Statistical Techniques in Robotics (16-831, F14) Lecture#7 (Tuesday September 16)

Undirected Graphical Models

Lecturer: Drew Bagnell *Scribes: ecappo, nrhineha* [1]

1 Review of Graphical Models: Bayesian Network

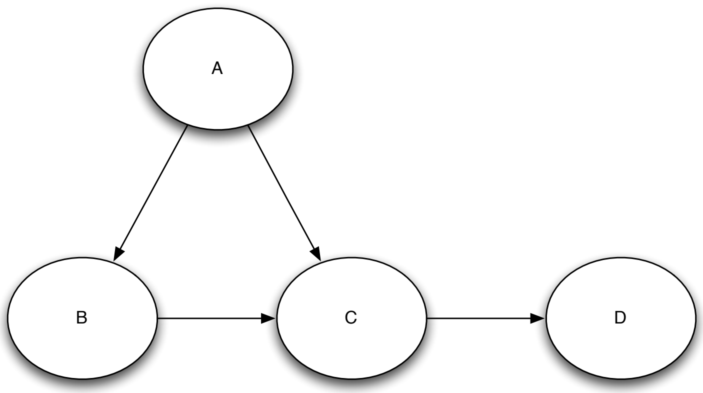

Figure 1: Bayesian Network

The joint probability of the graph shown in Figure 1 is

$$P(A, B, C, D) = P(A)P(B|A)P(C|A, B)P(D|C)$$

A Bayes' net is an example of a common graphical model. Bayes' nets are also known as belief networks, directed graphical models (DAGS), and Bayesian networks. The nodes represent uncertain quantities, and edges represent relationships between them. Frequently, these edges are written to indicate causal dependence between several variables, although this is not always the case.

The simplest Bayes' net is one in which all nodes are *completely disconnected*. Considering the case of removing the edges from the node structure of Figure 1, the joint probability can be decomposed as follows:

$$P(A, B, C, D) = P(A)P(B)P(C)P(D)$$

Another simple Bayes' net is a *Markov Chain*, e.g.: $X_1 \to X_2 \to X_3 \to X_4$. The Markov Chain is a commonly seen example of a Bayesian network; for example, it can be used to describe a natural progression in time, such as the evolution of a robot's pose. Or even the weather—for example, if X's are days, then this Markov Chain describes that the weather on Day 3 is based only on the weather the day before (Day 2), but is not dependent on Day 1 if weather on Day 2 is known.

Graphical models are useful for determining (conditional) independencies. Graphical models easily show that given a node's immediate parent(s), that node is independent of all "grand-parent"

[1]Some content adapted from previous scribes: Adam Stambler

nodes. In terms of time, we are basically saying that the present removes the connection between the past and the future. For example, in our Weather Markov Chain in the paragraph above, Day 3 is the child of Day 2 and therefore depends only on Day 2 and not on Day 2's parent, Day 1. (There is no direct arrow from Day 1 to Day 3.) Day 3 is therefore independent from Day 1 given Day 2. I.e., we can say that $x_3 \perp x_1 | x_2$.

Now, if we said that $P(x_1, x_2, x_3) = \prod_i P(x_i)$, is $x_1 \rightarrow x_2 \rightarrow x_3$ a valid answer? The joint probability of $x_1 \rightarrow x_2 \rightarrow x_3$ is written out as $P(x_1)P(x_2|x_1)P(x_3|x_2)$. So while this Markov Chain is a valid expression of $P(x_1, x_2, x_3)$, the arrows imply conditionality, and our answer is therefore *more* general than the original statement—the original statement told us that $P(x_1)$ was independent from $P(x_2)$, whereas our statement did not let us make such a specific conclusion. In Bayesian Networks, more arrows means less structure (fewer arrows implies more structure).

Graphical models help us answer questions like, "In Figure 1, is B independent of C given A? Or is C independent of B given A?"

