Technology Readiness Levels for Randomized Bin Picking

Jeremy A. Marvel
National Institute of Standards and Technology
100 Bureau Drive, MS 8230
Gaithersburg, MD 20899
+1 (301) 975-4592
jeremy.marvel@nist.gov

Roger Eastman
Loyola University, Maryland
Department of Computer Science
4501 North Charles Street
Baltimore, MD 21210
+1 (410) 617-2281
reastman@loyola.edu

Geraldine Cheok
National Institute of Standards and Technology
100 Bureau Drive, MS 8230
Gaithersburg, MD 20899
+1 (301) 975-6074
cheok@nist.gov

Kamel Saidi
National Institute of Standards and Technology
100 Bureau Drive, MS 8230
Gaithersburg, MD 20899
+1 (301) 975-6069
kamel.saidi@nist.gov

Tsai Hong
National Institute of Standards and Technology
100 Bureau Drive, MS 8230
Gaithersburg, MD 20899
+1 (301) 975-3444
hongt@nist.gov

Elena Messina
National Institute of Standards and Technology
100 Bureau Drive, MS 8230
Gaithersburg, MD 20899
+1 (301) 975-2661
elena.messina@nist.gov

ABSTRACT
A proposal for the utilization of Technology Readiness Levels to the application of unstructured bin picking is discussed. A special session was held during the 2012 Performance Metrics for Intelligent Systems workshop to discuss the challenges and opportunities associated with the bin picking problem, and to identify the potentials for applying an industry-wide standardized assessment and reporting framework such as Technology Readiness Levels to bin picking. Representative experts from government, academia, and industry were assembled to form a special panel to share their insights into the challenge.

Categories and Subject Descriptors
C.4 [**Performance of Systems**]: Performance Attributes; I.5.4 [**Applications**]: Computer Vision

General Terms
Measurement, Documentation, Performance, Experimentation, Verification

Keywords
Bin Picking, Technology Readiness

1. INTRODUCTION
Manufacturing technologies have witnessed a veritable boom in robot integration and improved sensing modalities for safety and task automation. Worldwide manufacturing initiatives stress the integration of robot technologies in modernized manufacturing facilities, and push the boundaries of both productivity and innovation in an ever-increasingly competitive market.

Despite years of considerable progress in 3D pose estimation systems and vision-guided robotics, one of the greatest challenges to manufacturing automation is the task of component acquisition from a randomized bin of parts. A special session was held at the 2012 Performance Metrics for Intelligent Systems workshop that focused on the state of the art and metrics of technology readiness levels (TRLs) for bin picking solutions that are robust against random pose and part variations. We addressed the indicators of maturity of approaches for overcoming shape variation, pose and orientation uncertainty, weak or no distinguishing image features, and limited grasping options. Presenters discussed both the TRL development process and the needs and challenges from the perspectives of both users and vendors regarding bin picking for manufacturing automation.

The principal goal of the special session was to establish a common understanding of how to match the robotic bin picking perception requirements of manufacturers against the current capabilities of vendor systems. Further, we intended to determine the best mechanisms for advancing the capabilities and greater deployment of robotic bin picking. This could be through an advanced perception TRL framework or other common set of metrics and evaluation criteria that can be developed by the user, vendor, research, and government communities through a consensus standardization process.

We discussed the requirements and processes involved with the grading of different levels of bin picking difficulty, and the feasibility of establishing a set of standardized artifacts for bin picking solution validation. Additional topics of discussion included the challenges inhibiting solution integration, and opportunities for advancement in next-generation manufacturing environments.

This report provides an account of the proceedings of the 2012 Performance Metrics for Intelligent Systems (PerMIS) workshop special session, and outlines preliminary action items for the development of a process for evaluating and documenting the maturity of technologies for bin picking. Section 2 presents an overview of the bin picking problem, and discusses the challenges

Disclaimer: Certain trade names and company products are mentioned in the text or identified in certain illustrations. In no case does such an identification imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products are necessarily the best available for the purpose.

and measurable properties of its use. In Section 3, we provide a summary of TRLs and discuss their applicability (and that of other technological maturity assessment scales) to bin picking. And Section 4 outlines the discussion topics from the special session regarding the use of TRLs within the bin picking problem domain.

