

National BILT Meeting Minutes
“A.I.” KSA Vote and Discussion

MEETING DATE: Tuesday, March 4, 2025	MEETING TIME: 10:30am-12:30pm Eastern	MEETING PLACE: Zoom
RECORDER: Mark Dempsey	RECORDING: Available upon request	PREVIOUS MEETING: Industry trends meeting – January 24, 2025

MEMBERS PRESENT

BILT:		
Brent Blawat, Mew Resources Consulting	Jeff Magnuson, SpartanNash Company	TrainOnQ LLC
Matthew Coatney, 10 Billion Ventures	Priti Malkan, consultant	Teresa Younkin, Mosaic Life Tech
Brian Cunningham, Strategies	Milton Maddoz, US Artificial Intelligence Institute	
Rob Garretson, CareSource	Mark Richter, Hitachi Digital Services	
EDUCATORS: College of Dupage, Collin College, Columbus State Community College, Daytona State, Foothill DeAnza, Lone Star College, Northern Virginia Community College.		
NITIC STAFF: Ann Beheler, Mark Dempsey, Christina Titus, Larry McWherter (NITIC principal investigator), Alie Hernandez, Stephanie Schuler, Rajiv Malkan, Deb Hecht, Diane Meza, Andie Bonskiwske		

Agenda items	Discussion
NITIC and BILT overview	<p>Larry provided an overview of the NITIC grant, then explained that the goal of today’s meeting is to prioritize job skills for entry-level A.I. technicians. Larry told the BILT members that their anonymized feedback is widely disseminated to educators nationwide. He also noted that educators on the Zoom call are invited to listen actively; educators can post questions in the chatbox.</p> <p>Ann next provided an overview of the BILT (Business and Industry Leadership Teams). BILTs are business advisory councils that put employers in a co-leadership role. NITIC BILT members are providing feedback to colleges across the country. The BILT is driven by a structured and repeatable process that efficiently uses employer time. All of the employer feedback is shared. BILTs have two goals – aligning curriculum to workforce needs and building and deepening the relationship between educators and employers. Ann noted that the discussion after the vote is just as important as the vote itself.</p> <p>Ann explained the 1-4 voting scale. Each job skill item will receive its own individual vote.</p> <ul style="list-style-type: none"> * “4” means absolutely must be in the curriculum * “3” means should be in curriculum * “2” means it’s a “nice to have” * “1” essentially means delete <p>Abstaining is okay if an employer feels uncomfortable voting on a specific item. Today the focus is knowledge and skills; abilities – soft skills – are prioritized in a separate meeting.</p> <p>Using a QR code (and URL in the chatbox), the SMEs voted electronically.</p>

	Anything scoring 2.6 or less will be discussed because, typically, items with scores that low are left out of the curriculum.
Knowledge items	<p><u>K5 - Knowledge of CRISP-DM (Cross Industry Standard Process for Data Mining), 2.25 average</u></p> <p>One employer didn't think this was an entry-level skill. Another employer explained she voted it high, but only because data mining and understanding that lifecycle is useful "at the foundational level." General knowledge of data mining is more important than specific knowledge of CRISP. CRISP is similar to a CDMP (Certified Data Management Professionals) designation; it's an industry standard to follow. The BILT seemed to prefer making K5 more of a general knowledge area for data mining.</p> <p>ACTION: consider revising to read "knowledge of data mining" and keeping K5 on the list.</p> <p><u>K9 - Knowledge of designing data dictionaries, 2.44 average</u></p> <p>One employer who voted this low likened it to the CRISP discussion – it's a good foundational element for data literacy and competency, but those skills are captured elsewhere on the list. He noted that specific techniques and frameworks aren't as crucial as foundational basics. This seemed more like an advanced topic. Another employer suggested that while knowledge of data dictionaries is good, designing the dictionaries isn't as relevant for entry-level workers. The BILT agreed.</p> <p>ACTION: consider removing the word "designing" and keeping K9 on the list.</p> <p><u>K26 - Knowledge of identity and access management, 2.20 average</u></p> <p>One employer called this foundational. It's good to understand the technologies, but this isn't something in demand for robotic processes or business process automation toolsets. Another employer voted it low because on a list of things students should be learning in a two-year program, K26 ranks pretty low. Another employer isn't sure how IAM fits – it's a separate logic level from the infrastructure level. K26, to him, feels like something outside of the scope of entry-level AI responsibilities. Others agreed.</p> <p>ACTION: consider removing K26.</p> <p><u>K30 – Knowledge of cybersecurity assessments, 2.60 average</u></p> <p>One employer isn't sure an entry-level worker would need to know about assessments. Other employers agreed. This seems more like a security task than an AI task. Another employer said that AI is used more and more in the security domain, but admitted that if this is for a generalist AI technician K30 may not be essential. One employer noted that the trend is more toward cybersecurity AI applications. That's what students need to know: how will you apply AI to cybersecurity business needs? But that's not how K30 is reading right now. One employer noted K30 would be critical for a "cybersecurity student learning AI," but AI is now applied across all domains. AI is a feature of products. He didn't think an AI generalist getting hired at a company needs to be focused on cybersecurity. Other employers disagreed somewhat: everyone should be concerned with security. Another employer wasn't sure about the assessment part of K30. The BILT wondered if the phrase should be "risk assessments" rather than "cybersecurity assessments." Another employer looked at it from the software development perspective – he would rather someone have a security "mindset" and know what's vulnerable from the beginning. Address security during the development process and ensure it's secure from the start. Another employer mentioned ethics. When do the ethics of</p>

