Doing a PhD in AI

José Alferes

Simpósio Doutoral em Inteligência Artificial

Aveiro, 12 de Outubro de 2009

What is a PhD about?

- Learning how to make research
  - Learning how to study state of the art work
  - Learning how to be critical about the current state of the art
  - Finding out problems in which advances can be made (and you are convinced that it is important, and you can make them)
  - Being creative in finding solutions
  - Testing your solutions and comparing to extant

What it gives you

What it gives you

- Innovation skills

What it gives you

- Innovation skills
  - A PhD is a much more a degree for innovation, than one for techniques or specialised knowledge (but it also gives you these two)

What it gives you

- Innovation skills
  - A PhD is a much more a degree for innovation, than one for techniques or specialised knowledge (but it also gives you these two)
- Writing skills

What it gives you

- Innovation skills
  - A PhD is a much more a degree for innovation, than one for techniques or specialised knowledge (but it also gives you these two)
- Writing skills
- Reading skills (and being critical about what you read)

What it gives you

- Innovation skills
  - A PhD is a much more a degree for innovation, than one for techniques or specialised knowledge (but it also gives you these two)
- Writing skills
- Reading skills (and being critical about what you read)
- Research habilitation

What it gives you

- Innovation skills
  - A PhD is a much more a degree for innovation, than one for techniques or specialised knowledge (but it also gives you these two)
- Writing skills
- Reading skills (and being critical about what you read)
- Research habilitation
- A degree that allows you to apply for academic/research positions

What it gives you

- Innovation skills
  - A PhD is a much more a degree for innovation, than one for techniques or specialised knowledge (but it also gives you these two)
- Writing skills
- Reading skills (and being critical about what you read)
- Research habilitation
- A degree that allows you to apply for academic/research positions
  - and also to impress (some) family and friends

What is required to begin

What is required to begin

- An MSc
  - Strong background

What is required to begin

- An MSc
  - Strong background
- Scientific Inquiry/Curiosity
  - A PhD is about advancing the state of the art; it is not routine work!

What is required to begin

- An MSc
  - Strong background
- Scientific Inquiry/Curiosity
  - A PhD is about advancing the state of the art; it is not routine work!
- Other qualities that help
  - Open mindedness
  - Self confidence
  - Perseverance
  - Argumentation

While doing it, provides

- Freedom for developing your own ideas
  - together with responsibility and very hard work
- Integration in a global research community
  - Research centres, summer schools, conferences, workshops
- Maturity and Recognition

Especially in AI

- A great area to innovate
  - fascinating new themes
  - relatively new area, where problems crying out for innovation may not be that difficult to find
- A scientific area at the cross/boundary of several disciplines

A word on multidisciplinarity

- More and more science is driven by global challenges
  - Global warming
  - Ageing society
  - Information society
  - Water / food shortage
  - Energy needs
  - ...
- Governments, funding agencies, and also industries strive to make scientists, and other professionals, attack these challenges globally

On multidisciplinarity

- Global challenges require multi-disciplinarity
- Global challenges require innovation
- A PhD in AI is privileged for this
  - A PhD is learning innovation
  - Studying on the boundary of several disciplines improves abilities to work in multi-disciplinary teams
  - AI studies several problems directly applicable:
    - planning; scheduling; rule and expert systems; knowledge representation and retrieval; learning; decision support systems; smart buildings; ambient intelligence; reactive systems; diagnosis; CSPs; pattern recognition; etc; etc; ...

Starting a PhD

- The first, and **very** important task in a PhD, is the choice of a topic
- The topic of your thesis must:
  - be one which you are enthusiastic about
  - be at *PhD level*
  - be do-able in 3-4 years
  - have some professor knowledgeable in the school, willing to supervise

Starting a PhD

- The first, and **very** important task in a PhD, is the choice of a topic
- The topic of your thesis must:
  - be one which you are enthusiastic about
  - be at *PhD level*
  - be do-able in 3-4 years
  - have some professor knowledgeable in the school, willing to supervise
  - be one which you are enthusiastic about

    After all, you'll have to work hard on it for several years!!

First steps in a PhD

- Choose the area (inside AI, of course!)
- Choose the supervisor
- Get his help on finding interesting topics
- Start preparing a thesis proposal
  - Follow your supervisor's advice for this. He has more experience than you
    - He already went through this (at least once!)
    - At this time, you don't know yet how to deal with your supervisor (later you may hope to learn this)

The thesis proposal

- What is it about?
  - You have an idea for solving some important problem, and you must **convince others to invest** resources on it.
  - Why should anyone put resources into your ideas?