2 Bayesian net rules for determining conditional independence

It is useful to determine whether or not an arbitrary node is conditionally independent of another node given a set of "blocked" nodes, or nodes for which the outcome has been observed. If, for all paths between two nodes, every path is "blocked", then one can infer that, based on the graph, the quantities are conditionally independent of the given set. To determine blockage between one node and neighbors of its neighbors, one must evaluate the configuration of the nodes and the given information. The configuration falls into one of 3 main categories, and 6 total categories.

Let's condition on C. Looking at Figure 1, is there any path that goes from B to D that isn't "blocked?" (More on what "blocked" means in a minute...) If all paths between variables of interest B, D are blocked, then $B \perp D | C$.

2.1

Rule 1: The Markov Chain Rule—observing x_2 blocks the path

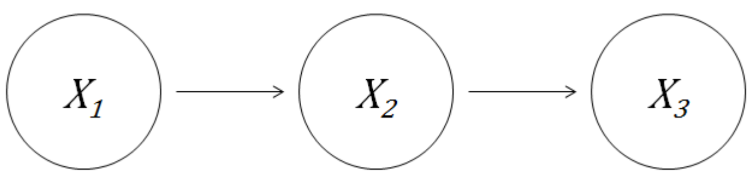

Figure 2: A Markov Chain.

If anywhere along the path you have $x_1 \rightarrow x_2 \rightarrow x_3$ and x_2 is known, then x_3 and x_1 are conditionally independent. We can say that x_2 *blocks* the path from $x_1 \rightarrow x_3$ or $x_3 \rightarrow x_1$.

2.2

Rule 2: The Converging Arrow Rule—not observing x_2 blocks the path

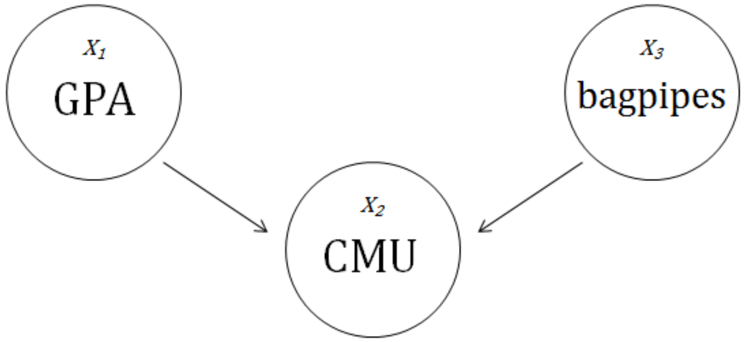

Figure 3: How a student might get into CMU, illustrating the converging arrows rule.

Suppose there are two ways for a student to get into CMU. They may have a high GPA or they may have a high "Kiltie Score" (i.e., be very skilled at playing the bagpipes). As there are no arrows drawn between GPA and bagpipes, Figure 3 implies that having a high GPA and having a high Kiltie score are independent. (Which could be plausible... if a student has finite time, they only have time to practice one skill.)

If we do not know the value of CMU—if CMU is unobserved—we must say that "GPA" is independent of "bagpipes": $GPA \perp bagpipes$.

Therefore, not observing CMU blocks the path. We must observe CMU in order to make any correlations between "GPA" and "bagpipes." For example, if we do know that the student got into CMU, then our belief in the student having a high GPA or a high Kiltiescore goes up. (Incidentally, if GPA were observable and we saw the student had a high GPA, then our belief in the student getting accepted to CMU goes up and our belief that the student has a high Kiltie score goes down.)

Further, the subtleties of converging arrows means you have to look at descendnts of the convergent node. **If the convergent node *and all children* are unobserved, then the path is blocked**. See Figure 4 for an example of this. If we observe that a police car drives up to our house, we can conclude that the alarm must have gone off. Therefore, we may as well have observed the alarm in the first place. Since the "alarm" node is the convergent node, we cannot conclude that $earthquake \perp burglar$. Because we observed a police car arrive at our house, we know the alarm went off and therefore we know there was either an earthquake or a burglar. This is why Rule 2 holds that a convergent node and all its children must be unobserved for the path to be blocked.

