

COSC451: Artificial Intelligence

Lecture 3: Vision for action: the reach and grasp motor pathways

Alistair Knott

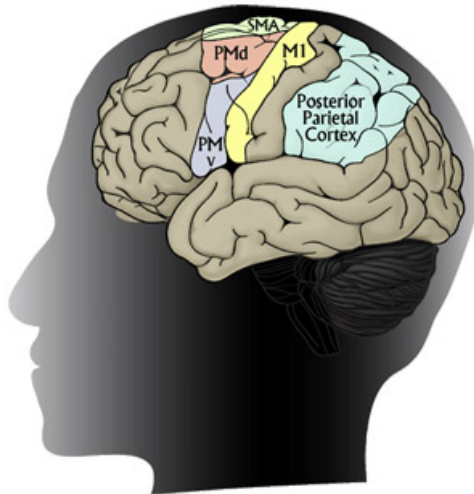
Dept. of Computer Science, University of Otago

Outline of the talk

- 1 Primary motor cortex
- 2 The 'reach' visuomotor pathway
- 3 The 'grasp' pathway
- 4 Endpoint of a grasp action: a stable grasp

Brain regions we'll be looking at today

The main areas (in a human brain)



Outline of today's lecture

- 1 Primary motor cortex
- 2 The 'reach' visuomotor pathway
- 3 The 'grasp' pathway
- 4 Endpoint of a grasp action: a stable grasp

Low-level motor control

Muscles are activated by **motor neurons**.

Motor neurons are activated by neurons/circuits in the **spinal cord**.

- These coordinate *groups* of muscles.
- There are many low-level motor circuits which operate without any contribution from the brain at all. (E.g. the knee-jerk (Patella) reflex).


Neurons in the spinal cord receive top-down inputs from the brain.

- Most of these come from the **primary motor cortex**.

Primary motor cortex (M1/F1)

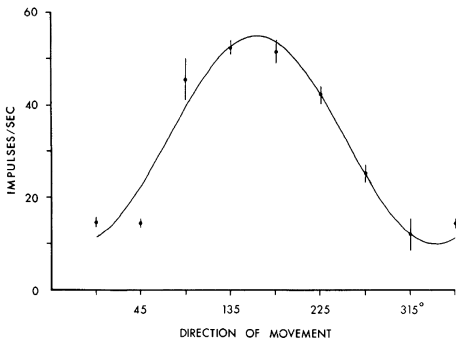
Primary motor cortex is **somatotopically** organised.

- Moving specific body parts activates specific parts of F1.
- Stimulating specific parts of F1 triggers specific body movements.



Population coding in primary motor cortex

Recording from individual neurons in F1 shows that they are broadly tuned to movement directions (Georgopoulos *et al.*, 1982).



But information from a *population* of neurons with different tuning curves can identify the direction of movement quite precisely.

Reaching to grasp: an overview

A reach-to-grasp movement has two separate components:

- A **reach** movement (involving the shoulder and elbow, and probably the upper body as well)
- A **grasp** movement (involving the fingers of the hand).

Obviously these need to be closely coordinated. But nonetheless, there seem to be two quite separate pathways involved in generating the two movements.

- The reach pathway is driven by representations of *object location*.
- The grasp pathway is driven by representations of *object shape*.

Each pathway proceeds through its own areas of parietal cortex → premotor cortex → motor cortex.

Visuomotor pathways

Any visuomotor pathway has to do three things:

1. Convert visual information into **movement goals**.
 - Movement goals are expressed in a **motor coordinate system**. (Locations are specified as *motor states*.)
So a complicated *coordinate system transformation* is involved.
 - An object can often be acted on in different ways. So a single object might evoke several different movement goals. These are called the **motor affordances** of the object.

Visuomotor pathways

Any visuomotor pathway has to do three things:

2. **Select** a movement from the available alternatives.
 - A movement will be chosen partly bottom-up, based on its *possibility*, and partly top-down, based on its *desirability*.

Visuomotor pathways

Any visuomotor pathway has to do three things:

3. **Control** a movement when it is under way.

- A movement is extended in time.
- We begin by generating a motor impulse. This moves the effector.
- At each point thereafter, we monitor where the effector is, and send a revised motor signal.
- When the movement goal is obtained, we stop.

Outline of today's lecture

- 1 Primary motor cortex
- 2 The 'reach' visuomotor pathway**
- 3 The 'grasp' pathway
- 4 Endpoint of a grasp action: a stable grasp

The reach pathway

The goal of the reach pathway is to select a target in the agent's **peripersonal space**, and bring the agent's hand to this target.

To discuss selection and execution of reach actions, I will use the concept of a **movement vector**.

- The movement vector is given in motor coordinates.
- It is a vector between the *current hand position* and the *goal hand position*.
- F1 motor signals (coarse-coded) can be generated from this vector.
(They'll be generated in the shoulder/elbow/torso areas of F1.)