2. BIN PICKING

The application of bin picking in manufacturing is an interesting problem in that the concepts are fairly ubiquitous and practically everyone understands the underlying concepts of acquiring objects from a bin of parts, but also that there are no consistent definitions of what the process actually entails. A major contributing factor to this ambiguity is the observation that there is no single bin picking application, but rather a spectrum of specific instantiations based on any number of unique constraints that are user-, process-, and product-specific.

In the classic literature, the process of bin picking has been frequently reduced to a three-step process consisting of isolating a specific object from the background, determining the pose of that object, and then creating a path trajectory to move a robot in and grasp the object (e.g., [1-3]). The base definition is necessarily vague given the broad spectrum of bin picking applications, and thus includes parts acquisition processes ranging from picking objects off a conveyor belt (though many would argue that a defining aspect of bin picking is extracting parts within a box or bin) to taking parts out of a randomized bin of multiple part shapes. The acquisition of a single part from a collection of parts is often considered an integral component of manufacturing, and can be considered a superset of many common industrial processes including kitting, palletizing, packaging, and assembly. Successfully integrating bin picking into a product line brings a number of inherent benefits, including higher throughput, reliability, and flexibility, but first the challenges of each particular picking scenario must be overcome.

When attempting to assess how difficult a particular picking problem is, it is helpful to have a common comparative scale. Just as there is no single bin picking problem, however, so, too, is there no single comparative metric for determining the complexity of a given problem. A common—and arguably over-simplistic—method comes from the Electrical Engineering Handbook, and utilizes a relative three-tier difficulty rating that assigns a complexity value based on the controllability of part position and appearance [4]. More recently, however, researchers have begun assessing problem difficulty based on the maturity of the component technologies required to address a particular bin picking application [5].

There are a number of categorical parameter spaces by which the difficulty of a bin picking application can be scaled. Commonly this spectrum is scaled according to the degrees of freedom of the parts to be acquired, i.e., ranging from X and Y axes variation ("2D"), to X and Y axes variation plus Z axis rotation ("2.5D"), to full X, Y, and Z position and rotation variation ("6D"). However, more recently, trends in describing the bin picking application have separated solutions for the problem domain according to image segmentation difficulty properties such as image feature strength [6]. There are a number of categorical parameter spaces by which the difficulty of a bin picking application can be scaled. These spaces include scenario complexity, part location or orientation, part or shape variation, image feature strength, part rigidity, and part overlap and interlock (i.e., when two or more parts become connected and require separation before they can be used).

Once a solution to a bin picking problem has been developed, there are three principal performance metrics by which that solution can be evaluated: speed, efficiency, and accuracy. Speed refers both to the time required to acquire an individual part from a bin (*picking time*), and the number of picks per given period of time (*bandwidth*). Efficiency is measured in terms of time utility (e.g., the time spent searching for parts to acquire versus the time actually spent picking them up), grasping quality and acquisition success, and robot trajectory optimization (i.e., how efficiently does the manipulator move into the bin and avoid collisions with parts or the bin?). Accuracy is the measurement error in object recognition and part pose estimation.

There are three primary challenge domains that may complicate the integration of a bin picking solution into the manufacturing process. The first domain, sensing, includes the inherent difficulties in sensor and algorithm development, but also includes components of process and workcell optimization. The types of challenges that an integrator must overcome include object identification issues due to lighting variations caused by surface reflectivity, shadows, and material transparency. Pose estimation algorithms may be further misled by shape and surface variations incurred during the manufacturing process, or by weak, inconsistent, or non-existent image features. Each effectively prevents an adequate fit of the detected part to a known model. Moreover, variations in the bin itself—such as position uncertainty and bin damage—may present additional challenges if the system does not know exactly where it should be searching for parts.