2.3

Rule 3: The Divergent Arrow Rule—observing x_2 blocks the path

In Figure 5, there are no arrows drawn between "police car" and "neighbor." This implies that these nodes are independent. However, they are correlated—if the alarm goes off, both nodes occur. The neighbor, irritated at the noise, calls to complain, and the police car responds to the alarm. We can also infer that the alarm must have gone off if we observe the police car pull up to the house, and we can therefore expect a call from the neighbor any instant.

If we observe the alarm occur, both the police car and the phone call are redundant; if we know the alarm went off, we don't need a nosy neighbor or the police telling us it went off. Therefore,

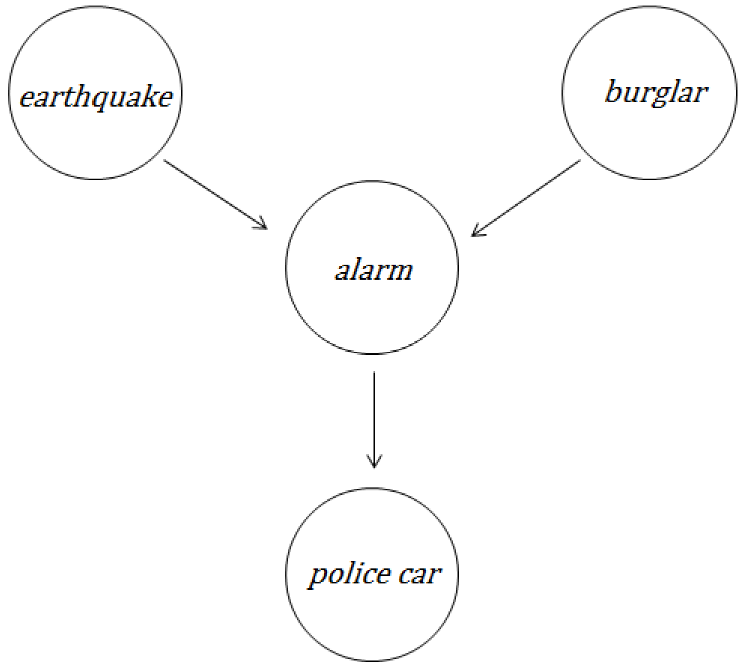

Figure 4: Judea Pearl's example illustrating the converging arrows rule with descendants.

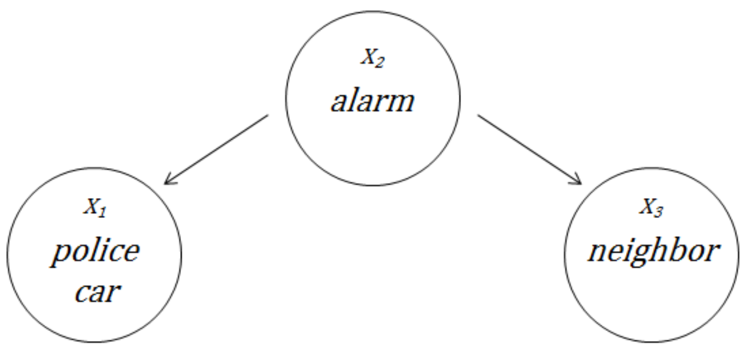

Figure 5: Illustration of diverging arrows.

Rule 3 states that if the divergent node is observed, the path is blocked.

3 Graphical Models of the Filtering Problem

Now let's use the rules we learned to look at graphical models for the filtering problem we have been looking at. As drawn in the book, the filtering problem looks like Figure 6, which is an open-loop representation of the controls. This model might represent an RC car, for example, which had pre-programmed controls.

