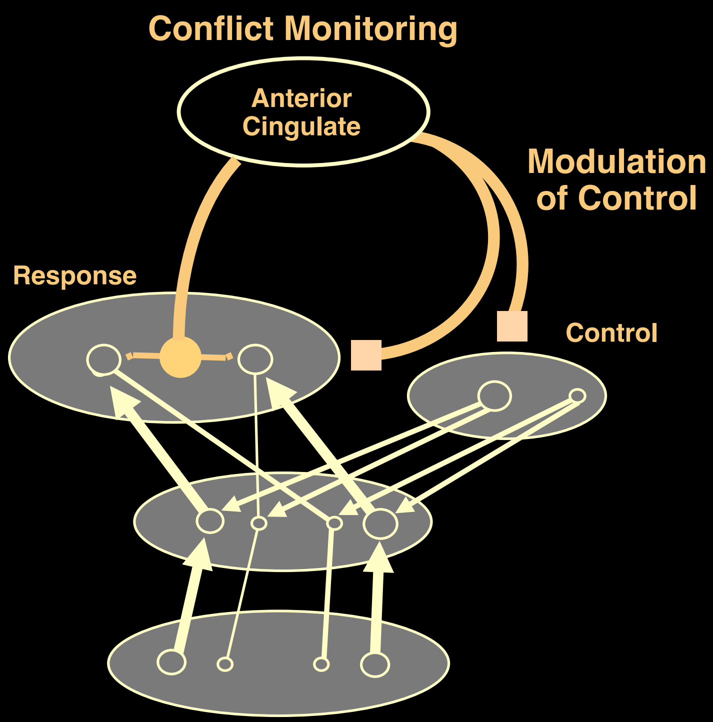

Dynamics of Conflict: Error trial

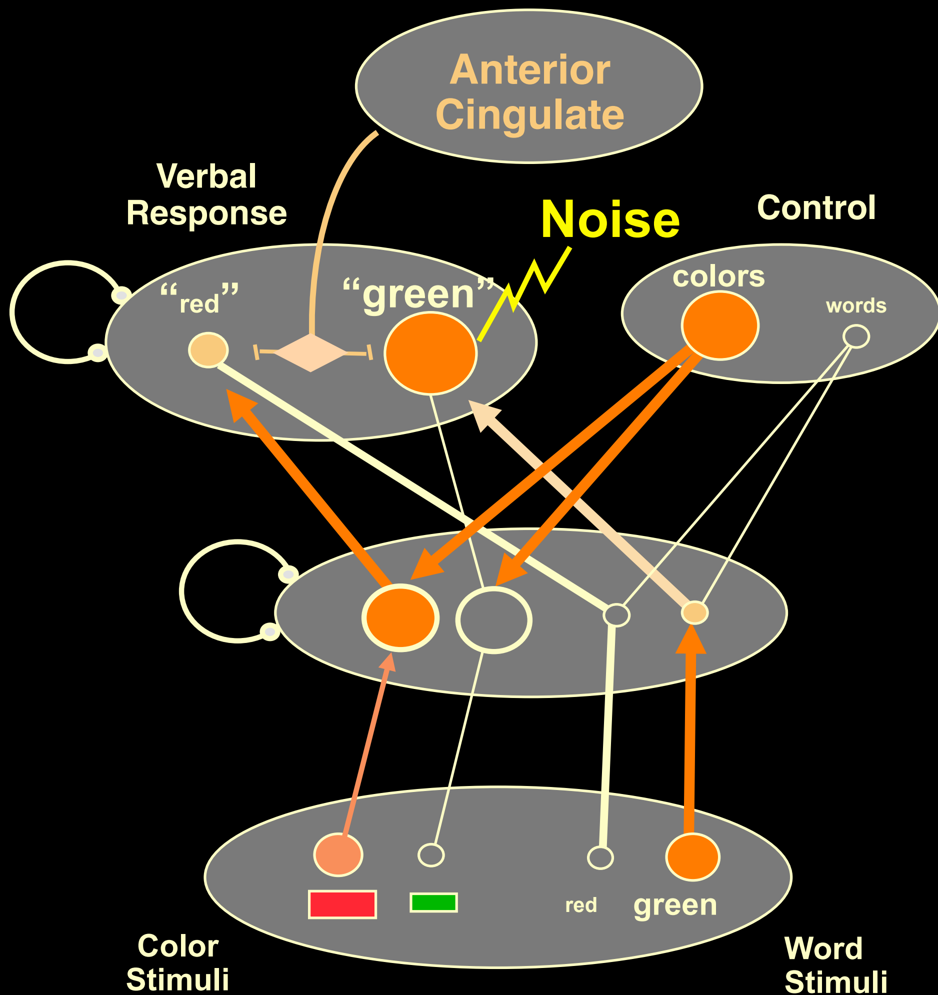

Stimulus: GREEN

Dynamics of Conflict: Post-processing and Conflict

Continued stimulus processing

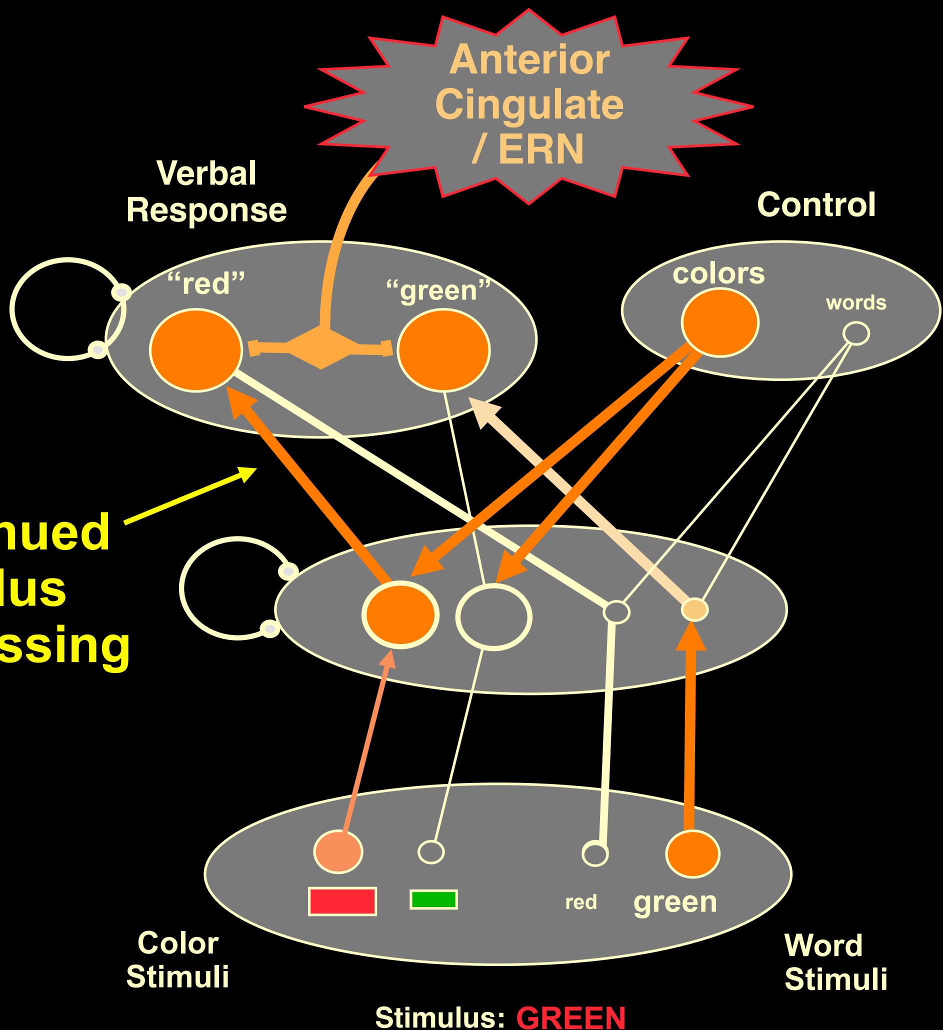

Conflict Monitoring and the ERN

Yeung, Botvinick & Cohen (under review)

Conflict Monitoring and the ERN

Yeung, Botvinick & Cohen (under review)

- **Model explains dynamics of electrophysiological findings:**
  - **Pre-response "negativity" for correct incompatible trials (N2C)**
    (Kopp, Rist & Mattler, 1996)

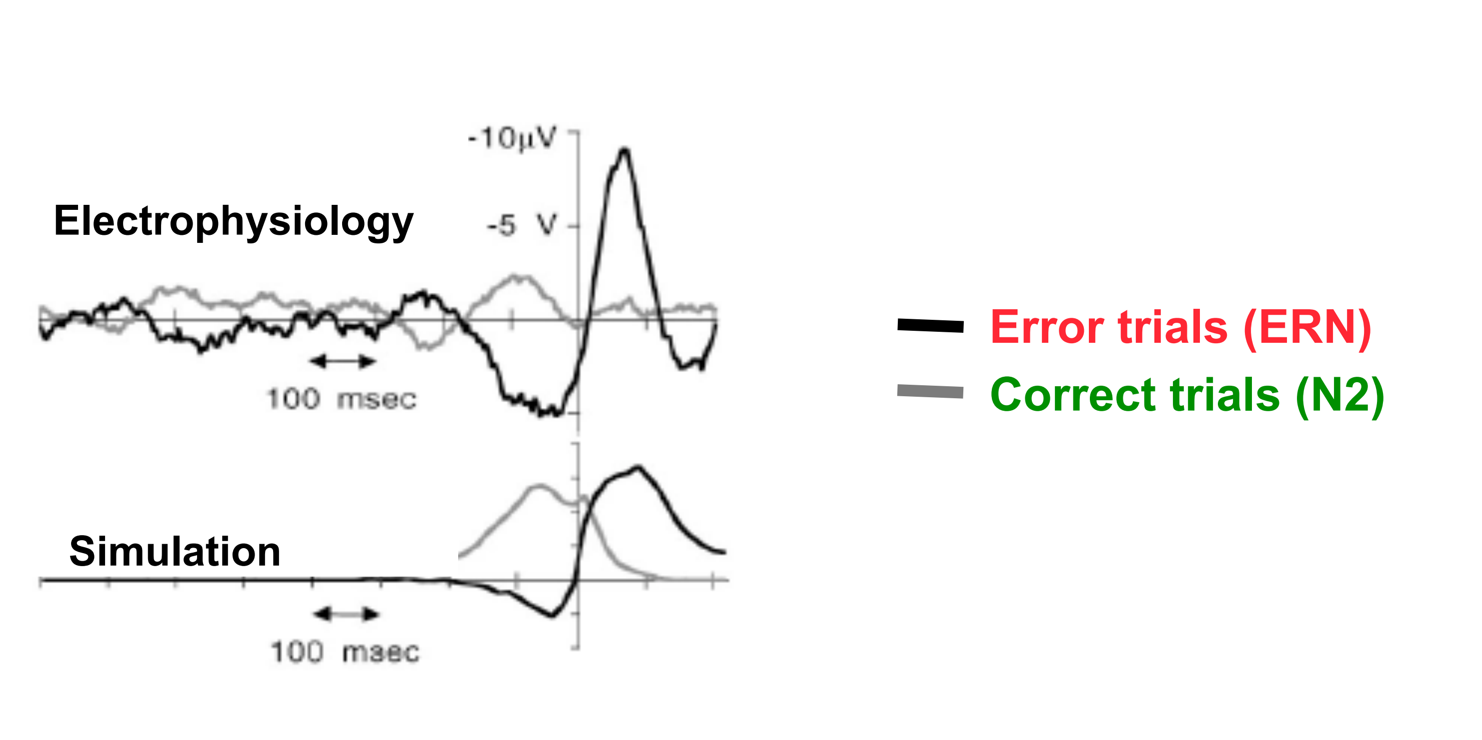

Conflict Monitoring and the ERN

Yeung, Botvinick & Cohen (under review)