The "must have"s

- A clear identification of the problem to solve
  - Why is this problem important
  - Why is it new
- What is your idea to solve it
  - Why are you the right team to solve it
- How do you plan to attack it
  - How can it be done in the prescribed time period

Clearly identify the problem

- What is the problem?
  - Don't be vague!
  - Wrong examples:
    - "I want to work on a better system for XPTO"
    - "I want to explore the area of XPTO"
- Be realistic when choosing the problem to solve
  - "I want to solve the problem of P=NP"
  - You must provide evidence that you stand some chance of solving the problem!

Impact

- Why is it an interesting problem?
  - Is it research at all? Or just routine application of known techniques?
  - Is it at PhD level?
- Why is it important?
  - Would anyone care if the problem is solved?
  - Which are the conferences that would publish the results of your PhD work?

PhD level projects

- Must expand the existing body of knowledge
  - It can't be something that is built solely on the basis of already existing knowledge
  - It is also not enough to be a different way of doing something. It must be a arguably better one, at least in some aspects
  - It must be recognised by the scientific community, measured by publications in leading scientific venues
- It's not that difficult. You just have to be creative!

Originality

- Have others tried to solve the problem?
  - **Important**: make sure that others haven't solved it already!
- How **exactly** does your proposal, if successful, advance the state-of-the-art
  - Of course, for showing this, you need to provide some state-of-the-art

Originality and breadth

- It may be easier to be original in narrow / very focused topics
- But the risk is higher
  - If your originality lies in solving very precise / focused problem, it falls apart if someone else solves it a short time after
- For a long period, typically 3-4 years, there should be enough breadth, with potential for generating original results

Your idea to solve it

- You **must** have some idea of how to solve the problem
- You **must** clearly explain your idea
  - Give real technical arguments to explain the idea
  - Give some evidence that it is a promising idea:
    - results of preliminary work
    - prototypes
    - publications

You're the right person

- You must show that you are the right person for this work
  - You have the right idea **but also**
  - You already know well the current state of the art, and foundational material
- You must have a state-of-the-art survey where you demonstrate that you totally know the field.
  - But more than that! You have a critical perspective of it, and would like to make things better / differently

Attacking the problem

- You have a problem to solve
- You have a starting promising idea for it
- You are the right person(s)
- And now what?
  - It is not enough to simply say that you'll work on it
- You must absolutely say what you are going to do, and how you plan to achieve it

Plan of the work

- Explain how (what steps) to go from your promising idea to the solution of the problem
  - Your promising idea is necessarily far from the solution (otherwise it wouldn't be a PhD plan)
  - Tell the story of how to go from there into the goal
  - In the end put (tentative) dates on each of the steps

Useless general plans

- In the 1st year I will study state-of-the-art
- In the 2nd year I will explore the idea and develop some theoretical results
- In the 3th year I apply the results and write the final version of thesis

- Such a general plan is useless!
  - It doesn't help the evaluator to check that it is do-able
  - It doesn't help you in guiding the work

Between the plan and the writing

Between the plan and the writing

- This will be funniest part!

Between the plan and the writing

- This will be funniest part!
  - It is when you do research

Between the plan and the writing

- This will be funniest part!
  - It is when you do research
  - It is when you are free to explore your ideas

Between the plan and the writing

- This will be funniest part!
  - It is when you do research
  - It is when you are free to explore your ideas
  - It is when you write papers about those ideas

Between the plan and the writing

- This will be funniest part!
  - It is when you do research
  - It is when you are free to explore your ideas
  - It is when you write papers about those ideas
  - It is when you go to conferences

Between the plan and the writing

- This will be funniest part!
  - It is when you do research
  - It is when you are free to explore your ideas
  - It is when you write papers about those ideas
  - It is when you go to conferences
  - It is when you meet researchers all over the world

Between the plan and the writing

- This will be funniest part!
  - It is when you do research
  - It is when you are free to explore your ideas
  - It is when you write papers about those ideas
  - It is when you go to conferences
  - It is when you meet researchers all over the world
  - It is when you gain recognition

The dark side of it

The dark side of it

- But it won't always be easy. Sometimes:
  - you'll feel lost
  - you'll feel isolated
  - you'll feel unfairly judged
  - you'll feel frustrate

The dark side of it

- But it won't always be easy. Sometimes:
  - you'll feel lost
  - you'll feel isolated
  - you'll feel unfairly judged
  - you'll feel frustrate

- And it is hard work
  - It is difficult to overestimate this!