	<p>cybersecurity and AI overlap? Rather than teach the entire security stack, a school could cover security topics during a responsible AI course.</p> <p>From the chatbox: “I would suggest that ethics relates to how we attempt to use AI, cyber security is how we prevent others from misusing the AI we design.”</p> <p>ACTION: consider changing “cybersecurity” to “risk” and keeping K30 on the list.</p> <p><u>K31 – Knowledge of applications in Elasticsearch, 2.40 average</u> One employer voted this high, but he thinks it’s too specific. K31 should be more general. Elastic can be lexical or vector, and he believed that vector search is more common. If a student understands vector search, he/she can use a variety of products, not just Elasticsearch.</p> <p>ACTION: consider revising to “knowledge of applications of vector search” and keeping K31 on the list.</p>
Skills items	<p><u>S9 – Skill in using fixed automation, 2.38 average</u> Ann noted that with an average that low, it’s likely this topic won’t make it into curriculum. One employer voted this low because it’s such a specific technique. It felt too narrow to be applicable. He also wasn’t sure how this would work in a classroom. Other BILT members agreed that this was too specific.</p> <p>ACTION: consider removing S9 from the list.</p> <p><u>S14 – Skill in nonparametric statistics, 2.56 average</u> The BILT had concerns similar to this one as they did with S9. Too specific.</p> <p>ACTION: consider removing S14 from the list.</p> <p><u>S17 – Skill in data framing and surrogate data, 2.56 average</u> One employer found this to be too specific. S17 would be good if you’re offering specializations in a data program. For a generalist AI program, S17 is too specific. Others agreed. One employer asked if we’re making a distinction between AI and machine learning (ML). Ann said no. That employer noted that he sees a lot of ML on this list – analytics and statistics. However, the AI he is experiencing is more about how to leverage the LLMs. That goes to an earlier point from one employer who noted that this list is missing mention of generative AI. One employer noted that AI is a “big umbrella” with gen AI, ML, and modeling. This job skills list covers all sides of AI, but it may be getting too specific with manufacturing and cybersecurity. Foundational skills are more important than knowing how to perform particular tasks. He also pointed out that fun AI things students like (such as facial recognition) require a certain level of expensive hardware. That’s not always easy for a two-year program to provide. The BILT seemed to agree that programs should focus on general foundational topics, then find ways to offer focus on specific areas. Two years isn’t enough time to get someone knowledgeable about specific AI topics.</p> <p>ACTION: consider removing S17 from the list.</p> <p><u>S23 – Skill in performing SQL integrations with AWS, Azure, GCP, 2.40 average</u> One employer asked whether students should learn about databases beyond SQL. They need a strong foundation in SQL, but can they also be aware of other forms of databases? Ann</p>

noted that some programs can only provide 18-24 hours of technical content. Space is limited. This led to a brief discussion of math requirements for AI programs. Ann works with some AI programs that offer limited “just in time” math classes for AI students. “Just in time” math takes much less time because it focuses just on the topics needed for AI; standard math courses cover topics not applicable to AI.

S27 – Skill in Persona design and creation for applications, 2.60 average

One employer said this felt like an advanced ancillary skill. Another employer, however, noted that this is soon going to be very important for companies looking to implement AI. He suggested keeping this in if it’s an introduction to Persona design to teach students how to create apps. The discussion then turned to whether this is better suited as a knowledge item students should be aware of rather than a skill they need to know how to do.