The second challenge domain reflects issues with the hardware involved, including the robot, the gripper, and the parts being acquired. Specifically, the robot's dexterity and reach may limit the number of parts that can actually be acquired. Challenges with the gripper's dexterity and design may restrict the number of possible grasp points as well as limit the grasp efficiency and quality. Similarly, the weight, durability, and separation of the parts may further restrict how they can be handled.

The third (and arguably most difficult) challenge domain to overcome includes the pragmatic issues of bin picking solution integration. This includes considerations such as cost, which is defined in terms of both financial burden and the times required to bring a system online, to train and tune the system for new parts, and to support the repurposing of an existing system for a new process. Further, issues concerning the bin picking problem application's uniqueness are often considerable. For example, when introducing a new part to the production line, what solution components can be recycled, and how well does the new solution actually fit the specific bin picking need? Conversely, can the new process be changed to be more congruent with the old bin picking system? Many times the old system must be shelved, and a new system built up from initial concepts. These considerations ultimately tie in to the understanding of the bin picking problem, itself: does the integrator know and understand the process well enough to be able to identify reusable components? Moreover, what level of understanding and awareness does the user have about bin picking in general? This final element is frequently characterized by users either not knowing what solutions are available, or having unrealistic expectations of the capabilities of robotic bin picking systems.

Table 1. Example TRL Description Summaries Based on the NASA [8], DOD [9], and DOE [10] Guidelines

TRL	Summary and Description
1	Basic principles observed and reported. Research begins to be translated into applied research and development (R&D)
2	Technology concept or application formulated. Practical (albeit speculative) applications can be invented after basic principles are observed.
3	Characteristic proof of concept. Active R&D is initiated, and includes analytical and lab studies for physical validation of analytical predictions of individual elements of the technology.
4	Laboratory validation of components. Basic technological components are integrated to verify they work together.
5	Target environment validation of components. Higher fidelity of component integration testing in a reasonably supporting environment to allow for simulated environment testing.
6	System/subsystem model in target environment. Models and prototypes demonstrating a significant technological readiness improvement are tested in a relevant environment.
7	System prototype in operational environment. Functional prototypes demonstrating the completed system in its approximate expected configuration are evaluated in an operational environment.
8	Final system qualified through demonstration. Technologies are proven to work in their final form and under expected conditions through test and demonstration.
9	Final system proven through vetting. Applications of technologies in their final form are proven through successful operations under mission conditions.

3. TECHNOLOGY READINESS LEVELS

Originally proposed by the National Aeronautics and Space Administration (NASA) [7, 8], the TRL structure describes a process for evaluating the maturity of technologies prior to their incorporation into deployable systems. The primary users of TRL scales are agencies and organizations, both domestic and international, with aeronautical and aerospace interests, but many users modify the language of NASA's TRL model to better suit differences in the user's production patterns, technologies, or management structures (e.g., the U.S. Department of Defense (DOD, [9]), and U.S. Department of Energy (DOE, [10]).

TRLs are used to measure maturity of technologies when determining the risks associated with inserting them into a mission (or mission component), and are critical to communication with partners, suppliers, and customers. The TRL structure is frequently implemented as a nine-stage hierarchy, as illustrated in Table 1. Generally speaking, TRL-6 is a desirable stage prior to any technology being integrated into a mission, and is considered the "go/no go" point. TRLs, however, are only one of several tools for the decision process. Key Decision Points (KDPs), for example, determine the readiness of a program/project to advance to the next phase, and are outlined in NASA's Procedural Requirements [11].

Despite its wide utilization in aerospace and aeronautics both within the U.S. and internationally, there is no standardized TRL structure or implementation. As a result, the TRL for a specified technology may not be identical for all missions or applications. Specifically, the readiness level for a given technology may be different depending on the considering agency, environmental factors, intended use, or even who within a given agency is assessing the technology. Similarly, there is a significant lack of clear exit criteria (i.e., conditions for moving from one level to another) for higher TRLs, and the guidelines for assessing TRLs are frequently vague or even conflicting.

