

The Art of the SVR Speech Seminar

Marcus Tomalin
University of Cambridge

November 19th 2002



Presentation: November 2002

Overview

- What are SVR Speech Seminars?
- Why give SVR Speech Seminars?
- Preparing a Seminar
- Talking and Pointing
- Answering Questions

What are SVR Speech Seminars?

The word 'seminar' can mean several different things:

- A small class at a university that meets for discussion and research
- A class meeting to study under a specified person
- A conference of specialists

The term comes from the Latin word *seminarium* which means 'a place where seeds grow'.

What are SVR Speech Seminars?

The SVR Speech Seminars enable various speakers to present a summary of their recent speech technology research to the SVR group.

The SVR Speech Seminars usually last for 1 hour and consist of:

- a 45 minute presentation given by a single speaker.
- 15 minutes during which the speaker answers questions about the topics discussed in the presentation.

Often food and drinks are offered to encourage people to attend :)

What are SVR Speech Seminars?

The SVR Speech Seminars usually focus upon speech technology (i.e., recognition, synthesis, web-applications), but presentations on other topics are encouraged if they have potential speech applications:

- Signal Processing (ambient noise, superimposed speech etc.)
- Acoustics (acoustic features, spectral analysis etc.)
- Linguistics (acoustic phonetics, syntactic theory etc.)
- Psychology (language processing, quantifying text readability etc.)
- Mathematics (statistical pattern recognition, chaos theory etc.)

Why give an SVR Speech Seminar?

There are many reasons why giving SVR Speech Seminars is good for you. The experience enables you

- to improve your communication skills
- to acquire pre-conference practice
- to obtain SVR Group feedback

Communication Skills

The best seminars are usually:

- concise
- precise
- coherent

Therefore, giving a seminar can improve your ability to talk about your research clearly.

Pre-Conference Practice

If you have a paper accepted for a conference, you will have to present the paper at the conference. Therefore giving an SVR Speech seminar enables you:

- to prepare a set of slides in advance
- to practice giving a seminar presentation in a friendly environment
- to practice answering questions about your research

Consequently, giving an SVR Speech Seminar *before* attending a conference will improve the standard of your conference presentation.

SVR Group Feedback

Giving an SVR Speech Seminar enables you

- to share your ideas/research with other members of the group
- to explain your research to someone other than your supervisor
- to answer questions from knowledgeable colleagues.

Attending an SVR Speech Seminar enables you

- to know who is working on certain research problems
- to support your colleagues so that they will support you when it is your turn to give a seminar

Preparing a Seminar

It is sensible to do the following when preparing to give an SVR Speech Seminar:

- plan your presentation carefully
- create a set of slides
- practice giving the seminar (in the language you will use!)
- think of possible questions you might be asked (in the language you will use!)

Planning the Presentation

A standard presentation could include

- an overview of the problem to be discussed
- an explanation of the underlying theory
- an analysis of implementational issues (if necessary)
- information concerning results obtained
- some discussion of the results
- some ideas for future research

Preparing Slides

The convention in the SVR Group is to use \LaTeX slides. A template for these slides can be found in:

`/home/tigger3/mt126/seminars/slides/templates/SVR_seminars.tex`

You can copy and use this template.

The slides can be projected from a laptop using `acroread`.

As a rough guide, plan to spend 2 minutes on each slide during the seminar, and try to avoid simply *reading* the slides.

Preparing Slides

Try to avoid creating slides that are

- too complicated
- badly set out on the page
- too small
- **TOO BIG**
- full of *unnecessarily* long and difficult equations

Always try and simplify while still conveying the substance of your work.

A Bad Slide

The equation for FCMs are

$$P(\mathbf{x} | \mathbf{y}) = \frac{|\Psi|^{-\frac{1}{2}}}{(2\pi)^{\frac{p}{2}}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \mu - \Lambda\mathbf{y})^T \Psi^{-1}(\mathbf{x} - \mu - \Lambda\mathbf{y}) \right\} \quad (1)$$

$$P(\mathbf{x}) = \frac{|\Psi + \Lambda\Lambda^T|^{-\frac{1}{2}}}{(2\pi)^{\frac{p}{2}}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \mu)^T (\Psi + \Lambda\Lambda^T)^{-1}(\mathbf{x} - \mu) \right\} \quad (2)$$

And this is self-explanatory.

A Better Slide

The multivariate Gaussian distribution involving Factored Covariance Matrices (FCMs) takes the form

$$P(\mathbf{x}) = \frac{|\Psi + \Lambda\Lambda^T|^{-\frac{1}{2}}}{(2\pi)^{\frac{p}{2}}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \mu)^T (\Psi + \Lambda\Lambda^T)^{-1} (\mathbf{x} - \mu) \right\} \quad (3)$$

where Λ is a $(p \times q)$ matrix of constants (called 'factor loadings') and Ψ is a $(p \times p)$ positive-definite diagonal matrix of variances.

Talking and Pointing

When giving a seminar, try to

- speak slowly
- speak clearly
- avoid specialist jargon (unless defined on the slides)
- point at the screen behind you (not the laptop screen!)
- finish on time

Answering Questions

After the 45 minute presentation, there will be questions from the audience. Usually the person who introduced the seminar will lead the question and answer session.

It is important to remember that

- the questions are not always audible
- the questions are not always well-expressed
- the questions are not always relevant
- the questions are not always interesting

Answering Questions

If you don't understand a question:

- ask the questioner to repeat the question
- ask the questioner for clarification
- ask the person leading the question and answer session for clarification
- say that you don't understand

Answering Questions

If you do understand a question:

- try to answer it fully but concisely
- try to avoid repeating yourself lots of times
- try to avoid introducing lots of jargon that has not already been introduced
- try to make sure that everyone can hear both the question and the answer

Potential Horrors

Some of the scary things that can happen include:

- no one turns up to listen
- the equipment fails
- the fire alarm sounds
- you forget what you wanted to say

Potential Happiness

Some of the pleasant things that can happen include:

- lots of people turn up to listen
- you feel more confident about discussing your work in public
- some of the questions give you new ideas to explore
- I won't ask you to give another seminar for many months!

Advertisement

The first proper SVR Speech Seminar of the year:

Linear Gaussian Models for Speech Recognition

Antti-Veikko Rosti (SVR PhD Student)

TIME: 1-2pm Tuesday 26th November

PLACE: LR5

Refreshments will be served!