OXFORD

SOT | Society of Toxicology
www.toxsci.oxfordjournals.org

TOXICOLOGICAL SCIENCES, 142(2), 2014, 312–320

doi: 10.1093/toxsci/kfu199
FORUM

FORUM

Adverse Outcome Pathway (AOP) Development I: Strategies and Principles

Daniel L. Villeneuve*[1], Doug Crump[†], Natàlia Garcia-Reyero[‡], Markus Hecker[§], Thomas H. Hutchinson[¶], Carlie A. LaLone*[,||], Brigitte Landesmann[|||], Teresa Lettieri[|||], Sharon Munn[|||], Malgorzata Nepelska[|||], Mary Ann Ottinger[||||], Lucia Vergauwen[#], and Maurice Whelan[|||]

*US EPA Mid-Continent Ecology Division, 6201 Congdon Blvd, Duluth, Minnesota 55804, [†]Environment Canada, Ecotoxicology and Wildlife Health Division, Ottawa, Ontario, K1A 0H3 Canada, [‡]Mississippi State University, Institute for Genomics, Biocomputing and Biotechnology, Starkville, Mississippi 39762, [§]University of Saskatchewan, School of the Environment and Sustainability and Toxicology Centre, Saskatoon, Saskatchewan, SK S7N 5B3, Canada, [¶]University of Plymouth, School of Biological Sciences, Plymouth, Devon, PL4 8AA, UK, [||]University of Minnesota, Water Resources Center, St. Paul, Minnesota 55108, [|||]European Commission, Joint Research Centre, Via E. Fermi 2749, 21027 Ispra, Italy, [||||]Department of Biology and Biochemistry, University of Houston, Houston, Texas, 77004, [#]Zebrafishlab, Veterinary Physiology and Biochemistry, Department of Veterinary Sciences, University of Antwerp, Universiteitsplein 1, 2610 Wilrijk, Belgium

[1]To whom correspondence should be addressed at U.S. EPA Mid-Continent Ecology Division, 6201 Congdon Blvd, Duluth, MN 55804-2595. Fax: (218) 529-5003. E-mail: villeneuve.dan@epa.gov.

ABSTRACT

An adverse outcome pathway (AOP) is a conceptual framework that organizes existing knowledge concerning biologically plausible, and empirically supported, links between molecular-level perturbation of a biological system and an adverse outcome at a level of biological organization of regulatory relevance. Systematic organization of information into AOP frameworks has potential to improve regulatory decision-making through greater integration and more meaningful use of mechanistic data. However, for the scientific community to collectively develop a useful AOP knowledgebase that encompasses toxicological contexts of concern to human health and ecological risk assessment, it is critical that AOPs be developed in accordance with a consistent set of core principles. Based on the experiences and scientific discourse among a group of AOP practitioners, we propose a set of five fundamental principles that guide AOP development: (1) AOPs are not chemical specific; (2) AOPs are modular and composed of reusable components—notably key events (KEs) and key event relationships (KERs); (3) an individual AOP, composed of a single sequence of KEs and KERs, is a pragmatic unit of AOP development and evaluation; (4) networks composed of multiple AOPs that share common KEs and KERs are likely to be the functional unit of prediction for most real-world scenarios; and (5) AOPs are living documents that will evolve over time as new knowledge is generated. The goal of the present article was to introduce some strategies for AOP development and detail the rationale behind these 5 key principles. Consideration of these principles addresses many of the current uncertainties regarding the AOP framework and its application and is intended to foster greater consistency in AOP development.

Key words: adverse outcome pathway; regulatory toxicology; predictive toxicology; extrapolation; knowledgebase