Applying TRLs to new problem domains such as manufacturing is complicated by the TRL structure's inability to handle certain factors that are important to these domains. For instance, though the TRL structure can readily be applied to the manufacturing domain, it does not address the requisite factors of throughput, profit, market needs, or the ease of labor and implementation issues. Because TRLs are typically applied to one-off or otherwise relatively small-scale production, applications requiring large-scale production or distribution are often incompatible. Instead, focus is placed on technological maturity, and consideration of factors such as the capabilities of processes or technologies would not be addressed using the current TRL structure.

As an alternative, the U.S. Department of Defense (DOD) introduced Manufacturing Readiness Levels (MRLs) [12], a 10-level administrative process focused on the actual production process. MRLs are used to quantitatively assess the maturity of technology components from a manufacturing perspective, and are used to determine the risks involved with bringing products to the production phase. This process involves an initial assessment of the basic needs for manufacturing products, and is used to document and demonstrate that given technologies are ready for wide scale manufacturing.

These deficiencies have thus prompted efforts to reassess the TRL structure. NASA, for instance, is reevaluating its TRL definitions and exit criteria, and efforts are being considered to create standards for assessing and reporting TRLs. Beyond NASA, the International Organization for Standardization (ISO) is coordinating space agencies and other stakeholders to develop an international TRL standard (ISO TRL work group, 14N665, *Definition of Technology Readiness Levels and their criteria of assessment*). Through this effort, ISO is also discussing the necessary steps to broaden the scope of the standard beyond aerospace, eventually encompassing other topics such as manufacturing.

4. PANEL DISCUSSION

Following presentations on TRLs and opportunities in bin picking by Karen McNamara from NASA, and Jeremy Marvel from the National Institute of Standards and Technology (NIST), respectively, a special panel of experts from government, industry, and academia was assembled to address the challenge of assessing and reporting technologies for addressing the bin picking problem domain. Alphabetically, these panel members were:

- Bob Bollinger, Procter & Gamble (P&G)

- Paul Evans, Southwest Research Institute (SWRI)
- Joyce Guthrie, United States Postal Service (USPS)
- Eric Hershberger, Cognex
- Carlos Martinez, ASEA Brown Boveri (ABB)
- Karen McNamara, NASA
- James Wells, General Motors (GM)

Roger Eastman from Loyola University, Maryland, moderated the discussion, and prompted dialogs based on topics relating to the development, utilization, and assessment of bin picking solutions.

The discussion began with an effort to expand the categorical classification of the user's perspective of bin picking. From a manufacturing perspective, there are three distinct and readily identifiable phases for which bin picking will be employed based on the stage of production in which the objects are being picked. As the manufacturing process nears a finished product, the level of care required to prevent damage increases. Early stages, for example, typically require the acquisition of raw (unfinished) materials frequently presented in randomized bins. In contrast, in-process and finished components require increasing levels of fixturing to prevent damage that would affect the functionality or aesthetics of the parts. The bin picking process varies accordingly based on the shipping or presentation method.

Improved inter-process component transfer is an impetus for production optimization, and the ability to handle material in a lean fashion is what is driving bin picking. One of the panelists described the production process as a series of transformations in which the components are transferred between robots, hoppers, bins, conveyor belts, dunnage, and so on. Intermediate transformation steps, e.g., moving parts from a hopper to a conveyor belt to be acquired by delta robots, add cost and complexity to the manufacturing process. The capacity for handling parts as they would naturally be presented in an unstructured form—particularly if the gripper does not have to be changed or the robot reprogrammed to handle the part changes—would thus improve process efficiency.

The distinction between structured and unstructured (i.e., random) presentation of parts within a bin plays a vital role in determining the complexity of the problem. As the strictness of fixturing decreases, the difficulty inherent in developing a bin picking solution increases. Structured bin picking (i.e., parts presented in known, repeatable positions and orientations) is largely considered to be a solved problem, and is addressed by simple matrix handling. In contrast, no general approach for addressing unstructured bin picking (where the locations, shapes, and identities of parts may not be known *a priori*) has been produced.