Overview: selection of a reach target

To begin with, several movement vectors are computed in parallel, for several objects in the agent's visual field.

- We can imagine that the saliency map supports this parallel calculation.

These movement vectors *compete* with one another.

- Behavioural evidence in humans (Tipper *et al.*, 1998): when we reach for a target object in the presence of a distractor object, we do plan a movement vector for the distractor, and then we inhibit this movement vector.

The winning movement vector is prepared, pending an internal 'GO' signal.

Overview: execution of a reach to the selected target

When the action is initiated, the movement vector takes on a dynamic role in determining the **trajectory** of the hand as the action unfolds.

- The trajectory of the hand is a specification of its **motor state** (position plus velocity) at each point during the action.

Computer motor planners tend to plan trajectories in advance, offline. But humans/primates probably generate trajectories 'on the fly':

- A movement vector is derived *in real time* from the current location of the hand and of the target.
- The movement vector causes a motor impulse, which moves the hand, and a new movement vector is derived. (Etc.)

Overview: the motor controller function

We can think of the circuit which computes the movement vector as implementing a **motor controller** function.

- At each moment, the function takes a movement vector, and computes a motor impulse.

It probably implements a mixture of **feedback** and **feedforward control** (see e.g. Kawato and Gomi, 1992):

- Feedback control is solely based on the movement vector.
- Feedforward control incorporates a learned **inverse model** of the dynamics of the arm.

Overview: forward models in motor control

During execution of a reach action, we need to know the hand's motor state.

- We can get hand *location* from the senses: stretch receptors in the muscles compute joint angles.
- However, it takes a while to get this **reafferent** sensory information. And it's not very accurate. And it doesn't give us hand *velocity* at all.

To get faster, completer, more accurate information about motor state, we use a learned **forward model** of the dynamics of the arm.

- This *predicts* the next motor state from the current motor state and the current motor signal.
(It uses an **efferent copy** of the motor signal sent to F1.)

Neural areas involved in reaching

Premotor areas F2 and F4 are involved in computing movement vectors.

- Cells in F2 and F4 compute a representation of the agent's peripersonal space (see Colby and Goldberg, 1999).
- F4 does head and body (Fogassi *et al.*, 1996); F2 does mainly body (Fogassi *et al.*, 1999).
- Location is coded in an **effector-centred** coordinate system. (So it changes as the effector moves.)
- Some cells respond only to visual stimuli. Some respond only to *tactile* stimuli. Some respond only to a combination of visual and tactile.

F2 and F4 both project to F1 (primary motor cortex).

Inputs to F2 and F4

F2 and F4 are *premotor* areas; they receive input from separate *posterior parietal* areas (see Geyer *et al.*, 2000).

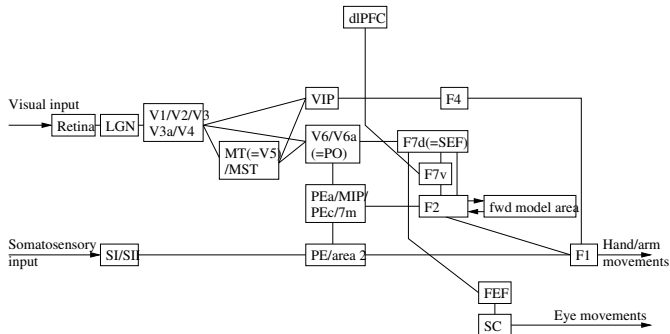
- F4 receives input from the ventral intraparietal cortex (VIP), whose neurons are similar to F4 (though with more bimodal neurons).
- The VIP-F4 circuit seems specialised for avoiding objects moving towards the head, or biting things which are close to the head (Colby *et al.*, 1993; Cooke *et al.*, 2003).
- The circuit involving F2 is more complex, and seems more specialised for voluntary grasping movements. I'll focus on this circuit.

Inputs to the F2 circuit

There are several inputs to the circuit leading to F2.

- Visual information arrives from V1-V4, and from MT/MST. This information is combined in a parietal/occipital area called **PO**.
- **Somatic** information (tactile and body position information) arrives from somatosensory areas **SI and SII**. It arrives in parietal areas **PE** and **area 2**.
- Visual and somatic information is combined in a set of adjacent parietal areas I will call the **PEa complex**.
- F2 also receives inputs from **prefrontal cortex (PFC)**, via frontal area **F7v**.

The reach sensorimotor pathway



Sensorimotor transformations in the reach pathway

Burnod *et al.* (1999) propose that there are two dimensions of variation in the reach pathway.

- A **visual-to-somatic** dimension. (Top-down in my figure)
- A **sensory-motor** dimension (Left-to-right in my figure).
 - Cells in sensory regions are time-locked to sensory signals.
 - Cells in F1 are time-locked to motor movements.
 - Cells in intermediate areas can be of either type.

Intermediate areas also contain two special types of cell.