The dark side of it

- But it won't always be easy. Sometimes:
  - you'll feel lost
  - you'll feel isolated
  - you'll feel unfairly judged
  - you'll feel frustrate

- And it is hard work
  - It is difficult to overestimate this!

- Nevertheless it is possible to survive
  - Look at those who have a PhD. They survived and regained mental sanity

The dark side of it

- But it won't always be easy. Sometimes:
  - you'll feel lost
  - you'll feel isolated
  - you'll feel unfairly judged
  - you'll feel frustrate

- And it is hard work
  - It is difficult to overestimate this!

- Nevertheless it is possible to survive
  - Look at those who have a PhD. They survived and regained mental sanity
    - Regarding this last part, you may have to search more carefully. But you'll surely find examples :-)

Some pitfalls you have to avoid

Some pitfalls you have to avoid

See Alan Bundy et al's *"Researcher's Bible"* for details

Some pitfalls you have to avoid

See Alan Bundy et al's *"Researcher's Bible"* for details

- Don't try to save the world with your thesis
  - If you really want to save the world, first do the thesis. After that, it will be easier...

Some pitfalls you have to avoid

See Alan Bundy et al's *"Researcher's Bible"* for details

- Don't try to save the world with your thesis
  - If you really want to save the world, first do the thesis. After that, it will be easier...
- Don't fall into the "ambitious paralysis"
  - If you're overly ambitious at the beginning, you'll never start

Some pitfalls you have to avoid

See Alan Bundy et al's *"Researcher's Bible"* for details

- Don't try to save the world with your thesis
  - If you really want to save the world, first do the thesis. After that, it will be easier...
- Don't fall into the "ambitious paralysis"
  - If you're overly ambitious at the beginning, you'll never start
- Good ideas don't fall from heaven
  - Read related work, talk to people, tackle simplified versions, write down half-backed ideas and present them locally

More pitfalls

More pitfalls

- Don’t be overwhelmed by implementations
  - You may need to implement things, but there is more to the PhD than the implementation. Don’t get lost with it

More pitfalls

- Don't be overwhelmed by implementations
  - You may need to implement things, but there is more to the PhD than the implementation. Don't get lost with it
- Don't get lost in abstraction
  - Abstraction is ok. But be sure to tackle the issues in concrete situations

More pitfalls

- Don't be overwhelmed by implementations
  - You may need to implement things, but there is more to the PhD than the implementation. Don't get lost with it
- Don't get lost in abstraction
  - Abstraction is ok. But be sure to tackle the issues in concrete situations
- Don't be trapped waiting for the expected / desirable experimental results

Even more pitfalls

Even more pitfalls

- Don’t get trapped in an Ivory Tower
  - There is life, and science, besides your PhD topic

Even more pitfalls

- Don't get trapped in an Ivory Tower
  - There is life, and science, besides your PhD topic
- Don't act as a misunderstood genius
  - Sometimes you'll have unfair criticisms
  - But always look at the criticisms and see how you can improving explaining your ideas
  - Most times critics are fair

An extra tip

- Learn to interact with your supervisor
  - Your life will be a lot easier if you have a good interaction with your supervisor
  - Your supervisor is willing to help
  - He will feel proud of your success
    - He'll feel that your success is a bit his success too
- But
  - He won't have as much time as you do
  - He is not able to read your mind
  - He won't know as much details of your work as you do
- See Toby Walsh's *"Managing your supervisor"*

Writing the thesis

- Writing a PhD thesis is a hard and painful task
  - Also don't underestimate this sentence. It is really painful to write a thesis!
- But you have to do it
  - There is no PhD without the thesis
- And you'll also learn from it
  - writing
  - preserverance

How to start

- Read at least a good PhD thesis before starting to write yours
  - E.g. ECCAI best thesis awards
- Start by the abstract
  - Write what is the message: what is the **thesis** after all, and how you defend it
  - Never forget the meaning of “thesis”

thesis |ˈθ ēsis|

noun (pl. **-ses** |-sēz|)

1 a statement or theory that is put forward as a premise to be maintained or proved