- **Model explains dynamics of electrophysiological findings:**
  - **Pre-response "negativity" for correct incompatible trials (N2C)**
    (Kopp, Rist & Mattler, 1996)
  - **Larger ERN for *congruent* than incongruent stimuli on *error trials***
    (Scheffers & Coles, 2000)

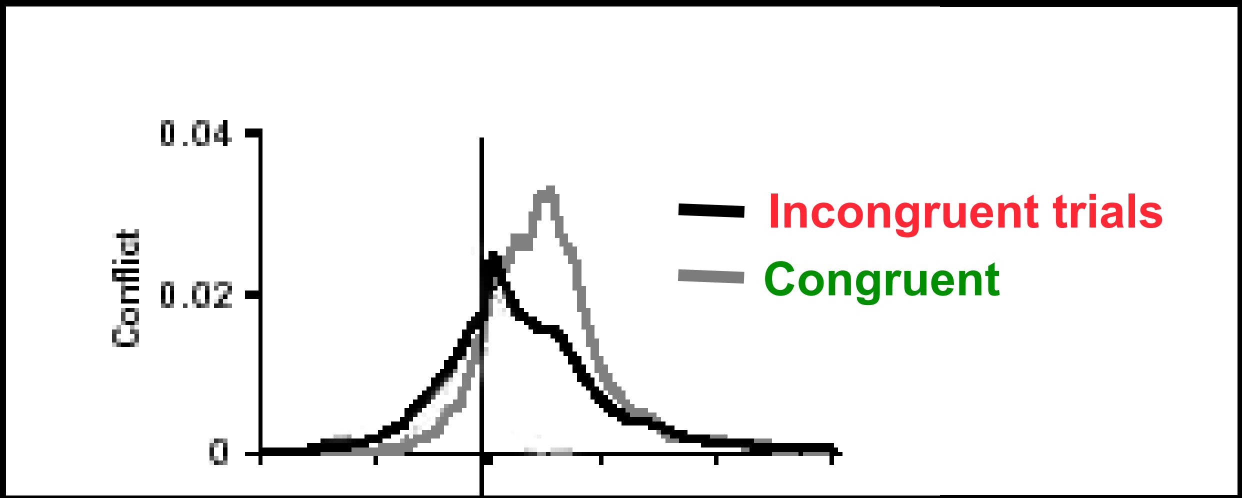

Control as Optimization of Performance

Control as Optimization of Performance

- **DDM:**
  - **Frequency: threshold adjustment**

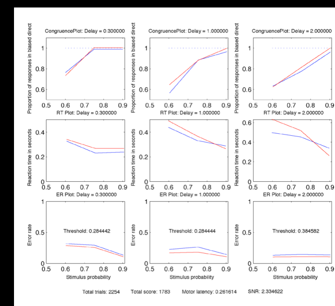

Control as Optimization of Performance

- DDM:
  - Frequency: threshold adjustment

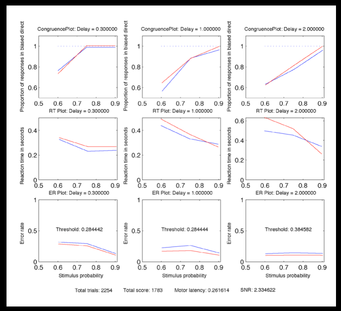

- Conflict monitoring:
  - Laming, Rabbitt effects: threshold adjustment

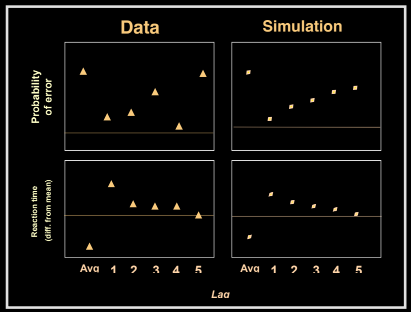

Control as Optimization of Performance

- **DDM:**
  - Frequency: threshold adjustment

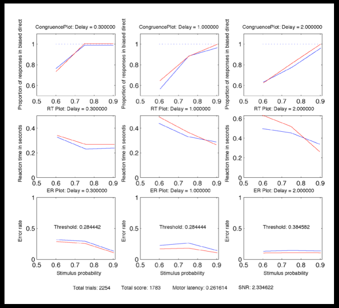

- **Conflict monitoring:**
  - Laming, Rabbitt effects: threshold adjustment

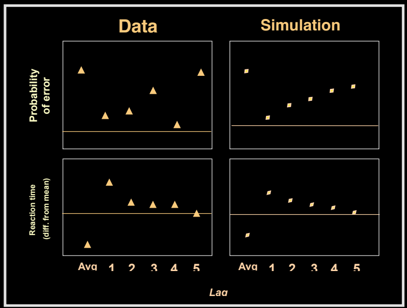

  - Gratton, Logan effects: attentional (internal drift rate) adjustment

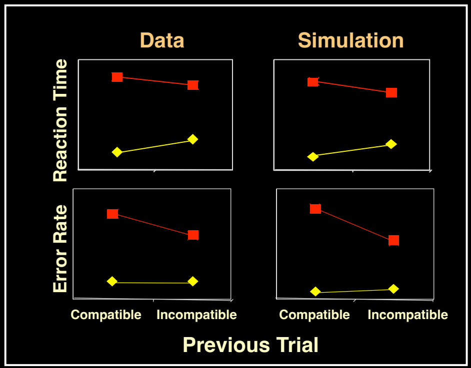

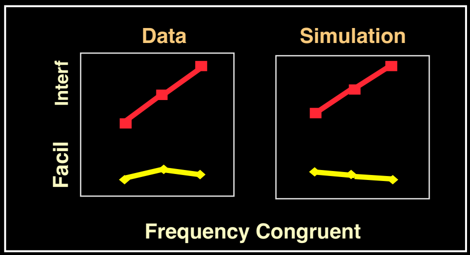

Control as Optimization of Performance

- DDM:
  - Frequency: threshold adjustment

- Conflict monitoring:
  - Laming, Rabbitt effects: threshold adjustment
  - Gratton, Logan effects: attentional (internal drift rate) adjustment

Control as Optimization of Performance

- DDM:
  - Frequency: threshold adjustment

    Speed-accuracy tradeoff: *optimize reward rate*

- Conflict monitoring:
  - Laming, Rabbitt effects: threshold adjustment
  - Gratton, Logan effects: attentional (internal drift rate) adjustment

Control as Optimization of Performance

- DDM:
  - Frequency: threshold adjustment

    Speed-accuracy tradeoff: *optimize reward rate*

- Conflict monitoring:
  - Laming, Rabbitt effects: threshold adjustment

    Speed-accuracy tradeoff: *optimize reward rate*

  - Gratton, Logan effects: attentional (internal drift rate) adjustment

Control as Optimization of Performance

- DDM:
  - Frequency: threshold adjustment

    Speed-accuracy tradeoff: *optimize reward rate*
- Conflict monitoring:
  - Laming, Rabbitt effects: threshold adjustment

    Speed-accuracy tradeoff: *optimize reward rate*
  - Gratton, Logan effects: attentional (internal drift rate) adjustment

    Attentional enhancement: *optimize reward rate*

Expected Value of Control Theory

Expected Value of Control Theory

- **Assumptions:**
  - **Control is allocated to *optimize performance* and maximize reward rate**

Expected Value of Control Theory

- **Assumptions:**
  - Control is allocated to *optimize performance* and maximize reward rate
  - Control allocation manages *limited capacity* (e.g., shared representations)

Expected Value of Control Theory

- **Assumptions:**
  - Control is allocated to *optimize performance* and maximize reward rate
  - Control allocation manages *limited capacity* (e.g., shared representations)
  - Allocation to a particular task carries an *intrinsic cost* $\propto$ *amount invested*

Expected Value of Control Theory

- **Assumptions:**
  - Control is allocated to *optimize performance* and maximize reward rate
  - Control allocation manages *limited capacity* (e.g., shared representations)
  - Allocation to a particular task carries an *intrinsic cost* $\propto$ *amount invested*

  $\therefore$ **Control is allocated based on a *cost-benefit analysis* of the**

Expected Value of Control Theory

- **Assumptions:**
  - Control is allocated to ***optimize performance*** and maximize reward rate
  - Control allocation manages ***limited capacity*** (e.g., shared representations)
  - Allocation to a particular task carries an ***intrinsic cost*** $\propto$ ***amount invested***

  ∴ Control is allocated based on a ***cost-benefit analysis*** of the

Expected Value of Control (EVC):

$$EVC(signal, state) = \left[\sum_i \Pr(outcome_i \mid signal, state) \times Value(outcome_i)\right] - Cost(signal)$$

Expected Value of Control Theory

- **Assumptions:**
  - Control is allocated to *optimize performance* and maximize reward rate
  - Control allocation manages *limited capacity* (e.g., shared representations)
  - Allocation to a particular task carries an *intrinsic cost* $\propto$ *amount invested*

  ∴ Control is allocated based on a *cost-benefit analysis* of the

  Expected Value of Control (EVC):

$$EVC(signal, state) = \left[\sum_i \Pr(outcome_i \mid signal, state) \times Value(outcome_i)\right] - Cost(signal)$$

  - ***Invest*** **in the task(s) that promise the** ***greatest returns*** **on investment**

Expected Value of Control Theory

- **Assumptions:**
  - Control is allocated to *optimize performance* and maximize reward rate
  - Control allocation manages *limited capacity* (e.g., shared representations)
  - Allocation to a particular task carries an *intrinsic cost* $\propto$ *amount invested*

  ∴ Control is allocated based on a *cost-benefit analysis* of the

  Expected Value of Control (EVC):

$$EVC(signal, state) = \left[\sum_i \Pr(outcome_i \mid signal, state) \times Value(outcome_i)\right] - Cost(signal)$$

  - *Invest* in the task(s) that promise the *greatest returns* on investment

The Law of Less Work

If two or more behavioral sequences, each involving a different amount of energy consumption or work, have been equally well reinforced an equal number of times, the organism will gradually learn to choose the less laborious behavior sequence leading to the attainment of the reinforcing state of affairs.