The question here is whether this graphical model of the problem makes sense, and in our homework,

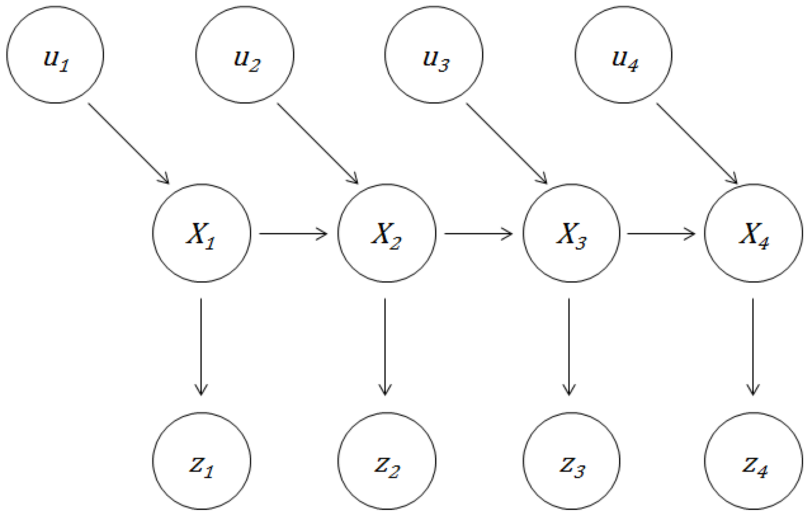

Figure 6: Filtering graphical model with open loop controls.

we worked to show the below statement:

$$P(x_3|z_{1:3}, u_{1:4}) \stackrel{?}{=} P(x_3|z_{1:3}, u_{1:3})$$

Are x_3 and u_4 conditionally independent given previous z's and u's? I.e., is $x_3 \perp u_4|z_{1:3}, u_{1:3}$? In Figure 6, the z's and u's are observed while the x's are not observed. Therefore, we can look at this as a case of Rule 2, a case of converging arrows, as shown in Figure 7. As x_4 (and its children) are unobserved, this chain is "blocked," and so $x_3 \perp u_4$.

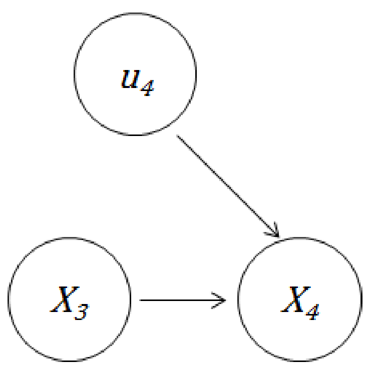

Figure 7: A "zoom-in" on a portion of the open loop control, filtering graphical model.

3.1 A second case of controls in the filtering model.

What if instead of having controls which were pre-programmed, we had controls which depended on the current state as shown in Figure 8? Would the derivation from the previous model still hold? Because x_3 is trivially connected to u_4—i.e., x_3 directly influences u_4—these states are no longer conditionally independent, and the previous derivation does NOT, in fact, hold.

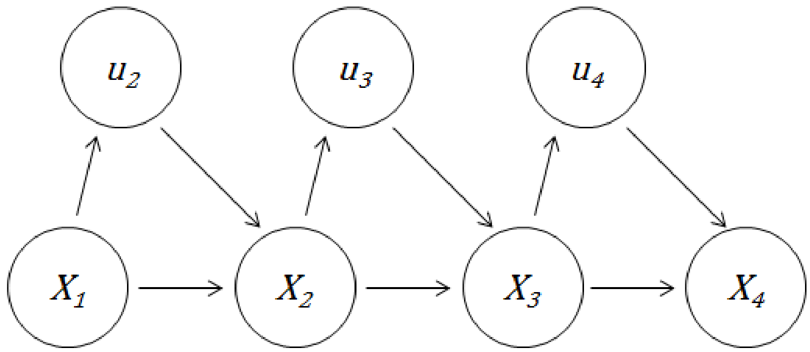

Figure 8: State dependent controls.

