20211210

LUND UNIVERSITY

AI-education @ Lund University

Report #3

LUND UNIVERSITY

MAGNUS EKBLAD | MIKAEL SUNDSTRÖM

8. AI Education suggestions

8.1 Strategic framing of courses

Courses can be, and are, of course developed based on many factors including teacher capacity or interest, timely topics, links to current research, and much else besides.

However, if we imagine this as a strategically important field, it might be helpful to set up some sort of mapping effort to be better able to contextualise and possibly structure development efforts and priorities. This is not in any way intended to challenge or replace intra-faculty development processes, but could be a useful *complement* when foci or target groups straddle faculty divides, or when target groups are external.

If we place our categories along a single (admittedly much simplified) continuum from core technological aspects (left) to societal impact and conceptualisations (right), we end up with something like this (table 8.1, below):

Table 8.1 – Framework categories as a single continuum

Links to society			Application		Fundamental techniques		
Theory foundation	**Techniques/ methods**	**Solution complexes**	**Application (science users)**	**Application (end users)**	**Impact on society**	**Governing AI**	**Perceptions of AI**

For each of the eight framework sub-categories we think it would be profitable to brainstorm about broad needs across the university and beyond. What follows below is *one* way of imagining the grander strategic picture aided by the framework (though of course not one with self-evident primacy we hasten to add).

8.1.1 Strategic concern: thinking about what to boost or scale back

A first consideration would be a "straightforward" strategy which, if any, of these eight areas should be boosted or indeed scaled back at Lund University. In this report (chapter 4), we for instance note the glaring lack of courses in Sweden pertaining to category 8, *perceptions of AI.* Maybe an area in need of a strategic boost? But we are not here advocating either way – category eight and the noted lack of existing courses is an *example* how the framework, and the findings of this report, can aid actual strategic discourse. Unless we *specifically* discuss category 8 and its possible "course viability" (and desirability), LU-wide strategy relating to that particular educational focus is simply a moot point.

We will from this point on use the overarching *three* dimensions, rather than the more granular eight. The reason is that it helps us maintain a strategic mindset – it is simply easier to get "strategically distracted" the closer we get to sleeves-up-oriented decisions and minutiae relating to actual course work. That is after all where capacity issues and political realities will naturally be major factors.

The aim basically remains the same, however: the (for now) trisected framework can be used to structure broad strategic thinking and resource allocation about what these students, whatever their background, will be focusing on… and what the University would *want* them to focus on. That leaves us with the following process (figure 8.1, below).

Figure 8.1 – Visualising strategy, part 1

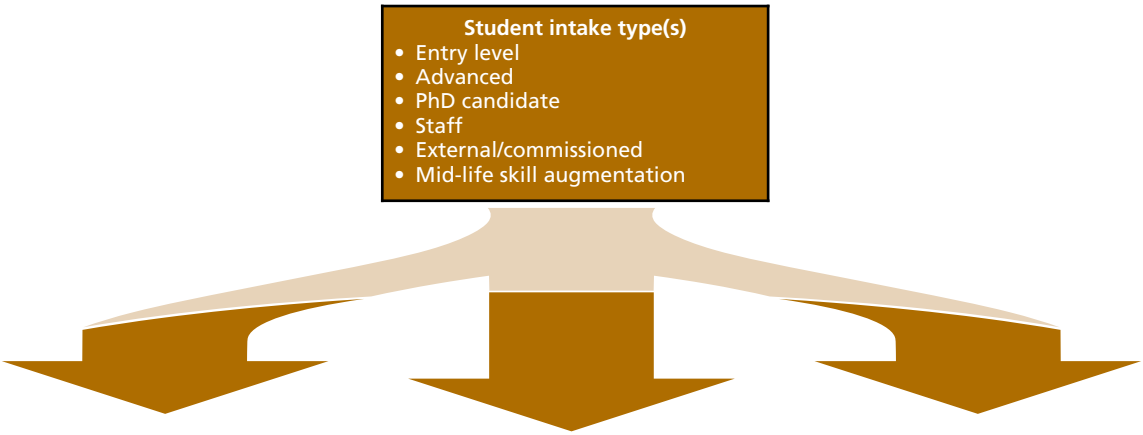

1. technical orientation			1. end-usage orientation		Socio-analytical orientation		
Fundamental techniques			Application		Links to society		
Theory foundation	Techniques/ methods	Solution complexes	Application (science users)	Application (end users)	Impact on society	Governing AI	Perceptions of AI

8.1.2 Strategic concern: student intake

Student intake is obviously an ever-present point of consideration. Some related questions are for that reason nothing out of the ordinary, such as: would courses *technically* cater to current students, mid-life complement studies students, or commissioned education initiative students? International or national students? Level of studies? How might HÅS distribution (full-time student equivalent) play out financially? How many can we cope with? We see no need to expand on such issues here as they are routinely discussed, and we could offer little in the way of value addition.