Regulatory toxicology in the twenty-first century faces numerous challenges such as the need for assessing an ever increasing number of chemicals to meet new legislative mandates, while reducing animal use, costs, and time required for chemical testing (Bradbury *et al.*, 2004; Krewski *et al.*, 2010). One concept that has been proposed to aid in addressing these challenges and the resulting regulatory needs is that of the adverse outcome pathway (AOP) (Ankley *et al.*, 2010). An AOP is a conceptual framework for organizing existing knowledge concerning the predictive and/or causal linkages (termed key event relationships; KERs) between measureable/observable biological changes that are essential (termed key events; KEs) to the progression from a molecular initiating event (MIE) to an adverse outcome (AO) considered relevant to regulatory decision making (Ankley *et al.*, 2010; Fig. 1 and Table 1). With respect to the aim of describing a series of measureable biological events that are critical to the induction and progression of a toxicological response and supported by robust weight of evidence, the AOP framework is conceptually synonymous with the Mode of Action framework developed for analyzing the relevance of toxicological effects observed in animals to human health risk assessment (Boobis *et al.*, 2006, 2008). Because AOPs are intended specifically to support regulatory decision making, it is important that the KEs be extended to an endpoint of regulatory significance (i.e., an AO). Likewise, because the concept and term AOP emerged, in part, from the quantitative structure activity relationship (QSAR)-development community (OECD, 2011) extending the relationships to the point of molecular interaction between a chemical stressor and a target biomolecule (MIE) was also viewed as a priority, to aid chemical category formation. Thus, an ideal AOP was conceptualized as a series of KEs and KERs that would link an MIE to an AO.

The desire to understand and describe toxicological modes and mechanisms of action is not new. However, following publication of a seminal National Research Council report on toxicity testing in the twenty-first century (Krewski *et al.*, 2010), there has been increasing focus on the need to define and enhance the relevance of measures of the initiation or early progression of toxicological insults for regulatory decision making. The intent is to make more effective use of mechanistic data, particularly those data that can be generated more rapidly and cost-effectively than apical outcomes measured in typical whole-organism guideline toxicity tests. Organization of existing knowledge into AOP descriptions provides a systematic and transparent assembly of the evidence that supports extrapolation from initial or intermediate measures of biological perturbation to useful predictions of potential hazard. With this in mind, the ultimate goal of AOP development is to describe key building blocks, KEs and KERs linking an MIE and AO, in sufficient detail to support the application of a wide range of mechanistically based data in risk assessment and regulatory decision making.

In 2012, the Organization for Economic Cooperation and Development (OECD) launched an international AOP development program (http://www.oecd.org/env/ehs/testing/adverse-outcome-pathways-molecular-screening-and-toxicogenomics.htm, Accessed October 9, 2014). In 2013, a *Guidance Document on Developing and Assessing Adverse Outcome Pathways* was published (OECD, 2013). A major purpose of the guidance was to introduce some standardization and rigor into what had been a largely *ad hoc* process of AOP development and description. The intent was to assure that AOP descriptions included the information required to facilitate assessment of the types of measurements and weight-of-evidence supporting an AOP (Becker *et al.*—https://aopkb.org/saop/workshops/somma.html, Accessed October 9, 2014) and the AOP's consequent suitability (or lack thereof) for various regulatory applications (Perkins *et al.*—https://aopkb.org/saop/workshops/somma.html, Accessed October 9, 2014). It was recognized that the guidance was a starting point that would require updating as more experience with the development and application of AOPs was gained.