3.2 A case of controls based on observations.

This time, what if the controls were memory-less and dependent on observations, as shown in Figure 9? Would the derivation still hold? Yes! The derviation still holds. (z_3 is observed, which

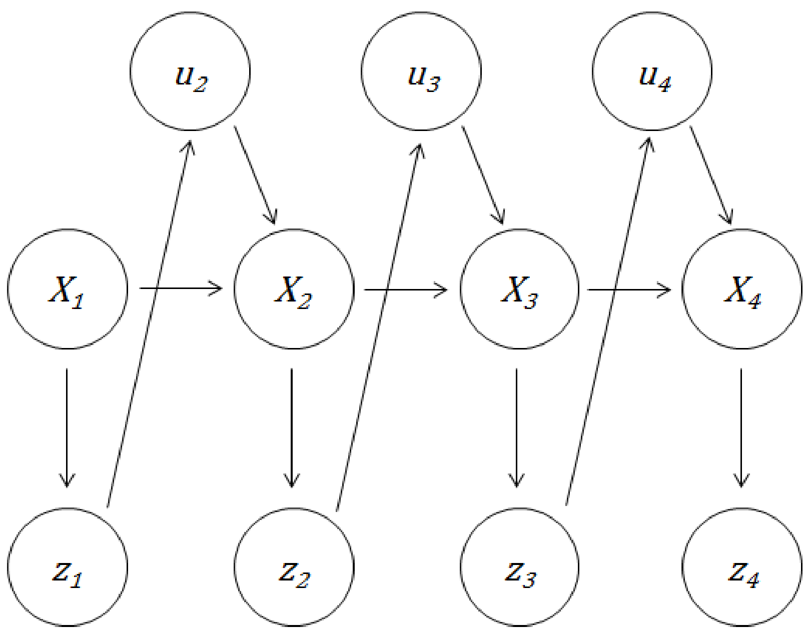

Figure 9: Measurement dependent controls.

blocks $x_3 \rightarrow z_3 \rightarrow u_4$, and x_4 and all children are unobserved).

3.3 Given a state estimation function...

In Figure 10 we show a state estimation function, f, which takes in observations and outputs controls. With this model, the derivation also still holds.

Given z_1 is observed, paths between x_1 and u_2 are:
$x_1 \rightarrow x_2 \rightarrow u_2$