What *does* provide value addition, however, is how we can now link envisaged/identified student groups to the overarching dimensions. What will mid-life students wishing to broaden their understanding about AI need *as categorised in the framework?* Technical fundamentals? Learning about apps powered by AI? Understanding impact on society and how AI can be governed?

What might LUCE's potential or actual partners need? What about early stage students? Late stage students? Doctoral candidates? Staff?

These kinds of broad but structured discussions will avoid an initial and interfering focus on what is already on offer or in the pipeline, and how to direct students to these existing options. Each of the potential student groups will instead be discussed, one after the other, based on perceived *needs* and *wants*, but separated from the on-the-ground situation.

Results of such discussions can then *eventually* be mapped against what we *do* offer (and this report will cover much ground in that respect), in order to identify where we already have readied options, and where we need to build new opportunities. A few examples will be given later in this chapter.

8.1.3 Strategic concern: reconciling wants and needs

What different types of student *want* will very likely direct their eagerness to invest by signing up to an offered course or programme. A perennial issue is that this might not coincide with what others think they actually *need.* In programmes, there is ample opportunity to link in components that students would not necessarily select of their own volition, because it is felt that they need these skills. For courses this is more difficult, and to try a high-handed didactic approach – "this is what you need" – is a dead end.

A solution is to consciously think about secondary *directionality.* Pushing in as a design parameter that even an ostensibly narrow "technical orientation" course focus should at least touch upon elements that are more "right-leaning" (in framework terms), would of course alter both development processes, course content and student exposure to different takes on AI. A similar requirement that a socio-analytical focus should offer glimpses under the technical hood would do the same. In both cases implementation of such thinking would broaden the offerings, and likely promote or even presuppose trans-faculty cooperation.

Obviously attempts at forcing any such changes on existing courses would engender political opposition in some quarters (strategic overrides always do), but the discussion itself is worthwhile. If, after actual discussions about repercussions and desirability stakeholders agree that a certain course is ill served by such a change, that is still an improvement over no change based on no reflection.

Anyway, to enter such secondary directionality in the visualised form would yield something like this (figure 8.2, below). In that figure we also add "technical course type" which must of course at some point be considered.