-- C. Hull, 1943

The Law of Less Work

If two or more behavioral sequences, each involving a different amount of energy consumption or work, have been equally well reinforced an equal number of times, the organism will gradually learn to choose the less laborious behavior sequence leading to the attainment of the reinforcing state of affairs.

-- C. Hull, 1943

The Law of Less Work

If two or more behavioral sequences, each involving a different amount of energy consumption or work, have been equally well reinforced an equal number of times, the organism will gradually learn to choose the less laborious behavior sequence leading to the attainment of the reinforcing state of affairs.

-- C. Hull, 1943

Control is “Costly”

Control is "Costly"

- **Empirical evidence:**
  - **Choose less demanding tasks:**
    - **Demand Selection Task (DST)** *(Kools et al., 2010)*

Control is “Costly”

- Empirical evidence:
  - Choose less demanding tasks:
    - Demand Selection Task (DST) *(Kools et al., 2010)*
    - Stroop DST *(Bustamante et al, 2023)*

Control is "Costly"

- **Empirical evidence:**
  - **Choose less demanding tasks:**
    - **Demand Selection Task (DST)** *(Kools et al., 2010)*
    - **Stroop DST** *(Bustamante et al, 2023)*
  - **Choose less demanding strategy** *(Lieder et al., 2014)*

Control is "Costly"

- **Empirical evidence:**
  - **Choose less demanding tasks:**
    - **Demand Selection Task (DST)** *(Kools et al., 2010)*
    - **Stroop DST** *(Bustamante et al, 2023)*
  - **Choose less demanding strategy** *(Lieder et al., 2014)*
  - **Exert more effort when provided more reward** *(Padmala & Pessoa, 2011)*

Control is “Costly”

- **Empirical evidence:**
  - **Choose less demanding tasks:**
    - **Demand Selection Task (DST)** *(Kools et al., 2010)*
    - **Stroop DST** *(Bustamante et al, 2023)*
  - **Choose less demanding strategy** *(Lieder et al., 2014)*
  - **Exert more effort when provided more reward** *(Padmala & Pessoa, 2011)*
  - **Pay to avoid more demanding tasks** *(Westbrook et al., 2013)*

Control is "Costly"

- **Empirical evidence:**
  - **Choose less demanding tasks:**
    - **Demand Selection Task (DST)** *(Kools et al., 2010)*
    - **Stroop DST** *(Bustamante et al, 2023)*
  - **Choose less demanding strategy** *(Lieder et al., 2014)*
  - **Exert more effort when provided more reward** *(Padmala & Pessoa, 2011)*
  - **Pay to avoid more demanding tasks** *(Westbrook et al., 2013)*
  - **Discount reward associated with more effortful tasks…**

Control is “Costly”

Botvinick et al. (2009)

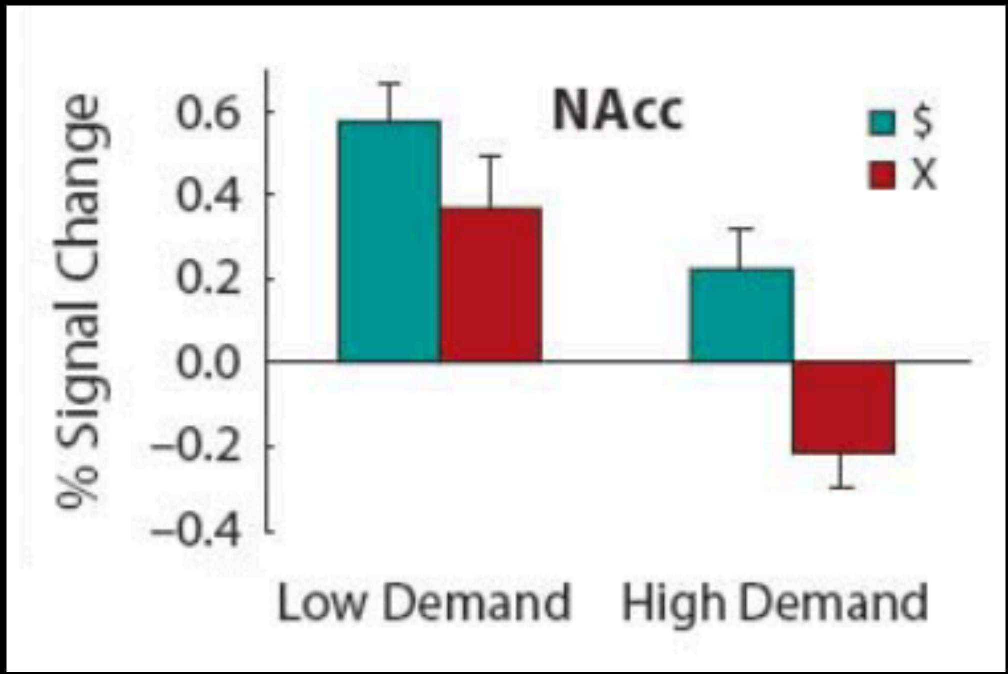

Expected Value of Control Theory

- 3 components to the control system (executive function):
  - ***Monitoring*** *(encode behavioral values)*
  - ***Specification*** *(select behaviors)*
  - ***Regulation*** *(implement selected behaviors)*

Expected Value of Control Theory

- 3 components to the control system (executive function):
  - ***Monitoring*** *(encode behavioral values)*
  - ***Specification*** *(select behaviors)*
  - ***Regulation*** *(implement selected behaviors)*

Expected Value of Control Theory

- **3 components to the control system (executive function):**
  - ***Monitoring*** *(encode behavioral values)*
    - ***evaluate outcomes*** **of performance of a given task**
      (e.g. ,verbal feedback, material rewards and/or punishments)
    - ***evaluate internal states*** **relevant to demands for control**
      (e.g., conflict, internally detected errors, task difficulty)
  - ***Specification*** *(select behaviors)*
  - ***Regulation*** *(implement selected behaviors)*

Expected Value of Control Theory

- 3 components to the control system (executive function):
  - *Monitoring* *(encode behavioral values)*
    - ***evaluate outcomes*** **of performance of a given task** *(e.g. ,verbal feedback, material rewards and/or punishments)*
    - ***evaluate internal states*** **relevant to demands for control** *(e.g., conflict, internally detected errors, task difficulty)*
  - *Specification* *(select behaviors)*
    - ***computation of EVC*** **for candidate control signals** *(i.e., those corresponding to current set of available tasks)*
    - ***selection of signal(s)*** **with maximum EVC, within "budgetary" constraints**
  - *Regulation* *(implement selected behaviors)*

Expected Value of Control Theory

- 3 components to the control system (executive function):
  - ***Monitoring*** *(encode behavioral values)*
    - ***evaluate outcomes*** **of performance of a given task** *(e.g. ,verbal feedback, material rewards and/or punishments)*
    - ***evaluate internal states*** **relevant to demands for control** *(e.g., conflict, internally detected errors, task difficulty)*
  - ***Specification*** *(select behaviors)*
    - ***computation of EVC*** **for candidate control signals** *(i.e., those corresponding to current set of available tasks)*
    - ***selection of signal(s)*** **with maximum EVC, within "budgetary" constraints**
  - ***Regulation*** *(implement selected behaviors)*
    - ***"top-down" biasing*** **of cortical pathways require to perform task** *(e.g., **"attentional selection"**, goal maintenance, motor control)*
    - ***systemwide changes*** **in parameters** ***(neuromodulation)***
      - **changes in learning rate, gating of new control signals into regulatory system**
      - **noise/gain modulation, explore/exploit tradeoff**
      - **threshold modulation**

Heart of EVC: Specification

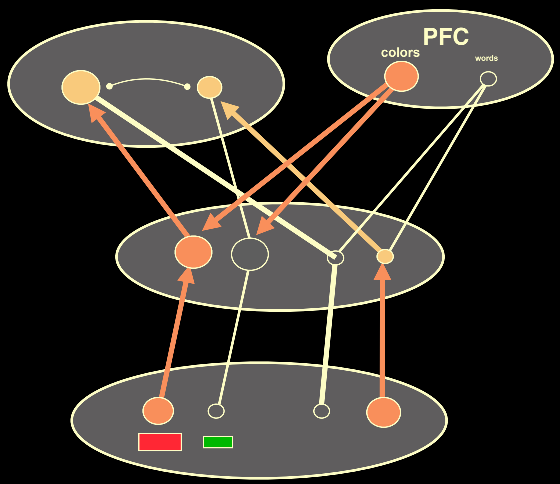

Heart of EVC: Specification

- Formalize EVC in terms of "control signals"
  - Control signals are defined by two parameters:

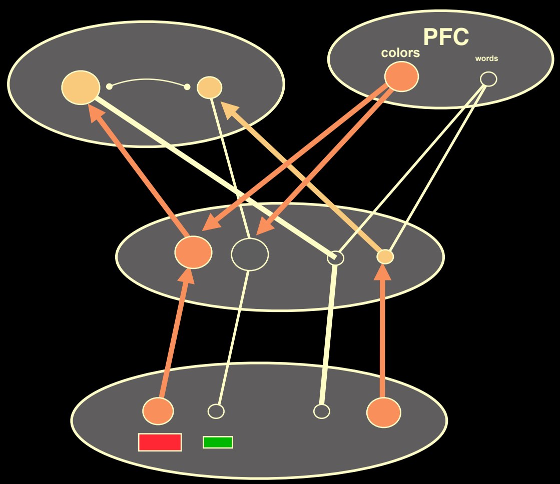

Heart of EVC: Specification

- Formalize EVC in terms of "control signals"
  - Control signals are defined by two parameters:
    - *Identity* — the task optimized by the signal
      (e.g., color naming or word reading)