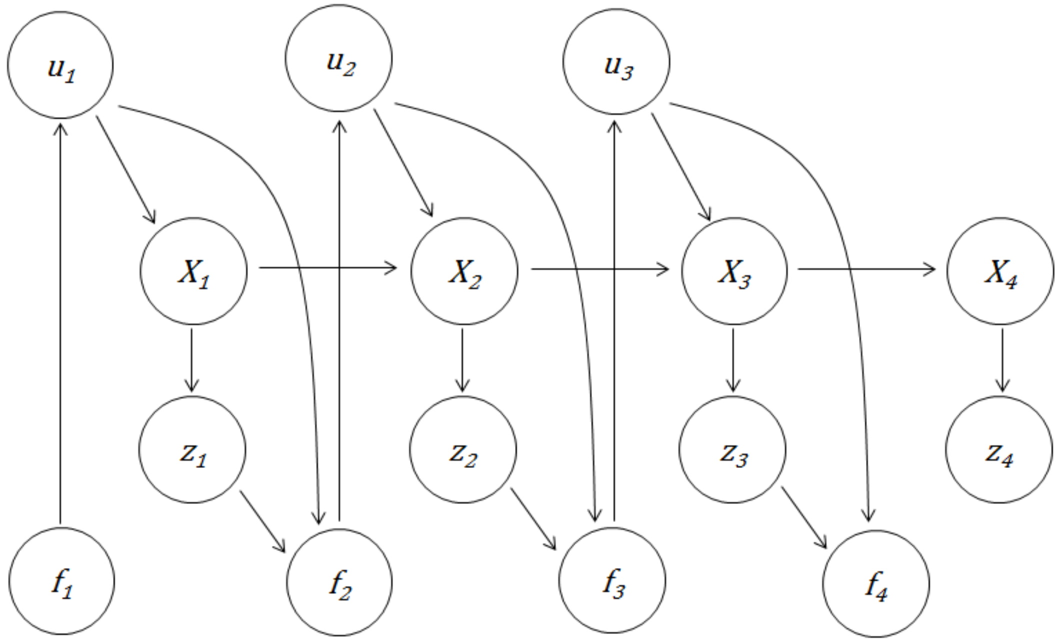

Figure 10: Filtering with a state estimation function.

$x_1 \to x_2 \to z_2 \to f_3 \to u_2$
$x_1 \to x_2 \to z_2 \to f_3 \to f_2 \to u_2$
$x_1 \to z_1 \to f_2 \to u_2$

All of these paths are blocked, and so $x_1 \perp u_2 | z_1$. (Or, for consistency with previous sections, $x_3 \perp u_4 | z_3$ with previous sections.)

3.4 All past controls and observations.

In Figure 11 we show the most extreme / general example, where the control is based on all past observations and controls. Again, with this model, the derivation still holds.

Given z_1 and z_2 are observed, paths between x_2 and u_3 are:
$x_2 \to z_2 \to u_3$
$x_2 \to x_3 \to u_3$
$x_2 \to x_1 \to z_1 \to u_3$
$x_2 \to u_2 \to z_1 \to u_3$

Any controller based on past observations and controls obeys the assumption that $x_t \perp u_{t+1}$. A thought to think about, then, is how would you correct the filter derivation in the book if u_{t+1} were based off of x_t?

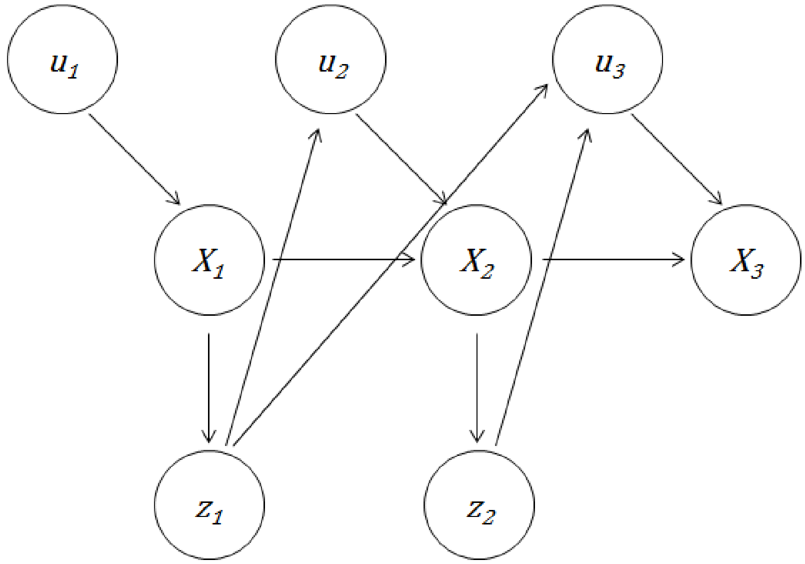

Figure 11: Controls based on all past measurements and controls.