Figure 8.2 – Visualising strategy, part 2

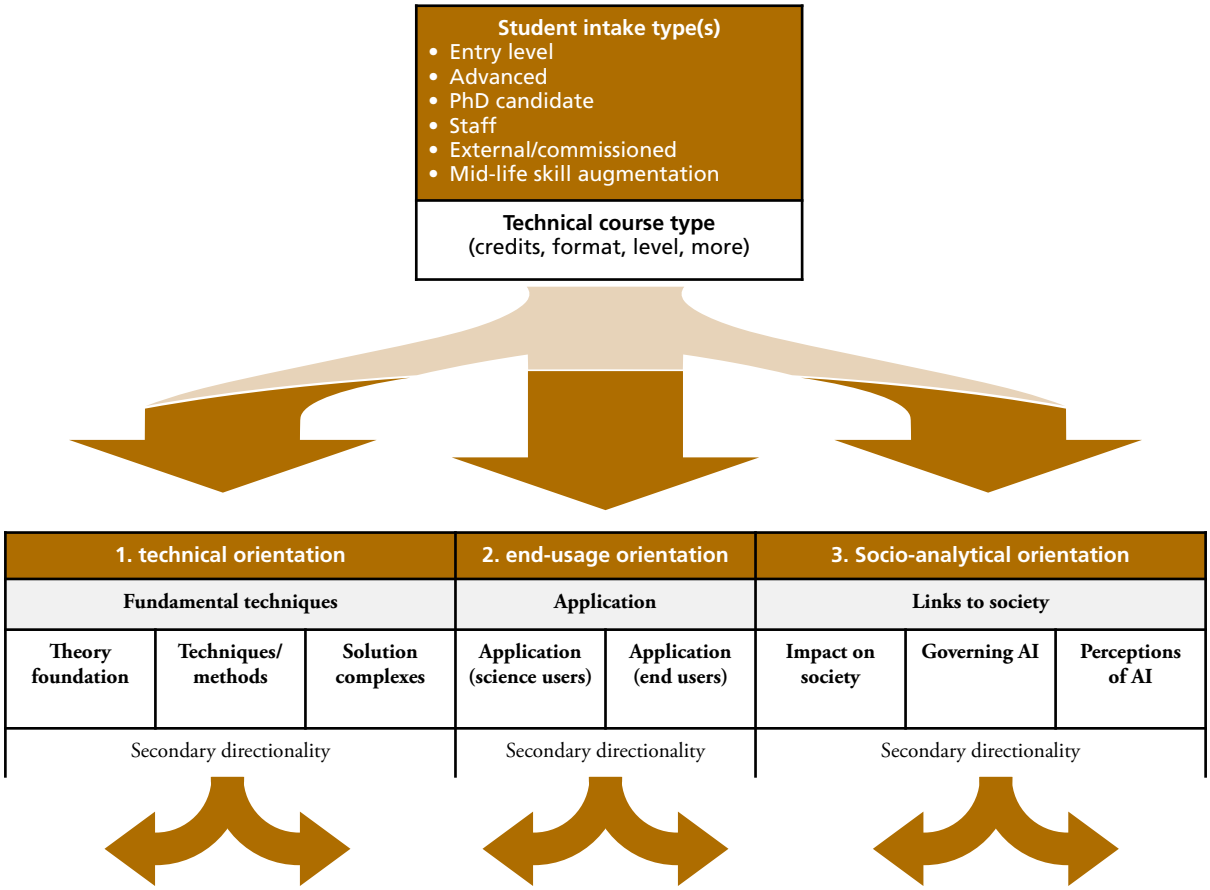

In the following, we'll put our money where our mouth is, and use some findings in this report plus the proposed way of thinking to guide us to a few very concrete suggestions.

8.1 Some proposed courses

8.1.1 End-Usage Orientation: Understanding AI-powered apps

We start with, and spend the most time on, the "end-usage" part of the framework, as it provides plenty of appealing, and easily "sellable" options. This part of the framework also offers a range of obvious options to demonstrate/realise secondary directionality.

To update skills about available AI-powered *applications*, recurring courses introducing such tools and how they can be used in practice are needed. Such courses should reasonably include discussions about how they complement conventional approaches, what ethical and methodological implications they have etc. The two-pronged dimension (Scientific use / "end-user") can be used to structure inward-looking (how can the various scientific disciplines use

these tools) as well as more general (how can tools be used in fields outside of research & engineering development). The former could for instance engender a course for PhD candidates across the University to make sure that they gain at least a basic understanding of some possibly relevant AI-powered complements. Needs for this will naturally differ between faculties and individual departments – but it seems problematic to us that many doctoral candidates can pass through an entire education without having at least a high-visibility option to take a look at AI tools which might conceivably have bearing on their field of research.

More general courses could include one tentative named *An introduction course in AI-powered decision-making tools: usage and ethics.*

Other courses might introduce AI-powered tools and methodologies to budding or more senior analysts or scientists, including PhD candidates and more senior staff at LU. This would be a way to infuse more knowledge about alternative or complementing methodologies in a range of fields.

Concrete suggestion

For *student intake,* we first went for a focus that would be as alluring as possible to a range of student types. Introducing and discussing actual AI-powered apps and how they can be employed by end-users would appeal to almost all conceivable student groups, but here we exclude scientific users, whose focus will by necessity be narrower and more bespoke.

We imagine the course as entry-level in nature, but with a ready-made path to more advanced options (see later in this text). The proposed course is envisaged as a 7.5 credit one – but links to "sibling" AI-focused courses should be planned out and highlighted. The course format should mean that a demand-driven extension to 15 credits should be relatively straightforward.

The *format* if the course is by no means self-evident if it should be a viable option for a wide assortment of potential students. Across the university there are diverging norms whether courses should be offered consecutively or in parallel but at half speed. To maximise structural compatibility, the best form would be… *both:* either in separate terms, or starting at the same time but then offering a fast and a slow completion track (technically this would likely demand separate syllabi, but it *might* be possible to offer a 50 % speed option as the default, but including duplicate events plus an examination opportunity in the first half to allow students to complete the course before the technical end date).