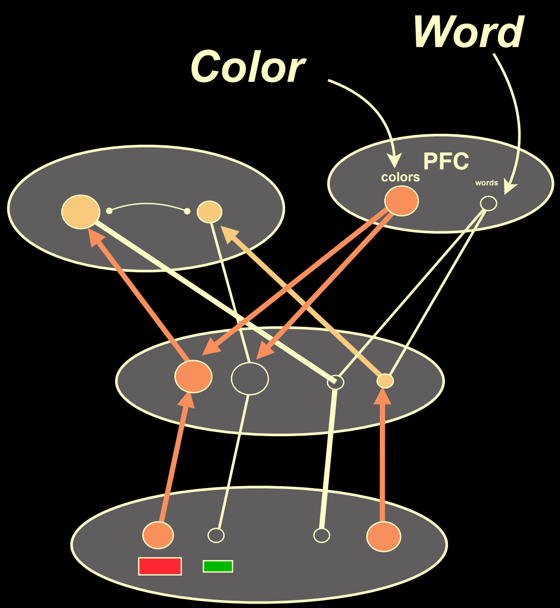

Heart of EVC: Specification

- Formalize EVC in terms of "control signals"
  - Control signals are defined by two parameters:
    - *Identity* — the task optimized by the signal
      (e.g., color naming or word reading)
    - *Intensity* — the strength allocated to a given task

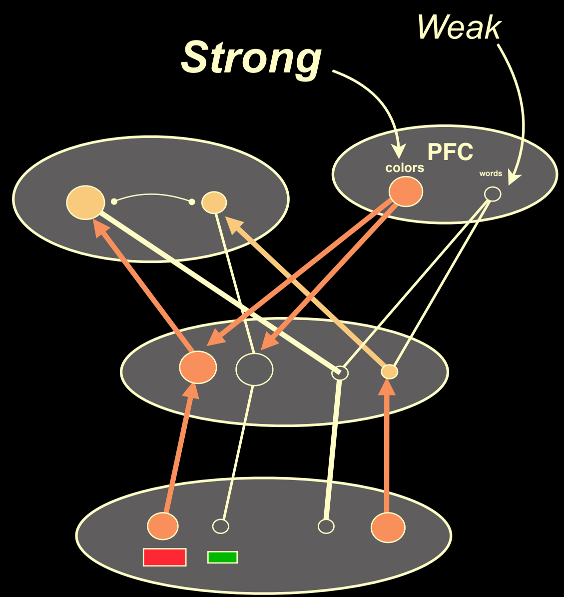

Toward a Formal Account

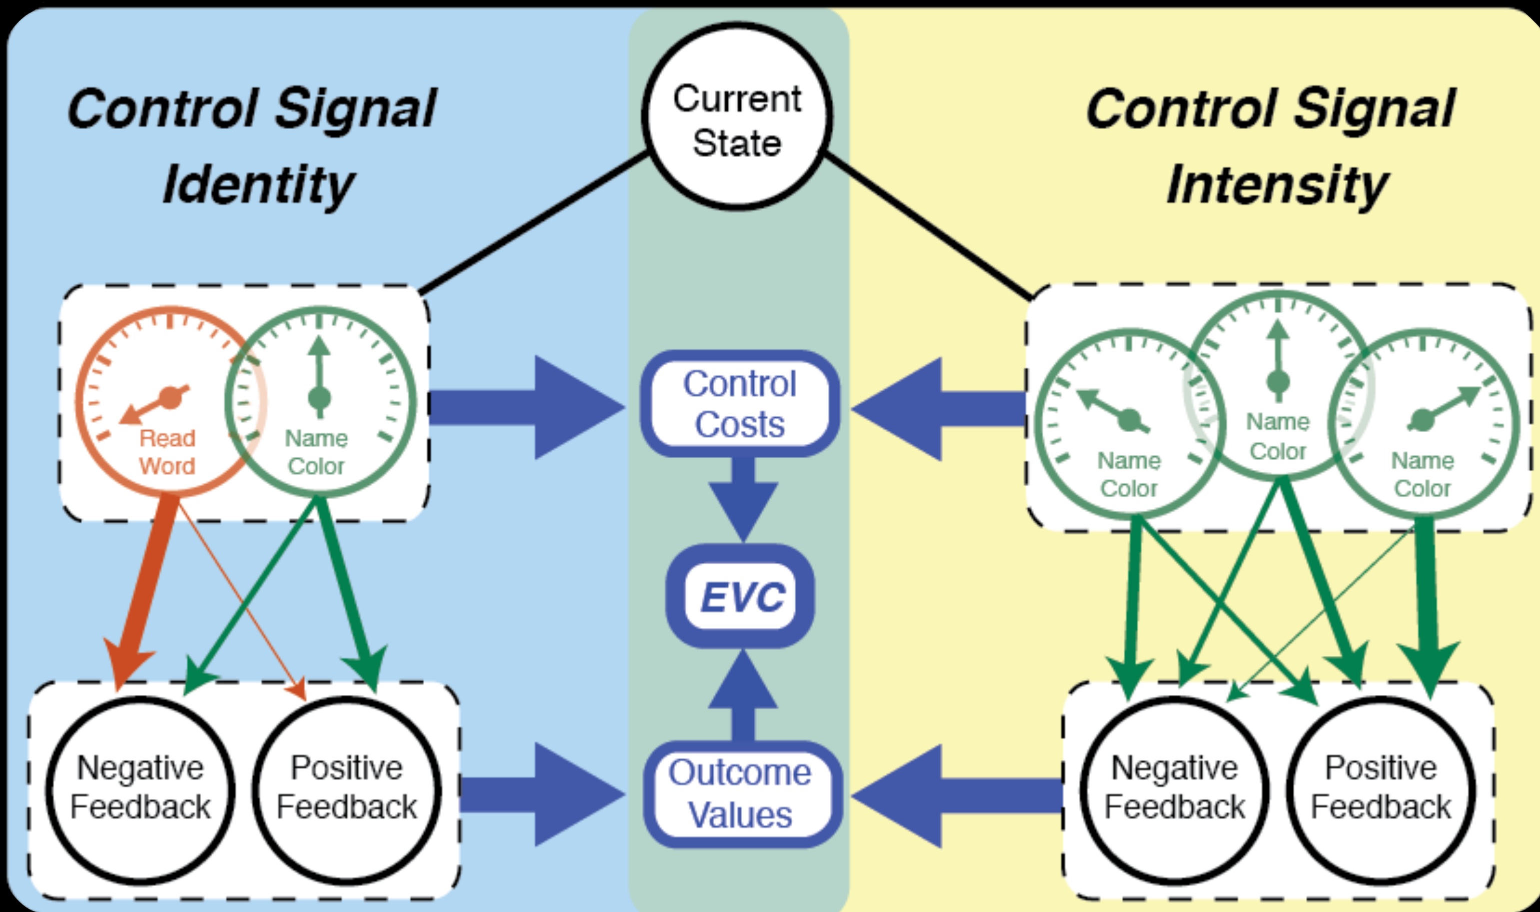

Heart of EVC: Specification

- **Formalize EVC in terms of "control signals"**
  - **Control signals are defined by two parameters:**
    - ***Identity*** **— the task optimized by the signal**
      (e.g., color naming or word reading)
    - ***Intensity*** **— the strength allocated to a given task**
  - **Reward (probability of success) and *cost***
    scale monotonically with intensity…

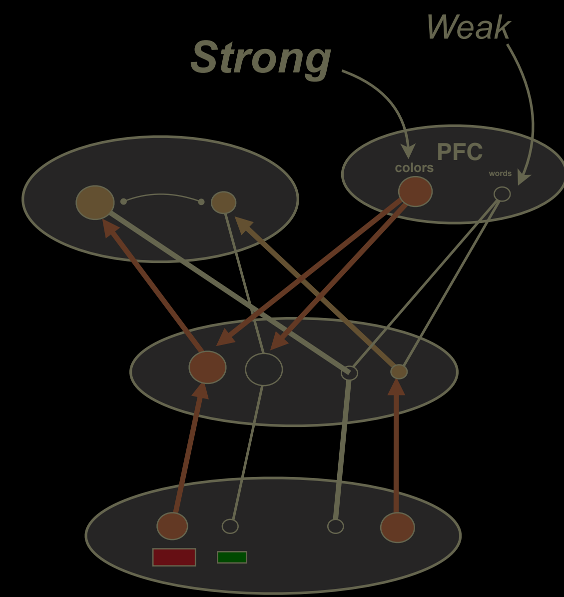

EVC

- Assume standard functional forms for:

EVC

- **Assume standard functional forms for:**
  - **COST:** ***accelerating*** *(e.g., exponential)*

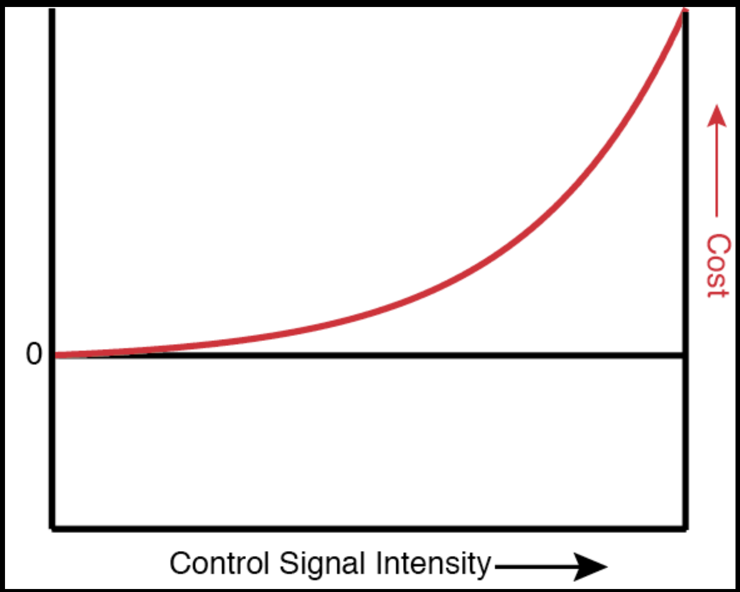

EVC

- **Assume standard functional forms for:**
  - **COST:** ***accelerating*** *(e.g., exponential)*
  - **REWARD:** ***saturating*****, diminishing returns** *(e.g., logarithmic)*