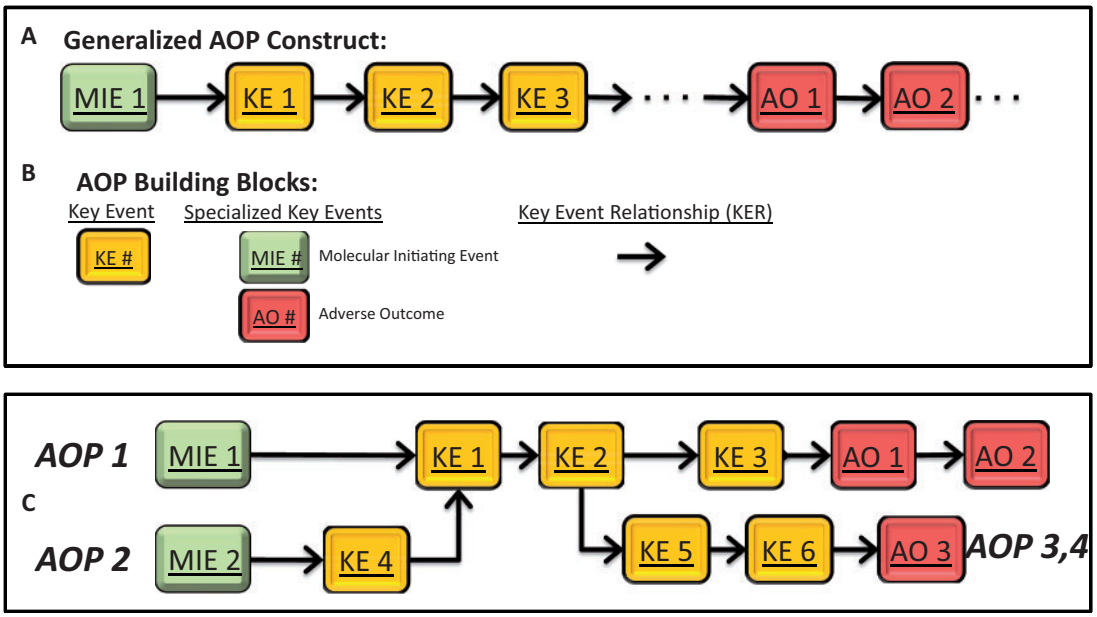

FIG. 1. Graphical representation of a generalized adverse outcome pathway (AOP; **A**). Each AOP is composed of two key components (**B**), key events (KEs) and key event relationships (KERs). Additionally, there are two specialized KEs, molecular initiating events (MIEs) and adverse outcomes (AOs) that anchor an AOP description. Individual AOPs sharing KEs or KERs can be represented as an AOP network (**C**). The AOP network depicted is composed of four individual AOPs, each representing a unique sequence of KEs linking an MIE to AO: AOP 1 [MIE1, KE1, KE2, KE3, AO1, AO2]; AOP 2 [MIE2, KE4, KE1, KE2, KE3, AO1, AO2]; AOP 3 [MIE1, KE1, KE2, KE5, KE6, AO3]; AOP 4 [MIE2, KE4, KE1, KE2, KE5, KE6, AO3]. Color image is available in the online version of the article.

TABLE 1. Primary Components of an Adverse Outcome Pathway (AOP)

Key event (KE)	• A measureable change in biological state that is essential, but not necessarily sufficient for the progression from a defined biological perturbation toward a specific AO. • Represented as nodes in an AOP diagram or AOP network. • Provide verifiability to an AOP description.
Key event relationship (KER)	• Define a directed relationship between a pair of KEs, identifying one as upstream and the other as downstream. • Supported by biological plausibility and empirical evidence. • Represented as a directed edge (i.e., an arrow) in an AOP diagram or AOP network. • Unit of inference or extrapolation within an AOP.
Molecular initiating event (MIE)	• A specialized type of KE. • Defined as the point where a chemical directly interacts with a biomolecule to create a perturbation—as such, by definition occurs at the molecular level. • Anchors the "upstream" end of an AOP.
Adverse outcome (AO)	• A specialized type of KE. • Measured at a level of organization that corresponds with an established protection goal and/or is functionally equivalent to an apical endpoint measured as part of an accepted guideline test. • Generally at the organ level or higher. • Anchors the "downstream" end of an AOP.

Additionally, early in the AOP development program it was recognized that generation of static documents was not an ideal way to either develop or disseminate AOP knowledge. Consequently, an AOP-Wiki (www.aopwiki.org, Accessed October 9, 2014) was developed as a user-friendly, open-source interface that facilitates both sharing of AOP knowledge and collaborative AOP development. The AOP-Wiki was designed with a series of structured and free-text fields that prompt users for the information requested by the guidance document and organizes it into a series of linked wiki pages (https://aopkb.org/aopwiki/index.php/Main_Page, Accessed October 9, 2014). Beyond text-based knowledge captured via the AOP-Wiki, a broader AOP knowledgebase (AOP-KB) is under development. The AOP-KB (www.aopkb.org, Accessed October 9, 2014) will include additional modules for: (1) visualization and analysis of AOP networks (AOP-XPlorer; http://www.aopxplorer.org/, Accessed October 9, 2014), (2) capturing intermediate effects data from toxicity studies (i.e., integration with OECD Harmonized Template 201; http://www.oecd.org/ehs/templates/, Accessed October 9, 2014), and (3) assembling data on quantitative relationships between KEs (i.e., through integration with Effectopedia; www.effectopedia.org, Accessed October 9, 2014). The ultimate goal is to assemble a comprehensive collection of accessible AOP knowledge using internationally accepted standards.