- **Matching cells** encode mappings between sensory inputs and motor commands, learned through *Hebbian association*.
- **Condition cells** encode mappings between sensory stimuli and motor commands, learned through *reinforcement*.

Matching cells and condition cells

Matching cells are particularly common in the **PEa complex**.

- They appear to encode *possible* motor commands.

Condition cells are particularly common in **F7v**.

- Recall that this is the area via which F2 is linked to PFC.
- Condition cells appear to encode *desirable*, or *prepared* motor commands.

F2 can be thought of as a place where alternative motor commands *compete* amongst one another.

- Tipper *et al.* (1992): humans compute several alternative movement vectors, which compete amongst each other.
- Cisek and Kalaska (2005): F2 neurons encode two possible reach actions before the animal has chosen between them.

Integration of motor and attentional actions

Reach actions have an *attentional component*: agents **saccade** to the target object at the start of a movement (Johansson *et al.*, 2001).

While the eye is **tracking** the target,

- the position of the eye tells us about **target location**;
- the shape of the target can be computed in detail.

F2 links to an area called **F7d**, which is involved in generating eye movements which are coupled to actions.

- Attention-for-action and attention-for-categorisation have common neural substrates (Schneider and Deubel, 2002).

Outline of today's lecture

- 1 Primary motor cortex
- 2 The 'reach' visuomotor pathway
- 3 The 'grasp' pathway**
- 4 Endpoint of a grasp action: a stable grasp

The grasp pathway

The grasp pathway can be thought of as a motor control circuit which runs in parallel to the reach pathway.

- The visual input in this case relates to the *shape* of a selected object, and its *orientation*.
- The output is a sequence of motor impulses to the hand and wrist.

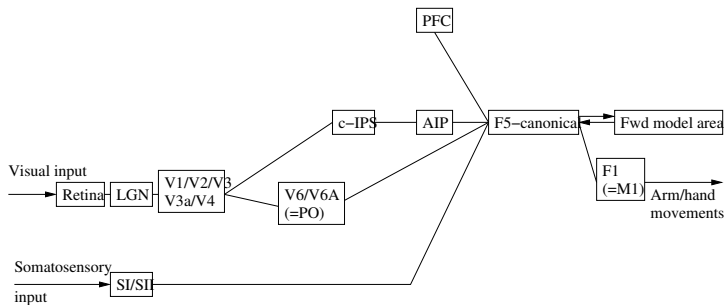
But while the reach circuit is involved in the process of selecting a target, the grasp circuit only becomes active once a target has been chosen.

Stages of the grasp pathway

The grasp pathway receives visual and somatosensory inputs.

- The main visual input comes from the **caudal IPS**, which computes the 3D shape of an attended object using **stereopsis** (Sakata *et al.*, 1998).
- c-IPS projects to **AIP**, which in turn projects to premotor area **F5**, which projects to F1.
- F5 also receives somatosensory inputs (from SI/SII).
- F5 also receives inputs from PFC.

The grasp sensorimotor pathway



AIP and F5 cells

A typical AIP cell shows sensitivity to one particular type of grasp, (e.g. 'precision pinch', 'power grasp' (Taira *et al.*, 1990).

- **Motor-dominant** cells respond to their associated grasp when it is *executed*.
- **Visual** cells respond to their associated grasp when the animal attends to an object which *affords* this grasp.

F5 'canonical' neurons similarly respond to particular associated grasps.

In Fagg and Arbib's (1998) model:

- AIP computes a range of possible grasps; F5 selects one of these.
- AIP then sends a *sequence* of commands to F5, specifying the different stages of the grasp action.

Outline of today's lecture

- 1 Primary motor cortex
- 2 The 'reach' visuomotor pathway
- 3 The 'grasp' pathway
- 4 Endpoint of a grasp action: a stable grasp**

Endpoint of a grasp action: the haptic interface

At the end of a (successful) grasp action, the agent is holding the target object in a stable grasp.

There are two ways of characterising this state.

- First, as with any action, the agent has brought about a change in the state of the world.
- Second, the agent has executed a special kind of **attentional action** to the target object.
 - The *location* of the target is now given by the position of the agent's hand.
 - The *type* of the object is now given by the position of the agent's fingers (see e.g. Goodwin and Wheat, 2004).

The grasp state provides an opportunity to **integrate** the haptic and visual attentional modalities (Driver and Spence, 1998).

Summary

- The agent computes a saliency map of all the objects in his visual field. (In LIP.)
- The agent computes a set of movement vectors in PO/PEa, representing the reachable objects.
- In F2, a movement vector is selected, based on bottom-up info from PO/PEa and top-down info from PFC/F7d.
- This biases *visual* attention towards the selected target, resulting in an improved representation of its shape.
- In the grasp pathway, AIP computes a range of possible grasps.
- F5 selects one of these (using input from PFC).
- The reach and grasp components of the movement are executed.
- A stable grasp is obtained, in which the agent can integrate visual and haptic object representations.