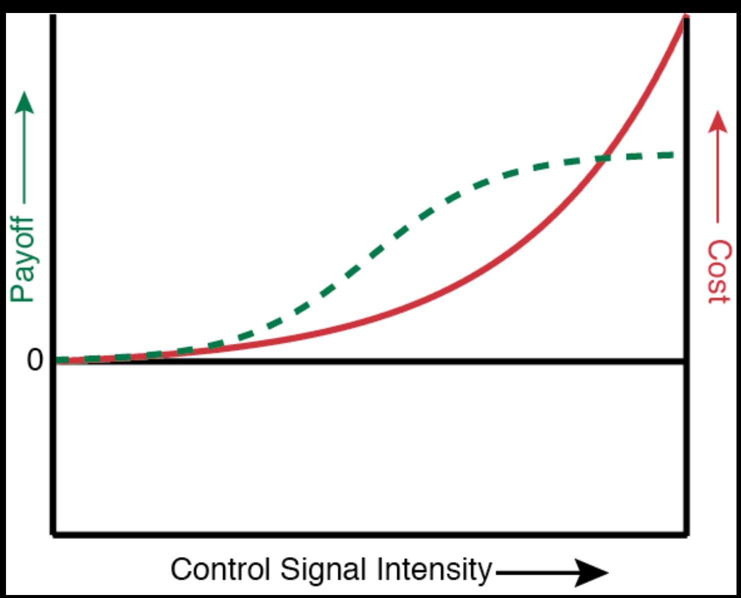

EVC

- Assume standard functional forms for:
  - COST: *accelerating* *(e.g., exponential)*
  - REWARD: *saturating*, diminishing returns *(e.g., logarithmic)*
- Yields *non-monotonic* function for EVC that

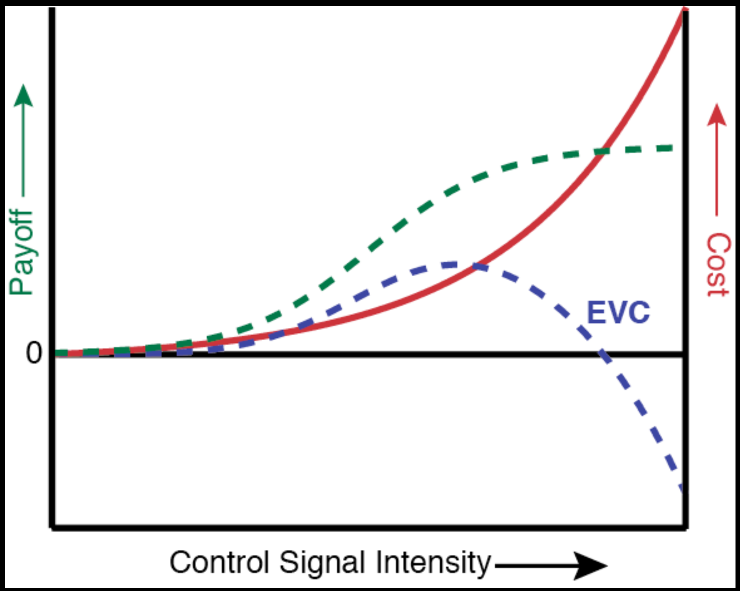

EVC

- Assume standard functional forms for:
  - COST: *accelerating* *(e.g., exponential)*
  - REWARD: *saturating*, diminishing returns *(e.g., logarithmic)*
- Yields *non-monotonic* function for EVC that obey *empirically observed* effects of incentives:

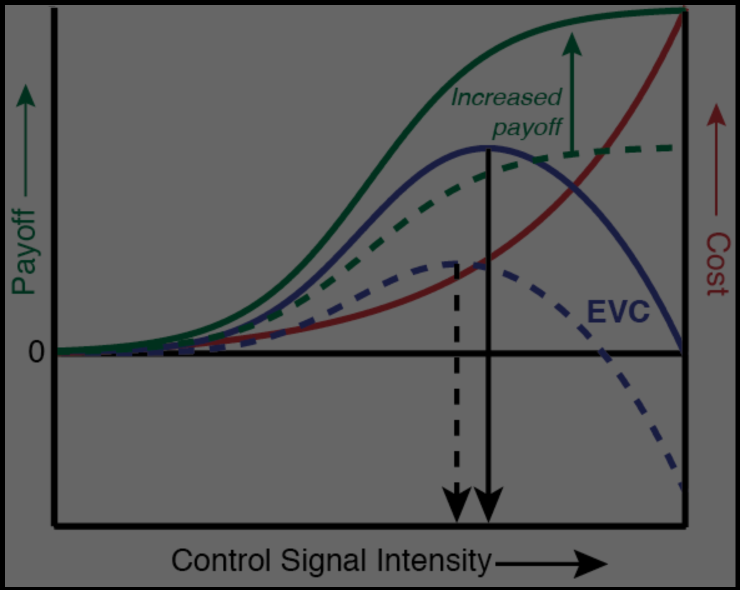

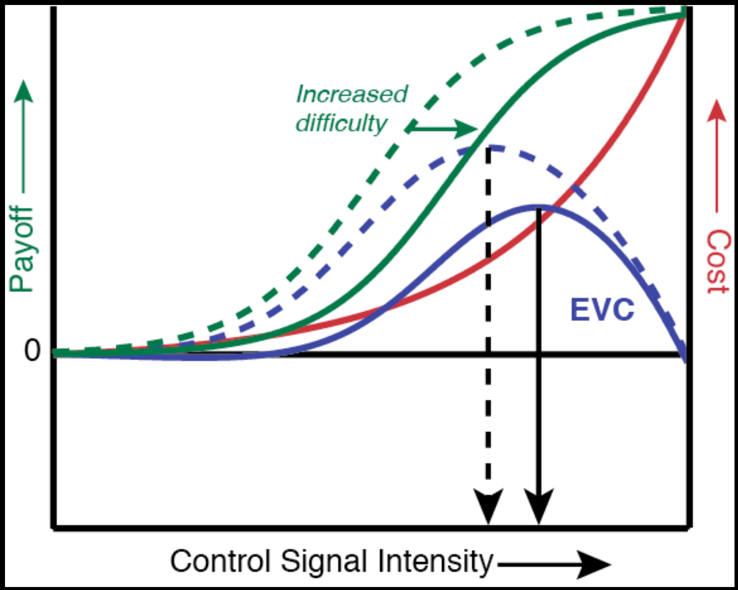

Heart of EVC: Specification

$$EVC(signal, state) = \left[\sum_i \Pr(outcome_i \mid signal, state) \times Value(outcome_i)\right] - Cost(signal)$$

Heart of EVC: Specification

- **Formalize EVC in terms of "control signals"**
  - **Control signals are defined by two parameters:**
    - ***Identity* — the task optimized by the signal**
      (e.g., color naming or word reading)
    - ***Intensity* — the strength allocated to a given task**
  - ***Reward* (probability of success) and *cost scale monotonically with intensity***
  - **The EVC for a given control signal is:**

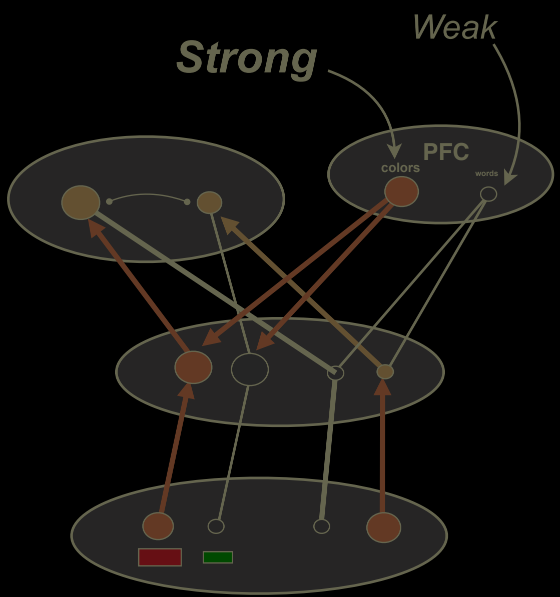

$$EVC(signal, state) = \left[\sum_i \Pr(outcome_i \mid signal, state) \times Value(outcome_i)\right] - Cost(signal)$$

Heart of EVC: Specification

- **Formalize EVC in terms of "control signals"**
  - Control signals are defined by two parameters:
    - *Identity* — the task optimized by the signal *(e.g., color naming or word reading)*
    - *Intensity* — the strength allocated to a given task
  - *Reward* (probability of success) and *cost scale monotonically with intensity*
  - The EVC for a given control signal is:

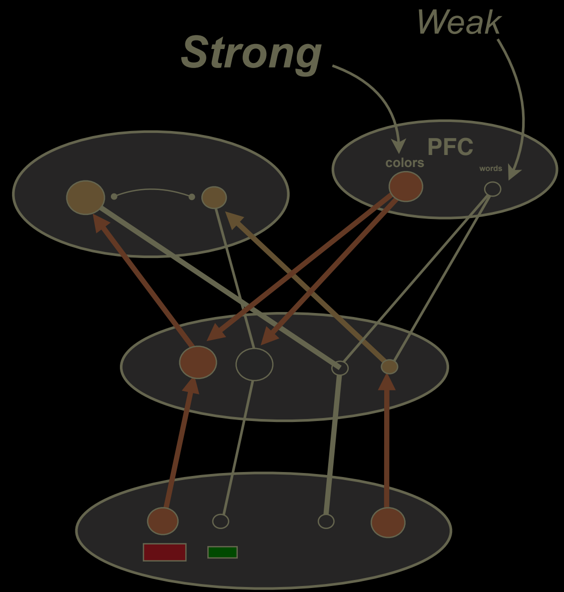

$$EVC(signal, state) = \left[\sum_i \Pr(outcome_i \mid signal, state) \times Value(outcome_i)\right] - Cost(signal)$$

  - Of the candidate set of control signals (≈ currently executable tasks), one(s) with *highest EVC* (that fall within budget) are *selected for allocation*