ACTION: consider moving S27 to the Ks.

S31 – Skill in understanding and applying linear regression, 2.60 average

The BILT discussed making this a knowledge area as well. One employer said that his low votes in this section of the list are motivated by the fact that he’s seeing fewer and fewer people building their own AI systems and training models. Instead, companies are just using the models that have already been built. Others agreed.

ACTION: consider moving S31 to the Ks.

S32 – Skill in developing Bayesian statistical models, 2.56 average

Employers wanted to make this a knowledge area as well. One employer said that students from a two-year program should know how to leverage models, not build them.

ACTION: consider moving S32 to the Ks.

S38 – Skill in K-means clustering, 2.57 average

S39 – Skill in building computer vision with software such as AWS, Azure, GCP, 2.33 average

Same argument as with S31 and S32.

ACTION: consider moving S38 and S39 to the Ks.

S41 – Skill in developing and applying Supervised Learning Applications, 3.11 average

S42 – Skill in developing and applying Unsupervised Learning Applications, 3.20 average

S43 – Skill in developing and applying Applications of Neural Networks, 2.90 average

S44 – Skill in developing and applying Artificial Neural Networks (ANN), 2.90 average

S45 – Skill in developing and applying Convolution Neural Networks (CNN), 2.60 average

S46 – Skill in developing and applying Recurrent Neural Networks (RNN), 2.70 average

One employer explained that it depends on the activities the students are performing. Are they in a class dedicated to training models? Or are they learning more general foundational skills? These skills are essential, but if the context is more general applications, all of these should be moved to the knowledge area Ks on the list. Employers agreed.

ACTION: consider moving S41, S42, S43, S44, S45, and S46 to the Ks.

S50 – Skill in using facial recognition, 2.22 average

	<p>One employer noted ethical concerns related to facial recognition in terms of bias and discrimination. Many companies are pulling back from facial recognition development. As a result, there is less of a need for this skill. Companies are not building these models. This may be better suited as a topic in an AI ethics course. Another employer noted that facial recognition is not much different than any other image recognition, so there's no reason to call this one out.</p> <p>ACTION: consider removing S50 from the list.</p>
<p>What should be included next time</p>	<p>Ann next asked the employer what items were missing from the list and should be included for the next job skills vote.</p> <p>The BILT agreed that generative AI knowledge and skills – including prompt engineering – should be included next time.</p> <p>One employer said that students need to understand that dedicated hardware components are needed to run AI models. This understanding needs to include specifics about the “cost of technology requirements,” such as how power and cooling needs are quadrupled when using GPU hardware. These are all critical considerations in designing data centers. There is a cost to running GPUs, which are very “power hungry.”</p> <p>Another employer mentioned the importance of students knowing how to “test the truthfulness” of AI responses. That is, learning how to ground the language model to factual information and then verifying the responses. You don't want the LLM to hallucinate.</p> <p>One employer mentioned agentic AI systems. This is still very new, so it's more about students being aware of the trend and that it's coming. This led another employer to wonder if students need to know about AI architectures, but the group seemed to think that was too advanced.</p> <p>Ann noted that it's fortunate that there's so much courseware available to schools in AI. Finding someone to teach AI content in community colleges can be very difficult.</p> <p>One employer explained that, in general, his preference with new hires is finding people with “good breadth.” He's less concerned with depth. He'd rather have someone with a passing knowledge of most of these than someone who's spent massive time learning neural networks. He wants a generalist. He can train the new hire on the specifics. The other BILT members agreed. One noted that this approach “future proofs” the list. Another said that with AI evolving so quickly, having someone adaptable and a self-learner is essential. In five years, this list will likely be obsolete.</p> <p>This led to a more extensive discussion of generalist versus specialist. Some companies want one; some want the other. Schools will need to decide what kind of program they want to offer. If you teach a generalist program, you'll have to pick and choose items off this list. Ann wondered how much of specialist skills would be covered by on-the-job training. Another employer asked if it's reasonable for a two-year program to teach an AI specialist degree. Ann knows of AI programs that feature applications and capstones that allow students to focus on more specialized areas.</p> <p>Another employer said one specialist skill that will continue to be in demand – and can be taught in two years – is generative AI and prompt engineering. How do you build and incorporate LLM like a Swiss Army knife to address business problems? The other BILT</p>

members agreed. The workforce has changed. Jobs that had been focused on knowing how to create models are now focused on leveraging those models using the correct command prompts. Another employer agreed – you can do a lot with broad knowledge, but the company often upskills employees and creates the specialist in-house.