The degree of randomization of the parts within the bin thus contributes to complexity. For example, a bin full of cast parts is considered to be an easier problem than a bin of irregularly shaped mail. Solutions to such problems have not been forthcoming, and some solution providers have enacted policies to decline requests for unstructured bin picking. Despite years of research in algorithms, robotics, and sensor systems, no unstructured bin picking solution has been developed that is reliable, small, cost effective, or widely applicable. Even within classes of parts (e.g., plastic container caps), the required flexibility of bin picking solutions has not materialized, and the capacity to compensate for product line changes requires hard automation (i.e., large, highly-fixtured, part-specific feeder and handler systems). The issue is further complicated by cases where such hard automation is impossible due to large variances in part shape and size.

In contrast with the hard automation solutions, the cost for robot bin picking solutions is not driven by the cost of the robot. Rather, it is the cost of integrating the robot into the manufacturing process that presents the largest hurdle. Specifically, handling safety and process-specific ancillary assembly line system requirements contribute the most to the price of the system, and thus hinder cost efficiency and flexibility. Specialized fixturing and dunnage to ease the burden on perception add additional cost to the system, and must be redesigned or repurposed as the products and processes change. The actual cost of the robot is comparatively small, as is the impact of the robot on the complexity of the bin picking solution. Though different bin picking classifications may require different robots, the control, repeatability, and reliability of robots in general are considered largely solved. Similarly, the gripping of the objects for process utilization, though considered a specialized component given the parts being acquired and subsequent utilization, is also considered solved.

If the physical aspects of the bin picking problem are considered solved, then what is the greatest hindrance? The panel agreed that perception (and associated sensing technologies) of the various components in the manufacturing setup is the limiting factor in the improvement of bin picking. For example, the USPS already has the technology to handle packages once they have been acquired, but reiterated that perceiving the locality of the materials as they come in presents the greatest barrier to full automation.

Similarly, the bin itself provides a challenge in a number of ways. Identifying variations of the bin in terms of placement, shape, and condition (e.g., due to incurred damage to the bin) add complexity to both the part location process and to collision-free trajectory generation. Recognizing when the bin is empty is a common challenge, as missed parts at the bottom of the bin lead to waste, and, in terms of mail delivery, loss of business functionality and reliability. Once a part has been acquired, if the robot needs to control or attach the part to a fixture or another part, the system will need to know exactly how the part is being held, which requires additional perception capabilities for process validation.

Another common theme expressed by the panel members was the desire to have robots and humans working collaboratively on the production line. This functionality requires an extension of the perception capabilities of the workcell to include robot safety, for which the panel discussed improved situational awareness of the workcell integrating multiple sensors and algorithms from multiple vendors. An additional consideration included a fundamental reconsideration of the requirements of the workcell, and a redesign of process components (e.g., the bins containing the parts) such that they are robot friendly rather than requiring robots to work within the confines of human accessibility.

The second part of the panel discussion focused on whether the development of an evaluative maturity measurement process like TRLs would aid in the advancement of bin picking technologies. Most of the larger manufacturers have internal processes similar to TRLs that they use to measure the maturity of technologies prior to integration. One company, for instance, has a management-integrated process for evaluating required technologies (e.g., technologies necessary for the design of a new car) and technologies that improve existing processes. Ultimately, the technology evaluation process is merely an input into the decision process and is not a goal in and of itself.

These internal processes, however, are typically proprietary, or otherwise unavailable for other users to utilize as either an example or as a means of benefiting from the larger manufacturers' experiences. The question was thus raised of the panel: how can small companies learn from the experiences of larger companies; what reporting processes other than TRLs are available? As an alternative, it was suggested that a new standardized test method or generalized competition format could be used in lieu of the application- and user-specific technology maturity scales. It was further suggested that trade organizations such as the Robotic Industries Association may be able to provide aggregated abstractions of the technological knowledge for dissemination.

When discussing the metrics by which different technologies could be evaluated, a number of metrics were suggested as being common to users of bin picking solutions. Beyond the expected metrics such as picking speed and throughput discussed in Section 2, the panel also recommended measurement concepts such as agility and repurposing. Agility is the capacity of a robot working with product A to quickly re-task to begin working with product B, and repurposing refers to the amount of time, effort, and skill required to have a robot perform a different task.