The course should in either case include clear-cut sub-components that can be presented as self-contained units by LUCE, whose clients tend to need shorter, more concentrated opportunities. That way LUCE clients can be attached on a component basis. A set of recorded teaching material aimed at students who are thus parachuted in without access to the "normal" course introduction should be devised.

Our concrete suggestion for a course is *Understanding AI-powered apps*, which should include and discuss a range of example apps aimed at decision-making support in different fields. It is conceived as an entry-level course, with open options for LUCE clients (see above). Core content should comprise usage, comparison to competing/older decision-making strategies and related ethical concerns. Depending on the teaching team composition (more specifically which faculties/competencies are represented) and promotion opportunities (not least LUCE's identified

needs), the syllabus should be set up in a way that allows some flexibility as to which apps are included.

When thinking about *secondary directionality,* we consciously add a more technical sub-focus than the course would "organically" include – a non-negligible course component devoted to improving a layman's understanding of how the AI engine actually reaches its conclusions (exemplified by one or more of the studied apps), with overt links to both material and further LU training opportunities in that area should be included in the course.

8.1.2 Technical orientation: Statistics & AI for beginners

Certain skills are needed to even approach the *technical orientation* of AI. Examples include python programming for beginners and statistics underpinning AI algorithms. We of course realise that such objectives are very much part of faculty-internal education activities at, for example, the School of Engineering and the Science faculty, but courses introducing and then building such skills could conceivably be useful to students – and staff – from all over the University, and may be appealing as commissioned education initiatives as well. Here too, we stay with such a pan-university perspective, and follow the framework pathway when deciding on a suggestion.

Concrete suggestion

Given the already noted relative scarcity of *entry-level* AI courses, we again focus on such a target group but with an option for staff to join as "quasi students" (this model has successfully been piloted at the Faculty of Social Sciences). The course is however primarily directed at students with some knowledge of statistics. In certain faculties such knowledge is far from certain, and we for that reason conceptualise a 7.5 credit "pure" course, and a second option where the 7.5 credit block is preceded by a preparatory stats block with strong tie-ins to the coming AI component. This stats block could conceivably be offered separately too in order to aid, for example, mid-life stats upskilling.

Looking for further synergies

We also note that students writing their theses, or working on equivalent projects sometimes need to brush up on stats in ways that they may not have been able to anticipate earlier in their education. A slim 7.5 credit stats course/track spaced out across an entire term (@ 25 % speed) could be a way to add such an opportunity, and a complementing way to offload some basic technical questions from the assigned project/thesis supervisor(s). Such a format may also hold appeal for complementing target groups, including university staff.

This is intended to crudely exemplify how there may be low-hanging "synergy fruits" ready for the picking across the university when planning consciously widens the scope to include a full range of possible target groups and scenarios. This is obviously easier to do when an interdisciplinary focus – like AI – naturally blurs boundaries between actual disciplines and/or units.

Indeed, this might also be a way to assuage political fears when new broader education strategies are launched (like, say, any educational activities connected to the currently discussed profile areas). Making sure that some new educational elements link to other target groups, offering something new that is useful in other contexts, will almost inevitably improve political receptivity.

Technical format issues resemble what we discussed under 8.1.1 but with the complication of the basic stats complement.

Orientation is technical (far to the left in the framework).

Secondary directionality would be right-leaning – at the very least including end-usage example apps, but conceivably with a component juxtaposing strict statistics with ethical concerns and thus an *impact on society* element. Secondary directionality can also be used to think about conscious lead-ins to complementing courses (such as the *Understanding AI-powered apps* discussed under 8.1.1).

As the report makes clear, Lund University offers a good number of courses with a technical bent, and we have a lot of relevant talent on hand, meaning that these kinds of courses should be relatively uncomplicated to establish – it is more a matter of repurposing and reconfiguring existing resources, and make sure that a widened student intake and secondary directionality can be maintained.

8.1.3 Socio-analytical orientation: Governing AI

At some point, it would be beneficial for technically-oriented AI students to get to grips with challenges pertaining to AI's impact on society – and related governance issues (including ethics). Such a course could also be of use for a range of "AI practitioners" that LUCE might target. LU has talent – albeit with limited spare capacity – on hand to offer courses of this nature.