Ann noted that with AI, new graduates are competing with incumbent older adult workers.

Another employer in the software development space sees AI as an “augmentation.” It’s okay for AI to be its own program, but he’s **implementing and leveraging AI in software development**. He’s especially interested in agentic AI, which will require critical thinking. He noted the truism that if someone can learn independently, you can teach them anything.

Larry asked about cybersecurity. Should we add specific cybersecurity elements to the list? While one employer said security concerns differ based on whether it’s a public platform (which has more needs and vulnerabilities) or an internal platform, another argued that both sides should be treated the same. The first employer countered that most organizations have finite budgets and have to make choices, but he agreed that architecturally both sides should be designed the same way. Everyone seemed to agree that foundational security topics should be added across all AI classes. One employer noted that a fully dedicated solely to cybersecurity is probably too much for an AI program. Another thought was that security topics were adequately covered in the existing job skills list.

One employer asked about programming and the **debate between C#, Python, and Javascript**. Only Python is mentioned on this list. Another employer explained that, to him, Python is good for tying together the ML side of things; Javascript is good for the user experience, and .NET C# is good for data movement.

From the chatbox: “C#: Best for Windows apps, Unity game development, etc.; Python: Ideal for data analysis, AI, and rapid prototyping; JavaScript: Essential for web front-end development.”

Ann noted that programs cannot teach all of those. One employer said they see more developers use C#; they don’t want to retrain on Python. Another employer said that at a law firm he works with, everything is developed in .NET, but the early work is in Python. So, while they’re limited in working with bleeding-edge options, once it reaches enterprise readiness, it moves into Azure and C#. One employer stated that Python is the default language of AI. It’s a firm requirement. Others agreed. But if someone is coming in with a skill set of C#, they can easily apply it to AI. Another employer restated the situation: this is an item to watch but not one to act on. Many .NET developers want to work in AI but want to do it in C# even though Python is defacto across all AI platforms. Ann noted that Python also has many libraries, and teaching it doesn’t require many pre-requisites.

Larry asked **what soft skills might be critical for AI**. The employers mentioned adaptability, lifelong learning, critical thinking, and curiosity. One employer called out conceptualization – if you can’t see it, you can’t build it. Another referenced “motivational interviewing.” Can you gather business-driven requirements from stakeholders? They can tell you what they want, but you have to tease out the exact needs rather than a laundry list of features. Another employer told the joke about a business owner receiving an application and saying, “It’s what I asked for, but not what I wanted.” Another mentioned communication skills. Students need to learn how to listen and organize their thoughts. Technical people like to argue and prove that someone else is wrong.

	<p>One employer asked about the barrier of posted jobs requiring four-year degrees. This is a concern, but it's starting to change in AI. Larry noted that this has been an issue in cybersecurity for a long time. One employer stated they hire people with all sorts of degrees (or no degrees), choosing to focus on the skills and experiences of the applicants. This has been a significant benefit to the company. Another said that degrees are necessary, but they're trying to look for more knowledge and skills rather than schooling.</p> <p>As for using AI on the job, one employer noted that many organizations haven't yet figured out their AI policies. Specifically, companies will need to start teaching ethical, responsible AI use as part of the new-hire onboarding.</p> <p>From the chatbox: "My approach to hiring is based on skills and abilities. seeing a degree to me suggests the ability and determination to 'complete'. I've rejected many 4 and 6-year degrees if they could not communicate actual hands-on skills. if you can graduate a 2-year who can walk in and start working, I'd consider them for a junior role."</p>
Conclusion	<p>Larry announced the next two BILT meetings. A cross-disciplinary workforce trends meeting is set for Tuesday, April 29. Later in the summer, NITIC will host a job skills prioritization and discussion meeting on entry-level software development job skills.</p> <p>Outcomes from these BILT meetings (minutes, prioritization worksheets, and trends summaries) are available for free here: https://www.nitic.org/industry/national-bilt/bilt-overview/</p>
<p>Next Meeting: Tuesday, April 29 (9:30am-10:30am Central/10:30am-11:30am Eastern) – IT industry trends</p>	