5. CONCLUSIONS AND ACTION ITEMS

In the TRL for Randomized Bin Picking special session of the 2012 PerMIS workshop, a panel of experts was organized to discuss the needs and challenges of unstructured bin picking, and to assess whether a TRL structure would help facilitate the documentation and advancement of bin picking technologies. The panel agreed that structured bin picking—situations in which objects are presented in a regular matrix such that parts acquisition requires little to no actual perception to locate a particular object—has been largely solved with a comparatively high level of maturity. In contrast, unstructured bin picking—situations in which presented objects have inconsistent or unknown pose or shape—is considered an immature technology, and that some form of communication structure is needed to help unite the research community in order to fully address the problem

It was also agreed that the creation of some form of taxonomy for assessing and documenting the technological readiness of core processes and technologies would greatly benefit their integration and application in manufacturing practices. However, it was not certain that the TRL structure is necessarily the best approach for describing maturities of application-targeted manufacturing technologies. It was recommended that future efforts attempt to identify the full spectrum of alternatives in order to discover the one that is best for capturing the problem domain.

Particular to the domain of manufacturing, the panel decided the logical next step in addressing the challenges of unstructured bin picking was to first assess the current state of the art in picking technologies. Two action items were thus discussed. The first was to form a task group to identify, create, and document metrics and test methods for evaluating bin picking solutions. This process would include, but is not limited to, the development of standardized artifacts and data sets, performance evaluation frameworks, and a standardized lexicon of bin picking metrics. The second action item involves the documentation of available technologies (including sensing, perception, trajectory creation, and grasping) and categorically assessing their capabilities as applied to the bin picking problem domain.

REFERENCES

[1] Ikeuchi, K., Horn, B.K.P., Nagata, S., Callahan, T., and Fein, O. Picking Up an Object from a Pile of Objects. In *MIT AI Memos*. 1983. Pp. 139-166.

[2] Fuchs, S., Haddadin, S., Keller, M., Parusel, S., Kolb, A., and Suppa, M. Cooperative Bin-Picking with Time-of-Flight Camera and Impedance Controlled DLR Lightweight Robot III. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2010. Pp. 4862-4867.

[3] Skotheim, Ø., Thielemann, J.T., Berge, A., and Sommerfelt, A. Robust 3D Object Localization and Pose Estimation for Random Bin Picking with the 3DMaMa Algorithm. In *Proceedings of the SPIE Three-Dimensional Image Processing and Applications*. Vol. 7526. 2010. Pp. 75260E-75260E-11.

[4] Dorf, R.C., Ed. *The Electrical Engineering Handbook*. 2nd Edition. Boca Raton: CRC Press LLC. 1997.

[5] Shafi, A. Bin Picking: Definitions, Success Rates and New Solutions with Random Positions, Object Shape Variation and Weak Imaging Features. In *Proceedings of the International Conference for Vision Guided Robotics*. 2011.

[6] Oh, J.-K., Baek, K.K., Kim, D., and Lee, S. Development of Structured Light based Bin Picking System Using Primitive Models. In *Proceedings of the IEEE International Symposium on Assembly and Manufacturing*. 2009. Pp. 46-52.

[7] Sadin, S.R., Pvinelli, F.P., and Rosen, R. The NASA Technology Push Towards Future Space Mission Systems. In *Acta Astronautica*. Vol. 20. 1989. Pp. 73-77.

[8] Mankins, J.C. Technology Readiness Levels: A White Paper. NASA, Office of Space Access and Technology. 1995.

[9] U.S. Department of Defense. *2012 Defense Acquisition Guidebook*. http://dap.dau.mil.

[10] U.S. Department of Energy. *DOE G 413.3-4, U.S. Department of Energy Technology Readiness Assessment Guide*. 2009.

[11] U.S. National Aeronautics and Space Administration. *NASA Procedural Requirements NPR7120.5D: NASA Space Flight Program and Project Management Requirements*. 2010.

[12] U.S. Department of Defense. *DoD Manufacturing Readiness Level Deskbook*, V2.01. July, 2011. http://www.dodmrl.com.