The guidance document and subsequent supplements to the guidance identify key information to include in an AOP description (OECD, 2013) and the AOP-KB (www.aopwiki.org) provides a structured, collaborative platform for assembling and disseminating AOP descriptions. Nonetheless, strategies and principles to guide the practice of AOP development have been lacking. Conventions and best practices have not been widely shared, leading to inconsistencies in how the guidance is applied and how AOPs are described. Therefore, broader dissemination of guiding principles and best practices of AOP development, as defined through the experiences of pioneering AOP practitioners involved in the OECD AOP development program or related research efforts, stands to benefit prospective AOP developers who may face some of the common questions, challenges, and uncertainties that other AOP developers have encountered.

Therefore, as part of a workshop focused on "*Advancing Adverse Outcome Pathways (AOP) for Integrated Toxicology and Regulatory Applications*" (March 2–7, 2014, Somma Lombardo, Italy; https://aopkb.org/saop/workshops/somma.html, Accessed October 9, 2014) the authors were charged with outlining a set of strategies, guiding principles, and best practices that would aid AOP developers in assembling AOP descriptions. Recognizing that OECD guidance on AOP development predated development of the AOP-Wiki, special attention was paid to practices and recommendations that would help developers make effective use of the AOP-Wiki, and ultimately the broader AOP-KB, as a structured, collaborative platform for development of AOPs. The objectives of the present manuscript were to introduce common AOP development strategies and detail core principles that guide the AOP development process. The overall goal was to expand upon the guidance document and provide guiding principles that address some of the common conceptual misunderstandings regarding the AOP framework and provide added consistency in the practice of AOP development.

AOP DEVELOPMENT STRATEGIES

One of the first questions that arises with regard to AOP development is "what is the best way to go about it?" Based on the experiences of a diversity of AOP developers, there is no single "one size fits all" step-wise AOP development process that is suitable for all AOP development scenarios. A number of different general strategies have been successfully employed (Table 2). For example, AOP development stemming from observations of an adverse phenotypic outcome at the individual level without clear knowledge of the underlying mechanisms would generally require a top-down strategy for AOP development (Table 2). In this case, effort would be directed toward making the causal linkages—via additional experiments, literature reviews, etc.—between the AO and the KEs that precede it (e.g., Ankley *et al.*, 2009; Kimber *et al.*, 2014). Alternatively, an MIE (e.g., a receptor-ligand interaction) may be well characterized with QSAR models and/or *in vitro* assays that allow chemicals to be rapidly screened for their affinity to the receptor (e.g., Farmahin *et al.*, 2012; Shiizaki *et al.*, 2014; Zhang *et al.*, 2013). However, the understanding of the toxicological significance of those receptor interactions needs to be established. A bottom-up strategy (e.g., Schmieder *et al.*, 2004) would aim to describe KEs that occur at the cellular, tissue, organ, and organismal

TABLE 2. Overview of Some Common AOP Development Strategies

Strategy[a]	Definition	Example
Top-down AOP development	Developer starts with an apical AO of interest and then delves down to progressively lower levels of biological organization in an effort to connect that outcome with a specific MIE (or multiple MIEs to construct a network of AOPs).	Investigators are interested in understanding the diversity of ways in which chemicals can adversely impact reproduction (Ankley *et al.*, 2009). An important human health issue is strongly associated with chemical exposure, but the causes are poorly understood and predictive assays are lacking (Kimber *et al.*, 2014).
Bottom-up AOP development	Developer starts with a well-defined MIE and begins linking to effects at higher levels of biological organization to develop the AOP.	A QSAR model or expert system for predicting chemical structures that will bind a receptor are available, but additional assays and endpoints are required to distinguish agonism from antagonism and link binding to hazard (Schmieder *et al.*, 2004).
Middle-out AOP development	Developer starts with an observable phenotype or biological measurement (i.e., a KE) that is neither thought to be directly perturbed by exogenous chemicals or stressors nor considered to have regulatory relevance in and of itself. The developer then starts developing connections to the mechanisms underlying change in that KE and to the significance of that event as part of a causal chain leading to an AO.	Investigators are interested in developing alternatives to a widely used chronic toxicity test. Major morphological events during development are used as a set of observable KEs from which to initiate AOP development. (Villeneuve *et al.*, 2013)
AOP development from case-study	Developer starts with a well defined sequence of biological events linking an MIE to AO for a single, well studied chemical, and then assembles evidence supporting generalization of that "motif of failure" for other chemicals/stressors that cause the same type of perturbation(s).	Detailed studies on the effects of 2,3,7,8-TCDD on cardiovascular development in fish was used as a basis for development of an AOP linking aryl hydrocarbon receptor agonism to early life stage mortality and/or impaired growth. (Volz *et al.*, 2011)
AOP development by analogy	Developer starts with an AOP that has been well defined in a particular animal model or particular class of organisms. Development focuses on evaluation of which KEs and KERs in the extant AOP are conserved in another organism/organism class of interest and develops alternative KEs and KERs for those that are not.	While not originally intended for AOP development, the basic concepts behind frameworks for analyzing the human relevance of animal model-based modes of action provide a prominent example. (Boobis *et al.*, 2008) Similarly, analogy between AOPs developed for fish and those that may occur in birds has been considered. (Lalone *et al.*, 2013)
AOP development from data-mining	Developers utilize high content and/or high-throughput data sets (e.g., "omics," Toxcast screening data) and other types of automated literature and data-base mining approaches to infer (generally statistically) relationships between KEs. This strategy is most often used for early stages of AOP development.	(Knudsen and Kleinstreuer, 2011) (Perkins *et al.*, 2011)