Neural Architecture of EVC

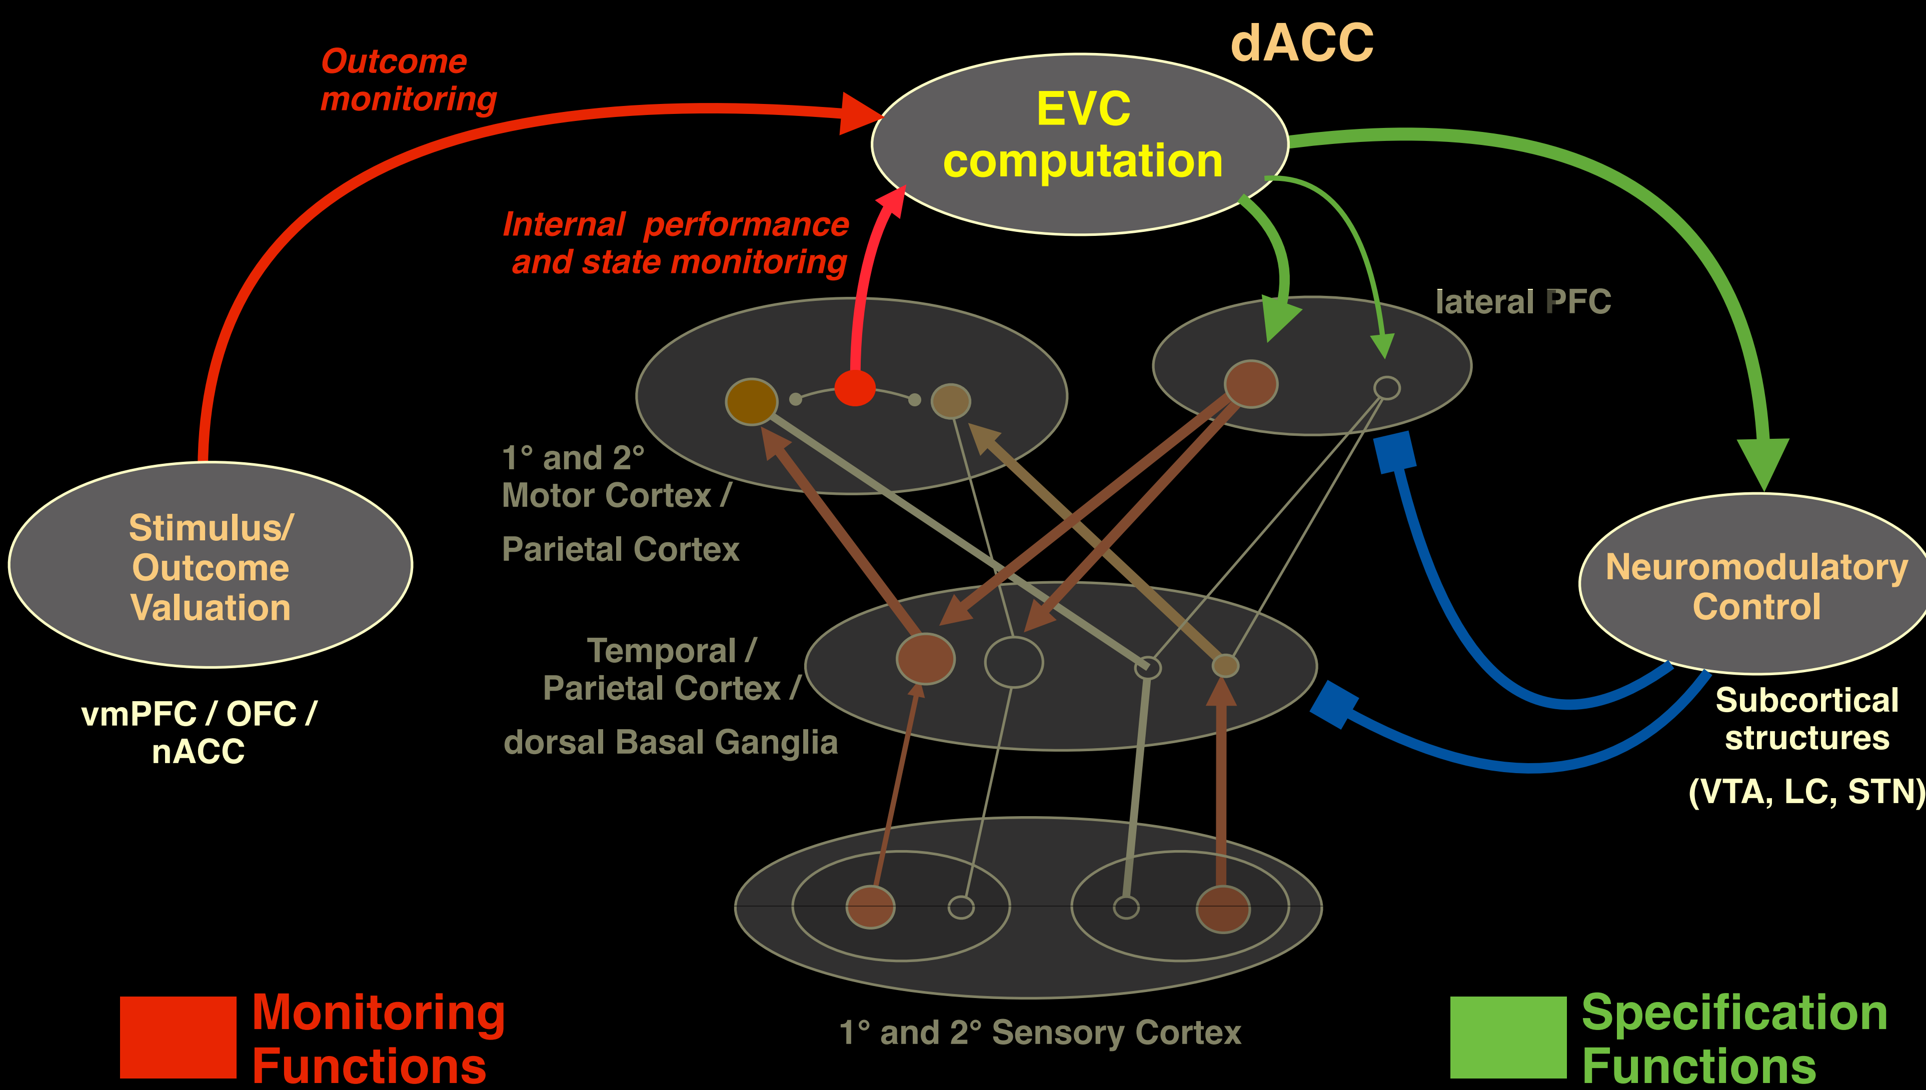

EVC and dACC

EVC and dACC

- Theory explains:
  - diversity of dACC responses:

EVC and dACC

- Theory explains:
  - diversity of dACC responses:
    - reward processing

EVC and dACC

- Theory explains:
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)

EVC and dACC

- **Theory explains:**
  - **diversity of dACC responses:**
    - **reward processing**
    - **aversive events (e.g., pain)**
    - **performance monitoring**

EVC and dACC

- **Theory explains:**
  - **diversity of dACC responses:**
    - **reward processing**
    - **aversive events (e.g., pain)**
    - **performance monitoring**
    - **explicit feedback**

EVC and dACC

- Theory explains:
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)
    - performance monitoring
    - explicit feedback
    - anxiety

EVC and dACC

- Theory explains:
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)
    - performance monitoring
    - explicit feedback
    - anxiety
    - adjustments in control

EVC and dACC

- Theory explains:
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)
    - performance monitoring
    - explicit feedback
    - anxiety
    - adjustments in control
    - action selection

EVC and dACC

- **Theory explains:**
  - **diversity of dACC responses:**
    - **reward processing**
    - **aversive events (e.g., pain)**
    - **performance monitoring**
    - **explicit feedback**
    - **anxiety**
    - **adjustments in control**
    - **action selection**
  - **close association of dACC activity with:**

EVC and dACC

- **Theory explains:**
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)
    - performance monitoring
    - explicit feedback
    - anxiety
    - adjustments in control
    - action selection
  - close association of dACC activity with:
    - valuation areas (vmPFC, OFC, nucleus accumbens)

EVC and dACC

- Theory explains:
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)
    - performance monitoring
    - explicit feedback
    - anxiety
    - adjustments in control
    - action selection
  - close association of dACC activity with:
    - valuation areas (vmPFC, OFC, nucleus accumbens)
    - control areas (dlPFC)

EVC and dACC

- **Theory explains:**
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)
    - performance monitoring
    - explicit feedback
    - anxiety
    - adjustments in control
    - action selection
  - close association of dACC activity with:
    - valuation areas (vmPFC, OFC, nucleus accumbens)
    - control areas (dlPFC)
    - neuromodulatory systems (VTA/DA, LC/NE, STN)

EVC and dACC

- Theory explains:
  - diversity of dACC responses:
    - reward processing
    - aversive events (e.g., pain)
    - performance monitoring
    - explicit feedback
    - anxiety
    - adjustments in control
    - action selection
  - close association of dACC activity with:
    - valuation areas (vmPFC, OFC, nucleus accumbens)
    - control areas (dlPFC)
    - neuromodulatory systems (VTA/DA, LC/NE, STN)
- Phenomenological experience…

Phenomenological Experience

Phenomenological Experience

- Electrical stimulation of dACC:
  - patients recounted a sense of “challenge” or “worry” (also known as foreboding) but remained motivated and aware that they would overcome the challenge

Phenomenological Experience

- Electrical stimulation of dACC:
  - patients recounted a sense of “challenge” or “worry” (also known as foreboding) but remained motivated and aware that they would overcome the challenge
  - feeling worried and anxious about something negative that was going to happen, but simultaneously knowing that he had to fight to make it through and not give up.

Phenomenological Experience

- Electrical stimulation of dACC:
  - patients recounted a sense of "challenge" or "worry" (also known as foreboding) but remained motivated and aware that they would overcome the challenge
  - feeling worried and anxious about something negative that was going to happen, but simultaneously knowing that he had to fight to make it through and not give up.
  - "like you're headed towards a storm that's on the other side, maybe a couple of miles away, and you've got to get across the hill and all of a sudden you're sitting there going how am I going to get over that… you have to keep going forward"

General Comments

General Comments

- ***Not*** **a single agent model:**
  - **may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system**

General Comments

- ***Not*** **a single agent model:**
  - **may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system**

  - **but:**
    - **no single, unified objective function, rather many different ones**
    - **individual components may respond differently in different environments**
    - **components may adapt at different rates to changes in the environment**

General Comments

- ***Not*** **a single agent model:**
  - **may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system**

  - **but:**
    - **no single, unified objective function, rather many different ones**
    - **individual components may respond differently in different environments**
    - **components may adapt at different rates to changes in the environment**

- **More easily described as:**
  - **amalgam of controlled + automatic processes ("agents")**

General Comments

- ***Not*** **a single agent model:**
  - may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system

  - but:
    - ♦ no single, unified objective function, rather many different ones
    - ♦ individual components may respond differently in different environments
    - ♦ components may adapt at different rates to changes in the environment

- **More easily described as:**
  - amalgam of controlled + automatic processes ("agents")
  - each fully rational with respect to its parameters (e.g., objective function for maximizing reward rate, time scales of adaptation)

General Comments

- *Not* a single agent model:
  - may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system

  - but:
    - no single, unified objective function, rather many different ones
    - individual components may respond differently in different environments
    - components may adapt at different rates to changes in the environment