Concrete suggestion

Student intake would here focus not on novices, but on student groups with some previous AI education exposure. This said, the core content (how [different takes on] AI can and perhaps should be governed is nevertheless introductory in character. For that reason the course is conceptualised as a Bachelor-level course, but with hooks to the Master level (and LUCE clients). There are several technical ways such hooks can be realised, e.g., generally accepted equivalence; or a dual syllabi solution (with the master version requiring, say, an extra examination component).

Technical format: for the most part similar to 8.1.1, with a 7.5 credit baseline, optionally extensible to 15 credits.

Orientation is "right-leaning" (framework-wise), with a clear focus on legalistic and philosophical/ethical matters.

Secondary directionality would be left-leaning. In this case that would mean introducing concrete examples in the form of apps and actual expert systems with AI-engines, but not necessarily more technically oriented than that, as the student intake profile presupposes some pre-existing knowledge of AI fundamentals.

8.1.4 Reflection: *end-usage orientation* as springboard

It seems to us that the app-centric nature of *end-usage orientation* is an excellent attractor for a range of, if not all, potential student groups. To learn about actual or coming AI-powered apps or services that may be useful in one's specific field of endeavour is an enticing proposition. Across Tertiary Education Sweden, this area is not really prevalent, maybe because it is not perceived as of immediate *academic* interest (too… *practical)* – and there are other education actors who are both skilled and active here.

But viewed as a way to *initially* attract different student groups, and then plan organised paths for them to continue to broaden and deepen their AI knowledge at Lund University, it looks like a key strategic tool to employ.

8.2 A Bachelor-level AI programme

8.2.1 Strategic question #1: what student type(s)?

As indicated earlier in this report, there is a relative dearth of opportunities at the bachelor level for prospective AI students in Sweden – and *particularly* so when thinking about an entire programme with that kind of focus. This is perhaps understandable, as "organic" development of a brand new field will naturally proceed from burgeoning research through PhD courses and master level studies to eventually make an impact on the bachelor level.

Strategic decisions that promote lower-level studies can significantly speed up that process, however, in order to cater sooner to more junior students with an interest in AI-related matters. For these reasons, the focus here is on entry-level students, and *the fundamental suggestion is to fast-track the development of a BSc (or BA) in Artificial Intelligence.*

Such a programme, however designed, would be an almost unique element in the Swedish AI-education context, and could also conceivably set off LU master-level complements at a higher level than is currently possible.

But what kind of a programme? The easiest solution would be a faculty-internal programme, which particularly the School of Engineering (LTH) or the Faculty of Science (N) could set up with little difficulty as they have all the resources in place to make it happen.

An alternative is an *interdisciplinary*, indeed interfaculty, programme allowing both a broad intake, and several different exit options allowing students to pursue a more technical or a more non-technical master level continuation depending on inclination. A full-scale University like Lund has the expertise to organise such a programme, a caveat being the sometimes prevailing difficulty to organise cross-faculty education opportunities.

The report also indicates that, relatively speaking, Lund offers fewer "wide-spectrum" courses than some other seats of learning. We get fewer answers how courses are bound together in programmes, but the scarcity of "right-leaning" courses, as well as encountered internal discussions about the need to tie in such elements in tech-oriented education lead us to believe that this is relatively rare. Result: we suggest that the bachelor-level programme is intentionally made wide in terms of scope (this also renders the conceptualisation compatible with the framework's *secondary directionality* prodding). Succinctly put, we want the programme to open up future education opportunities with a technical orientation through a socio-analytical orientation.

Translated to actual bachelor programme terms, this (see below) might be one result, starting at the non-technical end to have time to evaluate and possibly impact students' knowledge of for example stats and programming:

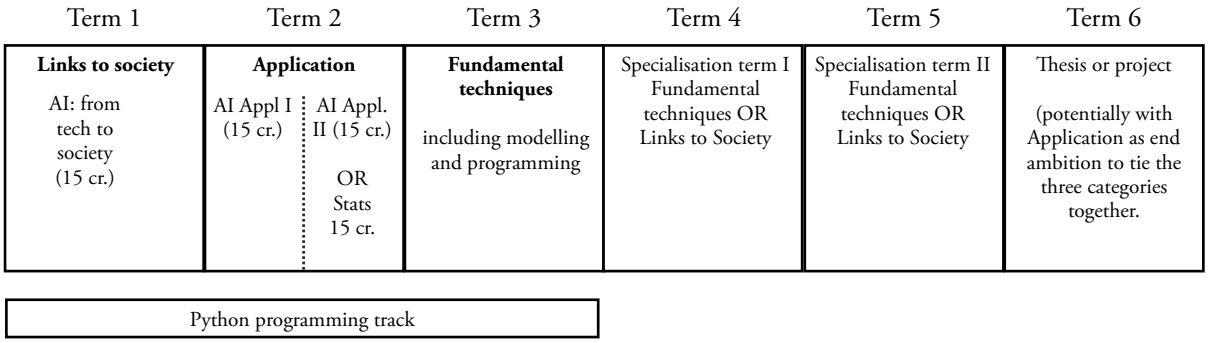

Obviously this example is again primarily intended to get going some thinking, but certain design parameters are probably worth highlighting.

- The potential intake in the example is intentionally made broad, with students with a tech bent as well as a social scientific/humanities background able to apply.
- Terms four to six can be designed to direct students towards a more tech-y or social scientific/humanities-y programme exit – ideally with a few target master level programmes included in the thinking – including a bespoke one (see below).
- Particularly terms one to three could be used as shared building blocks in other, bespoke faculty-internal AI programmes at the Bachelor level. That would make it more resource-effective to develop such complementing programmes. This can open new routes to switch focus mid-education.
- Terms four and five could conceivably be designed to also allow *external* applicants in – a way to maybe bring in mid-career people or offer LUCE options, and mix these with primary programme students. This is a way to bring about synergies and optimise the use of scarce teaching resources.
- Although not a primary design parameter in this model, it could be possible to design terms four to six to comply with *discipline* requirements to proceed to the master level. If these terms were, say, following a first cycle political science track, potentially with bespoke polisci + AI bachelor thesis requirements, such students could then pursue further polisci studies or master-level studies where AI remained the predominant focus.

The overarching "architectural" idea, then, is to establish building blocks that can be joined in many different ways with other education activities to make the most of scarce AI teaching resources.

8.3 Master programme

This short section is merely intended to complement the discussion about a bachelor-level programme above. If there is a clear default pathway for graduating students from the Bachelor

programme, LU can *guarantee* students with an AI bent a complete and well structured AI-focused education from novice-level to master…

…and beyond. We indicated that the bachelor-level programme might optionally include disciplinary elements to ensure eligibility for discipline-specific master programme applications. Here, at the Master level, there must be an equivalent pathway for students to progress to the PhD candidate level if they so desire. Options:

- Making sure that the master programme includes a range of majoring options to ensure eligibility to existing disciplinary PhD candidate positions.
  Benefits: relatively uncomplicated to set up, provided that relevant departments agree to support the majoring option (which would draw on their teaching capacity, via bespoke or existing courses).
  Challenges: very hard indeed to offer more than a limited number of disciplinary majoring options, as each department partnership requires constant servicing.
- Setting up or linking to PhD options which are interdisciplinary by design. The COMPUTE research school could conceivably expand to offer such an option for example – or new initiatives could be established to offer PhD programmes, and linking these to LU research and researchers. Such programme(s) too could be set up to "lean" either rightward or leftward in framework terms.
  Benefits: Lund already has a good infrastructure in e.g. AI Lund and COMPUTE that can be leveraged to realise such ambitions, given strategic funding and direction (designation of AI as a strategic profile area might be one such example).
  Challenges: traditional departments remain the core nodes in the University organisation, and network solutions of various kinds are more cumbersome when it comes to handling routine administration of education (although it can – manifestly so – be done).

For each alternative, the master programme must be prepared accordingly, and make sure that eligibility requirements are met.

It is however also important to fully internalise that many, likely most, students will not be intending to pursue an academic career after the completion of their master-level studies (this may sound obvious in the extreme, but is still an occasionally under-appreciated aspect when master programmes are planned). A concrete suggestion is to design the final 30 credits as a dual-option term: *either* a thesis *or* a project (possibly interning at partner companies or organisations).

8.4 From embryonic ideas to practical implementation

To get the ball rolling, we propose that the *AI Lund* network is tasked with, and funded to, develop some possible education content scenarios, and that possible organisational homes and funding models for such an initiative are meanwhile urgently explored by the University leadership (or group designated by them).