[a]The strategies described are not necessarily mutually exclusive.

level that can result from that initial receptor-ligand interaction, if the concentration and duration of exposure are sufficiently severe. A "middle-out" strategy would be most appropriate in situations where an intermediate KE is characterized, initially without an anchor to an MIE or an AO (e.g., Villeneuve *et al.*, 2013). For example, transcriptomics data can provide a profile of mRNA expression changes as a result of chemical exposure without an understanding of the MIE that caused the expression changes or the potential effects of those cellular modulations on tissue- or organismal-level changes. Overall, the appropriate strategy for a given AOP development process is dependent on a number of factors including the purpose for which the AOP is being developed (i.e., intended application) and the amount and type of supporting information available to the developer. In this respect, AOP development is similar to risk assessment in that a problem formulation step, in which the starting point and purpose for the AOP development is considered, can aid the identification of a general strategy that will be most appropriate to the situation.

KEY PRINCIPLES OF AOP DEVELOPMENT

Although there is no universal strategy for AOP development, there are a number of core principles that underlie AOP development, regardless of the specific strategy employed. Keeping these basic principles in mind will lead to more consistent AOP description and ultimately a more useful AOP-KB.

AOPs Are Not Chemical Specific

The ultimate goal of the AOP is to serve as a reliable predictive tool for chemical risk assessment in the context of environmental and human health. What separates a predictive science from a strictly empirical one is that the science has developed to the

point that generalized patterns of response, or, in the case of AOPs, motifs of system failure, are sufficiently understood that when one observes event A, one can reasonably predict the subsequent behavior of the system significantly better than by chance alone. In this context, the conceptual starting point of an AOP, i.e., the MIE, is defined as a specific type of interaction of a toxicant with a biological target (e.g., receptor, enzyme, and DNA) that represents the first step in a directed cascade of dependent biological processes (KEs) that lead to a defined AO. An implicit assumption underlying the AOP framework is that "*any*" chemical or stressor that triggers the MIE has the potential to elicit the chain of downstream KEs represented in the AOP, assuming that the magnitude and duration of perturbation at the MIE was sufficiently severe. Consequently, AOPs, by definition are not chemical specific.

This is not to suggest that chemical properties are not an important determinant of which AOPs a chemical will activate and with what potency. Indeed, identification of specific chemicals, chemical classes, and/or chemical properties known to confer a strong probability of interaction with a MIE is an important part of AOP description (see Villeneuve *et al.* 2014). Rather, the point of describing an AOP is the notion that the sequence of biological events that can be triggered by a specific type of chemical interaction at the MIE is similar regardless of the chemical that triggered it. There is no need to describe a separate AOP for every chemical that interacts with that MIE to produce a similar profile of downstream responses. Rather, if we directly measure a chemical as having biological activity at an MIE (e.g., via an *in vitro* screening assay), or predict its interaction at an MIE based on its chemical structure or properties, we posit that one can have some confidence in predicting the types of downstream responses it may elicit. This is particularly true if the chemical's potency at the MIE and the structure-dependent properties that influence its absorption, distribution, metabolism, and elimination (ADME), and thus probable dose at the target *in vivo* are understood. An example of the non-chemical-specific nature of AOP description was provided by Russom *et al.* (2014) who described an AOP linking acetylcholinesterase inhibition to acute toxicity in fish. This AOP is known to be relevant to dozens of organophosphate and carbamate insecticides, among other chemicals. A separate AOP is not described for each compound.