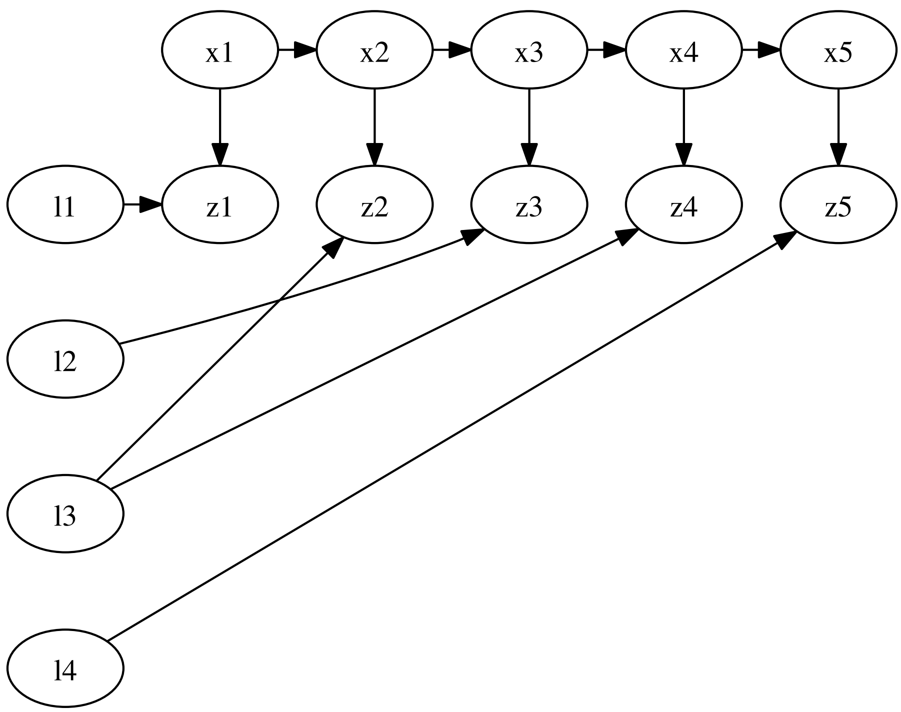

Figure 12: Landmark mapping graphical model.

4 Feature / Landmark Mapping

In landmark mapping, we assume that all landmarks are uniquely identifiable. Therefore, what are the random variables?

- The robot's pose: $x_1 \to x_2 \to x_3 \to x_4$
- The location of landmarks (in this example, we only see one landmark at every timestep): l_1, l_2, l_3, l_4
- Observations: z_1, z_2, z_3, z_4

An example of a landmark-based graphical model is shown in Figure 12.

In Figure 12, all z's are observed. If this is all the information we have, then we are solving the simultaneous localization and mapping (SLAM) problem.

We can ask the question, "Are l_2 and l_3 independent?" Although they may start out independent, as a robot explores, nearby landmarks become dependent based on the fact that they are close together. If they are seen in rapid succession, this gives us a relationship between them. This is illustrated in Figure 13. As the robot travels forward, both landmarks shift in the robot's view relative to the robot's position.

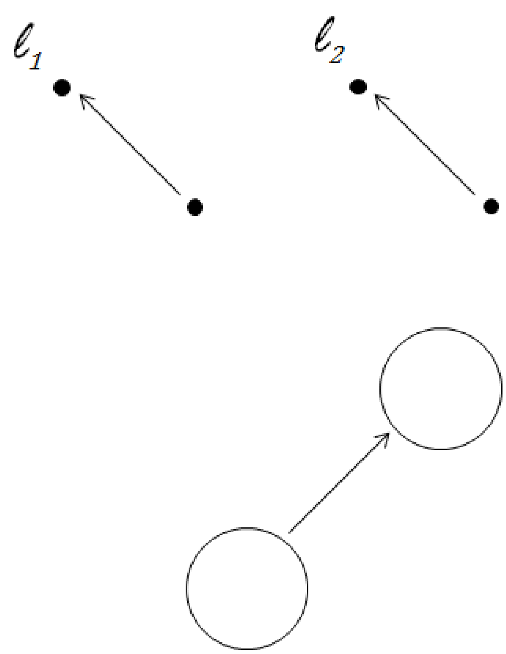

Figure 13: Related landmarks— as robot moves (circle), both landmarks 1 and 2 move together relative to the robot.

However, one last question to ponder, is "Why is Figure 12 a really weird graph?"