- More easily described as:
  - amalgam of controlled + automatic processes ("agents")
  - each fully rational with respect to its parameters (e.g., objective function for maximizing reward rate, time scales of adaptation)
  - compete with each other and are selected among: *(Dulberg et al., 2024)*
    - using heuristic approximations to optimal computations
    - based on intrinsic biases, experience and exigencies of environment

General Comments

- ***Not*** **a single agent model:**
  - **may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system**

  - **but:**
    - **no single, unified objective function, rather many different ones**
    - **individual components may respond differently in different environments**
    - **components may adapt at different rates to changes in the environment**

- **More easily described as:**
  - **amalgam of controlled + automatic processes ("agents")**
  - **each fully rational with respect to its parameters (e.g., objective function for maximizing reward rate, time scales of adaptation)**
  - **compete with each other and are selected among:** *(Dulberg et al., 2024)*
    - **using heuristic approximations to optimal computations**
    - **based on intrinsic biases, experience and exigencies of environment**

Personality

General Comments

- ***Not*** **a single agent model:**
  - may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system

  - but:
    - no single, unified objective function, rather many different ones
    - individual components may respond differently in different environments
    - components may adapt at different rates to changes in the environment

- **More easily described as:**
  - amalgam of controlled + automatic processes ("agents")
  - each fully rational with respect to its parameters (e.g., objective function for maximizing reward rate, time scales of adaptation)
  - compete with each other and are selected among: *(Dulberg et al., 2024)*
    - using heuristic approximations to optimal computations
    - based on intrinsic biases, experience and exigencies of environment

⇒aggregate behavior is not be reducible (easily, at all?) to a simple expression

General Comments

- ***Not*** **a single agent model:**
  - may (ultimately) be able to write down a set of equations that (nearly) fully describe (almost) the entire system

  - but:
    - no single, unified objective function, rather many different ones
    - individual components may respond differently in different environments
    - components may **adapt at different rates** to changes in the environment

- **More easily described as:**
  - amalgam of controlled + automatic processes (“agents”)
  - each fully rational with respect to its parameters (e.g., objective function for maximizing reward rate, time scales of adaptation)
  - compete with each other and are selected among: *(Dulberg et al., 2024)*
    - using heuristic approximations to optimal computations
    - based on intrinsic biases, experience and exigencies of environment

⇒aggregate behavior is not be reducible (easily, at all?) to a simple expression

Longterm Adaptation

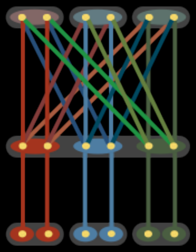

Longterm Adaptation

- Control-dependent processing:
  - Cost: *serialization* of processing

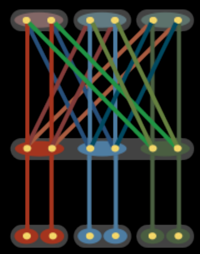

Longterm Adaptation

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation

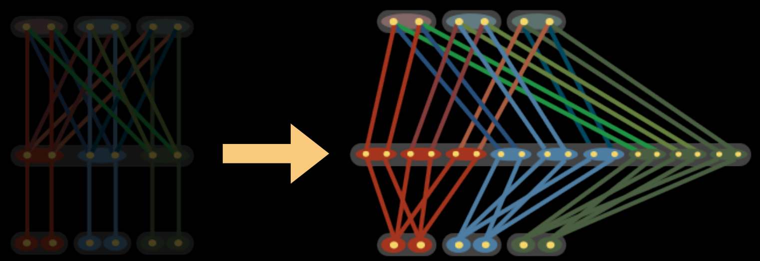

Longterm Adaptation

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Poses a higher level problem:
    - *intertemporal choice,* when to…

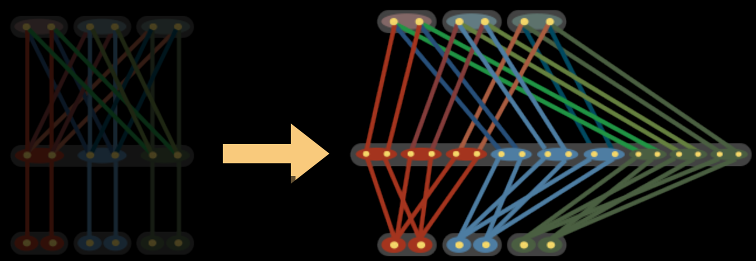

Longterm Adaptation

- **Control-dependent processing:**
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Poses a higher level problem:
    - *intertemporal choice,* when to…
      - ♦ reap *immediate rewards* at cost of serial processing

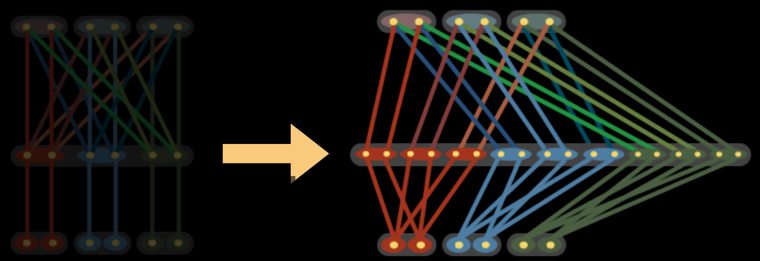

Longterm Adaptation

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Poses a higher level problem:
    - *intertemporal choice,* when to…
      - reap *immediate rewards* at cost of serial processing
      - invest in *longer term reward* of
        more efficient processing *(multitasking)*
        at cost of time and effort required for automatization

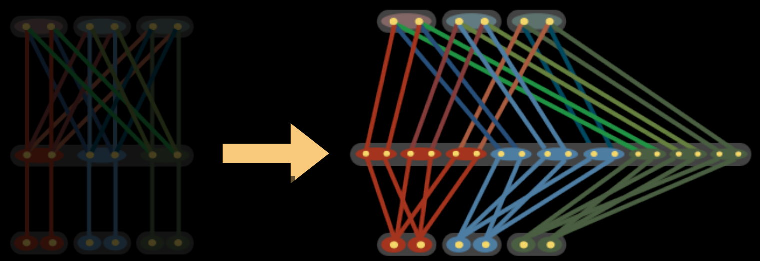

Longterm Adaptation

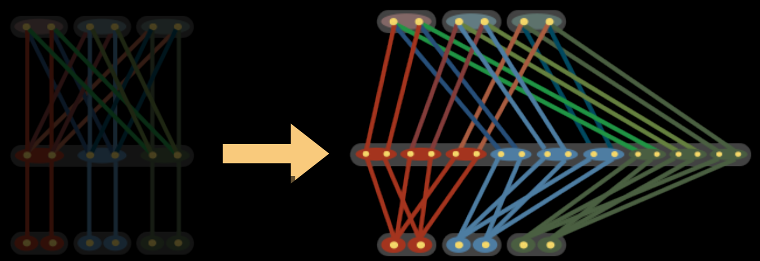

- **Control-dependent processing:**
  - Cost: ***serialization*** **of processing**
  - **Overcome through** ***automatization:***
    - **learn separated, dedicated task representation**
  - **Poses a higher level problem:**
    - ***intertemporal choice,*** **when to…**
      - ♦ **reap** ***immediate rewards*** **at cost of serial processing**
      - ♦ **invest in** ***longer term reward*** **of**
        more efficient processing *(multitasking)*
        at cost of time and effort required for automatization
  - **This too can be approached** ***normatively:***

Longterm Adaptation

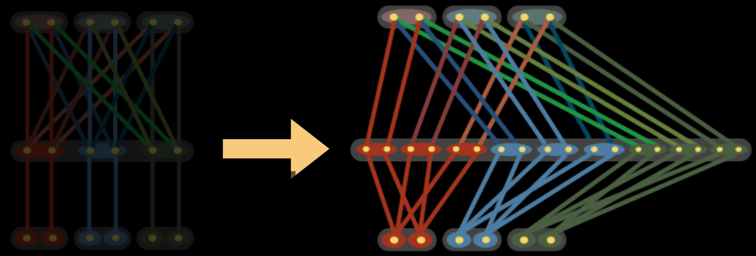

- **Control-dependent processing:**
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Poses a higher level problem:
    - *intertemporal choice,* when to…
      - reap *immediate rewards* at cost of serial processing
      - invest in *longer term reward* of
more efficient processing *(multitasking)*
at cost of time and effort required for automatization
  - This too can be approached *normatively:*
    - *rational self-reconfiguration….*

Rational Self-Reconfiguration

- **Formal analysis of learning speed vs. processing efficiency**
  (Musslick, Saxe et al., 2017)

$$(training\ time)^2\ \alpha\ \frac{(multitasking\ capacity)}{(\%shared\ representations)}$$

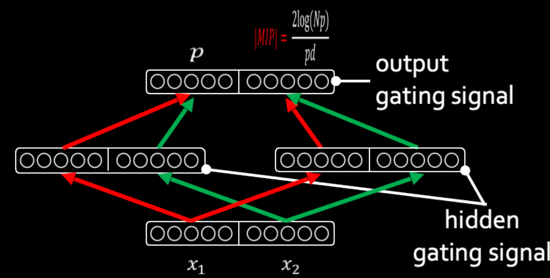

Rational Self-Reconfiguration

- **Formal analysis of learning speed vs. processing efficiency**
  (Musslick, Saxe et al., 2017)

$$(training\ time)^2\ \alpha\ \frac{(multitasking\ capacity)}{(\%shared\ representations)}$$

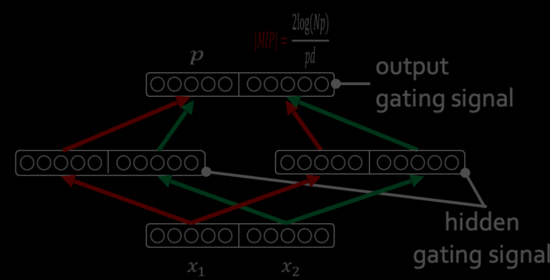

- **Bayesian optimal process model**
  (Sagiv, Musslick & Cohen, 2018)

$$\mathbb{E}_B[R|t] = \sum_{i=1}^{\min\{N,K\}} \mathbb{P}(\alpha = i) \sum_{j=0}^{i-1} \mathbb{P}_B(\text{success on task } j)(1 - jC)$$

$$\mathbb{E}_T[R|t] = \sum_{i=1}^{\min\{N,K\}} \mathbb{P}(\alpha = i) \sum_{j=0}^{i-1} \mathbb{P}_T(\text{success on task } j)(1)$$