Confusion around the point of whether or not AOPs are chemical specific comes from two major sources. The first being that many AOPs are informed by data derived from experiments focused on just one or a few prototypic chemicals, particularly in the case of ligand-receptor driven pathways (Table 1). Along those lines, a number of publications on the AOP concept have included these prototypical chemical initiators as part of the AOP diagrams they present (Ankley *et al.*, 2010; Volz *et al.*, 2011). This has fostered the misperception that specific chemicals and their specific properties are part of the AOP. As a practical counter to this thinking, it is important to note that one of the primary applications of AOPs, in an OECD and regulatory context, is in chemical category formation (OECD, 2011, 2013). If AOPs were chemical specific, they would have no value for this application.

The second major source of confusion comes from the conflation of AOP development/description with an AOP application. AOPs describe biology and relationships between biological events. They describe motifs of failure, assuming that the magnitude and duration of perturbation at the MIE (the point of interaction between the biological system and the chemical) are sufficiently severe to drive the responses to the AO. However, a critical aspect of actually *applying* these predictive motifs in chemical-specific assessments involves understanding of chemical-specific properties including potency and pharmacokinetic factors (i.e., ADME) that ultimately define the magnitude and duration of perturbation at the MIE (even if measurements are actually made at one of the downstream KEs). Because chemical-specific properties need to be considered when using an AOP to predict outcomes of a particular exposure, there is a tendency to want to define AOPs in a chemical-specific manner. This actually limits the utility of the AOPs. When developed in a chemical-specific manner, AOPs become little more than a summation of empirical evidence. In contrast, AOPs developed in a manner linking a specific molecular perturbation to an AO, regardless of the specific initiator, have predictive utility. Consequently, as a guiding principle of AOP development, developers should keep in mind that AOPs are not chemical specific.

AOPs Are Modular

For the AOP framework to be useful in the context of chemical risk assessment and fundamental research there needs to be a streamlined, practical, and functional approach to AOP development. On the one hand, AOPs need to be clear, transparent, and easy to understand and apply. On the other hand, the AOP framework must provide a large degree of flexibility and accommodate varying levels of detail as appropriate to the available supporting information. This is accomplished through the use of a modular structure (Fig. 1). Each AOP can be broken down into two fundamental units (Table 1), KEs and KERs. In the traditional AOP diagram, KEs are represented by boxes whereas KERs are represented as the arrows connecting a pair of boxes (Fig. 1). In a graph theory context, KEs represent nodes and KERs represent edges (Pavlopoulos *et al.*, 2011). In the AOP-Wiki (www.aopwiki.org, Accessed October 9, 2014) and Effectopedia (www.effectopedia.org, Accessed October 9, 2014) each of these building blocks are represented as separate pages with different types of information content which can be linked together to form an AOP description. KEs are, in essence, measurements of biological state or change in state with regard to a control or reference. Because KEs are measurements or observations of state, the confidence one has in a KE is dictated by the accuracy and precision with which that biological state can be measured. KERs, in contrast, are a unit of inference or extrapolation. They are defined by the biological plausibility and evidence that provide a scientifically credible basis for inferring or predicting the state of a downstream KE based on the known state of an upstream KE and the confidence in that inference or prediction is defined by the weight of supporting evidence. Thus, KERs give AOPs their predictive utility, whereas KEs provide verifiability.

A critical consideration with regard to AOP development is that "KEs and KERs are not unique to a single AOP." For example, a case study presented by Ankley *et al.* (2010) depicted three separate AOPs that shared a common KE, reduced vitellogenin production, leading to another common KE, impaired oocyte development, linked together by a common KER. Similarly, Crofton (2008) showed how multiple independent MIEs could converge at the common KE of reduced serum thyroid hormone concentrations, with the down-stream events leading to a life stage-dependent AO, shared in common. From a practical standpoint, in terms of AOP development, this means that these fundamental building blocks are reusable. If done appropriately, once the information needed for a KE or KER description is assembled, it does not need to be regenerated independently for every new AOP. Effectively, one can simply reuse that

information in the new AOP description. In the context of the AOP-KB, this is accomplished by simply linking to a KE or KER that has already been described in the knowledgebase. This concept of reusability has to be considered throughout the AOP development process, in order to take full advantage of this modular aspect of AOP development.