Next step: TOC

- Set a tentative table of contents
  - This will give you a draft of the logical structure of the document
- Example template:

1. Introduction
2. Background
3. Theoretical or algorithmic developments
4. Empirical results
5. Analysis
6. Related work
7. Conclusion

Pitfalls in thesis structure

- A thesis is **not** a report of what you did in the last 3 years
  - I started by doing this; then I did that; then I backtracked; ...
- A thesis is **not** a dump of everything you've done
  - Dead ends won't appear
  - You'll need to fill gaps while writing the thesis
- A thesis **is** a logical reconstruction with a coherent message, that is defended

Timetable

Timetable

- Make a tentative timetable
  - This will give you targets to meet
  - It will allow you to know how far you are from the end

Timetable

- Make a tentative timetable
  - This will give you targets to meet
  - It will allow you to know how far you are from the end
- Be realistic about the timing
  - Otherwise you'll feel frustrated for not keeping the pace set by it

Timetable

- Make a tentative timetable
  - This will give you targets to meet
  - It will allow you to know how far you are from the end
- Be realistic about the timing
  - Otherwise you'll feel frustrated for not keeping the pace set by it
  - "A thesis is like a baby. Give it 9 month"
    - Toby Walsh. *How to write a thesis*

Filling in the thesis

- Start writing the middle chapters
  - It is where you'll have most of the material
  - In them, you'll still get to do some research
  - Among these, start where you feel more comfortable
- Leave the conclusions and introduction to the end
- Never forget your thesis' message, while filling in the document
- Seek feedback while you write
  - Leaving all feedback to the end is definitely not a good idea

Know when to stop

- You'll always find more things that could be done in the thesis
- It is very important that you know when to stop
  - Define limits for what can go in the thesis
  - Your supervisor will surely help here
- In the final phase, finding new typos is an endless process

Know when to stop

- You'll always find more things that could be done in the thesis
- It is very important that you know when to stop
  - Define limits for what can go in the thesis
  - Your supervisor will surely help here
- In the final phase, finding new typos is an endless process
  - Be sure that, no matter how long you take this process, you'll find a new typo the day immediately after making the copies, at the latest

You’re almost there

You're almost there

- After doing the research and writing the thesis you're almost there!
- But you still have to defend your thesis

You're almost there

- After doing the research and writing the thesis you're almost there!
- But you still have to defend your thesis
- Don't panic!
  - Most likely, nobody in the world knows more about the theme than you

You're almost there

- After doing the research and writing the thesis you're almost there!
- But you still have to defend your thesis
- Don't panic!
  - Most likely, nobody in the world knows more about the theme than you
- Enjoy it!
  - It is not often that you get a group of expert to really read careful what you do
  - And (this is even less often) they really will want to talk to you about your work!

Prepare your defence

- By then, you already have prepared several talks
  - So it won't be such a big deal to prepare a presentation
  - But prepare this one with special care
    - Rehearsing it in from of colleagues is not a bad idea
- Be prepared for the discussion
  - Ask your supervisor, what examiners look for in thesis defences
  - See how that applies to your thesis, and be prepared for the most obvious questions

Conclusion

Conclusion

- It is worth doing a PhD

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot
  - You'll have a lot of fun

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot
  - You'll have a lot of fun
  - You'll have a lot of work

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot
  - You'll have a lot of fun
  - You'll have a lot of work
  - You'll be much more equipped after it

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot
  - You'll have a lot of fun
  - You'll have a lot of work
  - You'll be much more equipped after it
    - for innovation, and for facing the challenges of the future

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot
  - You'll have a lot of fun
  - You'll have a lot of work
  - You'll be much more equipped after it
    - for innovation, and for facing the challenges of the future
  - You'll have a PhD diploma

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot
  - You'll have a lot of fun
  - You'll have a lot of work
  - You'll be much more equipped after it
    - for innovation, and for facing the challenges of the future
  - You'll have a PhD diploma
    - Good, even essential, for some jobs

Conclusion

- It is worth doing a PhD
- It is worth doing a PhD in AI
  - You'll learn a lot
  - You'll grow a lot
  - You'll have a lot of fun
  - You'll have a lot of work
  - You'll be much more equipped after it
    - for innovation, and for facing the challenges of the future
  - You'll have a PhD diploma
    - Good, even essential, for some jobs
    - Nice to show to family and friends

Questions?