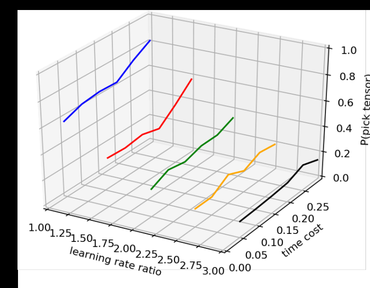

Rational Self-Reconfiguration

- **Formal analysis of learning speed vs. processing efficiency**
  (Musslick, Saxe et al., 2017)

$$(training\ time)^2\ \alpha\ \frac{(multitasking\ capacity)}{(\%shared\ representations)}$$

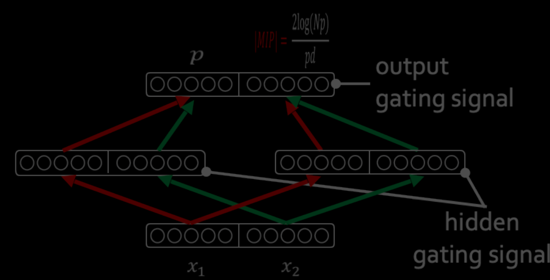

- **Bayesian optimal process model**
  (Sagiv, Musslick & Cohen, 2018)

$$\mathbb{E}_B[R|t] = \sum_{i=1}^{\min\{N,K\}} \mathbb{P}(\alpha = i) \sum_{j=0}^{i-1} \mathbb{P}_B(\text{success on task } j)(1 - jC)$$

$$\mathbb{E}_T[R|t] = \sum_{i=1}^{\min\{N,K\}} \mathbb{P}(\alpha = i) \sum_{j=0}^{i-1} \mathbb{P}_T(\text{success on task } j)(1)$$

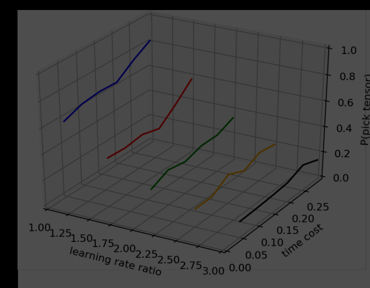

- **Deep learning applications**
  (Ravi, Musslick & Cohen, under review)

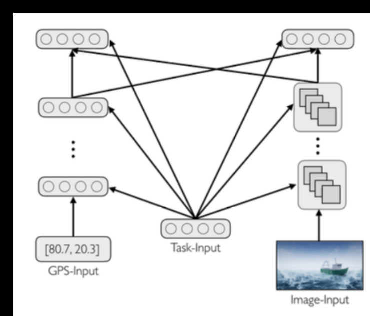

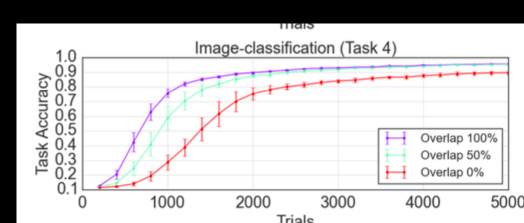

Longterm Adaptation

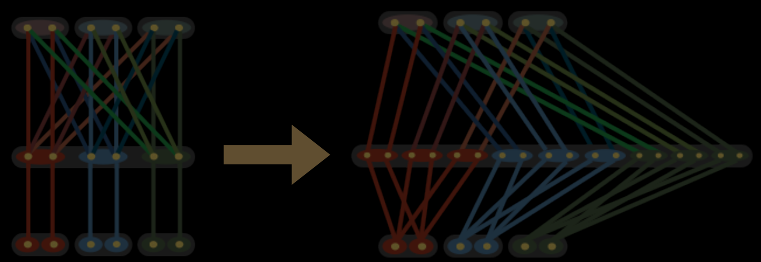

Longterm Adaptation

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Focused on simple stimulus-response mapping tasks

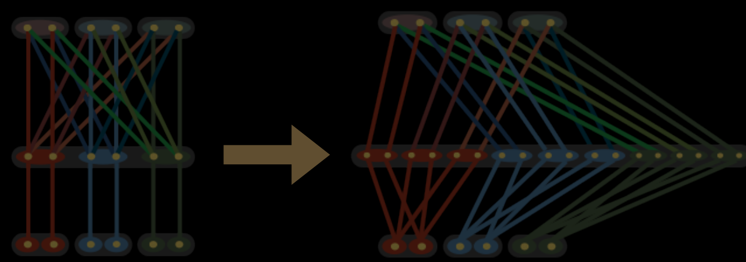

Longterm Adaptation

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Focused on simple stimulus-response mapping tasks
    - What about tasks involving *sequences* of states and actions?

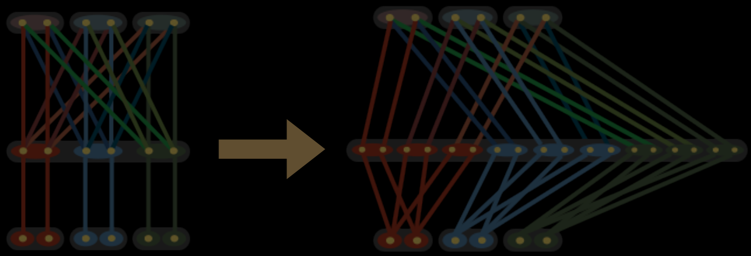

Longterm Adaptation

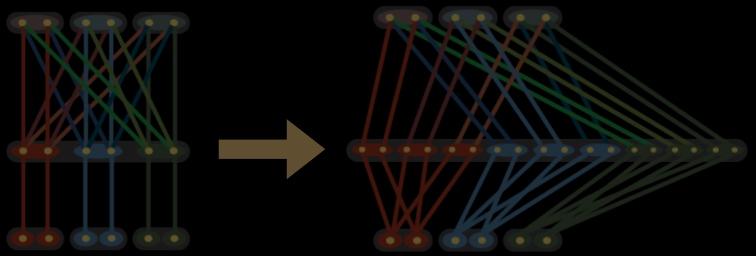

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Focused on simple stimulus-response mapping tasks
    - What about tasks involving *sequences* of states and actions?
    - Poses same problem of *intertemporal choice,* when to…
      - reap immediate rewards of flexible but *serial processing*
      - invest in longer term reward *compilation*

Longterm Adaptation

- **Control-dependent processing:**
  - Cost: *serialization* of processing
  - Overcome through *automatization:*

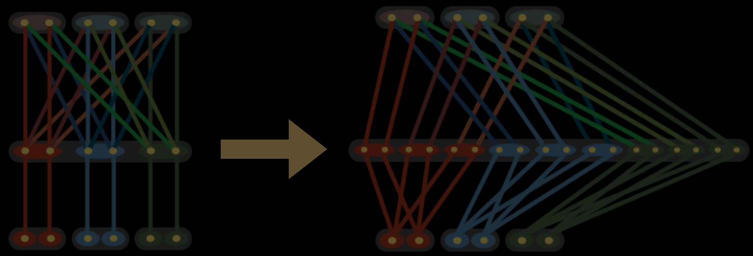

    - learn separated, dedicated task representation
  - Focused on simple stimulus-response mapping tasks
    - What about tasks involving *sequences* of states and actions?
    - Poses same problem of *intertemporal choice,* when to…
      - reap immediate rewards of flexible but *serial processing*
      - invest in longer term reward *compilation*
  - Now, though, the representational problem is in *time* rather than *"space"*

Longterm Adaptation

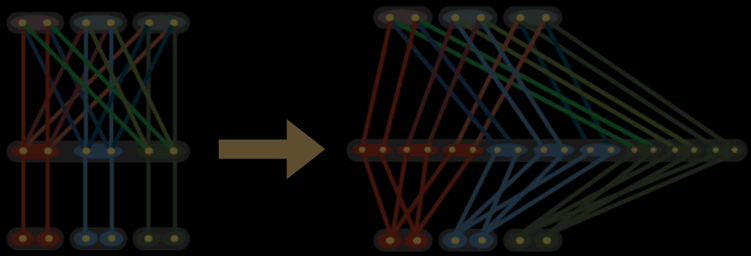

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Focused on simple stimulus-response mapping tasks
    - What about tasks involving *sequences* of states and actions?
    - Poses same problem of *intertemporal choice,* when to…
      - reap immediate rewards of flexible but *serial processing*
      - invest in longer term reward *compilation*
  - Now, though, the representational problem is in *time* rather than *"space"*
  - This too can be approached normatively with rational reconfiguration:

Longterm Adaptation

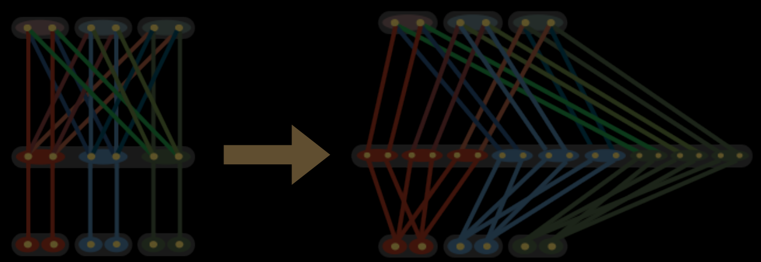

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation
  - Focused on simple stimulus-response mapping tasks
    - What about tasks involving *sequences* of states and actions?
    - Poses same problem of *intertemporal choice,* when to…
      - reap immediate rewards of flexible but *serial processing*
      - invest in longer term reward *compilation*
  - Now, though, the representational problem is in *time* rather than *"space"*
  - This too can be approached normatively with rational reconfiguration:
    - **use of EM to "flatten" time…**

Longterm Adaptation

- Control-dependent processing:
  - Cost: *serialization* of processing
  - Overcome through *automatization:*
    - learn separated, dedicated task representation

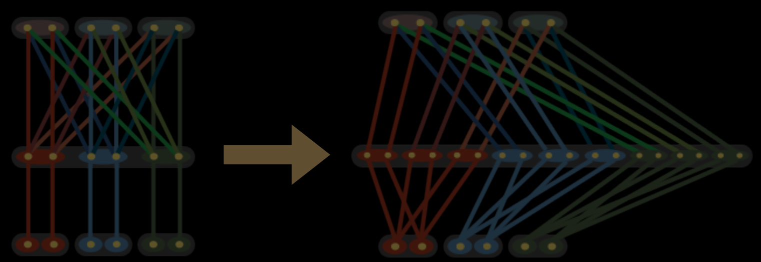

  - Focused on simple stimulus-response mapping tasks
    - What about tasks involving *sequences* of states and actions?
    - Poses same problem of *intertemporal choice,* when to…
      - reap immediate rewards of flexible but *serial processing*
      - invest in longer term reward *compilation*
  - Now, though, the representational problem is in *time* rather than *"space"*
  - This too can be approached normatively with rational reconfiguration:
    - **use of EM to "flatten" time… EGO and ISC-CI models!**