An Individual AOP Is a Pragmatic Unit of Development and Evaluation

The idea that KEs and KERs can be shared by multiple AOPs and that perturbation of a variety of MIEs can converge at certain KEs and subsequently share downstream KEs and KERs is relatively intuitive to most toxicologists. Such convergence is both a fundamental principle of systems biology (Csete and Doyle, 2004) and a function of the hierarchy of biological levels of organization (i.e., individuals are composed of multiple organ systems, organ systems of multiple organs, organs of multiple tissues, tissues of multiple cells, cells of multiple subcellular compartments, etc.). The question frequently raised then, is: why define individual AOPs as a single, non-branching sequence of KEs, linked by KERs, connecting a single MIE to an AO (e.g., Fig. 1)?

It is acknowledged that most biological processes, including mechanisms of failure resulting from exposure to toxic chemicals, operate in a systems context where interaction and crosstalk with other pathways is the norm, not the exception. It is for purely pragmatic reasons that individual AOPs are defined as a single chain of KEs. Specifically, inclusion of all possible branches of KE strings under one AOP would render it difficult or even impossible to "finish" an AOP description. Allowing for some degree of branching in the chain of KEs would lead to inconsistencies in AOP description that would both complicate evaluation of the predictive relationships it defines and make it intractable to define a discrete unit of AOP development. Fundamentally, AOPs are not intended to be a complete representation of complex biological processes but rather provide a structured and simplified way of organizing toxicological knowledge in a manner that enhances its utility for decision support and chemical risk assessment. Further, AOPs are not intended to represent every possible way one could get from a particular perturbation to an AO, but rather define one of the ways that can occur. By analogy, each AOP is a single set of directions describing how to get from point A to point B, not a comprehensive map of every possible route one could take. In this context, an individual AOP is the simplest functional unit of prediction that is represented through a finite assembly of KEs that describes a set of predictive relationships connecting a specific type of MIE (i.e., a type of chemical-biological interaction) to an AO. For a theoretical "pure ligand," i.e., a compound that is specific for a target and one target only, an individual AOP is also a functional unit of prediction which defines the maximum amount of predictive certainty that can be achieved based on the relationships and supporting evidence underlying that AOP. As a result, individual AOPs represent a practical unit upon which to conduct a weight of evidence evaluation (Meek *et al.*, 2014; Becker *et al.*—https://aopkb.org/saop/workshops/somma.html, Accessed October 9, 2014) which in turn defines the suitability of a set of predictive relationships, encompassed by the AOP, for supporting different types of regulatory decision making.

For Most Real-World Applications, AOP Networks Are the Functional Unit of Prediction

Although individual AOPs can be viewed as discrete, pragmatic units for AOP development and evaluation, hypothetical "pure ligands" which interact with a single MIE are likely to be uncommon. Similarly, most real-world exposure scenarios involve exposure to complex mixtures, not individual chemicals, let alone "pure ligands." Finally, even if perturbation of a single MIE can be assumed, that perturbation can potentially lead to different motifs of failure, and thus be linked to different AOPs, depending on the magnitude and duration of perturbation and the biological context (i.e., cell/tissue/organ, sex, life stage, taxa) in which it takes place. Consequently, in practice, prediction of AOs based on mechanistic or pathway-based data will often require consideration of multiple AOPs, many of which may share KEs and KERs. Systems of multiple interacting AOPs sharing one or more common KEs or KERs form AOP networks (Fig. 1C).

AOP networks are envisioned to be a more realistic representation of the complex biological interactions that would occur in response to chemical mixtures or single toxicants exhibiting multiple biological activities (i.e., perturbing multiple MIEs). Consideration and analysis of AOP networks have potential to provide important information regarding the interactions among multiple AOPs, and represent an interface between the specific toxic outcome captured in a single AOP and modulation of those outcomes due to interactions occurring in a systems biology context. Additionally, analysis of the intersections (shared KEs and KERs) among AOPs that make up an AOP network can reveal unexpected or under-appreciated biological connections.

De novo construction of a comprehensive AOP network would be a daunting endeavor akin to mapping all known biology. However, by assembling individual AOPs using modular building blocks of KE and KER descriptions that can be shared by and linked to multiple AOPs represented in a common AOP-KB, the AOP development community contributes to *de facto* AOP network development. Thus, although individual AOPs may not capture the biological complexity needed to make effective predictions of toxicological outcome for many real-world scenarios, the collective endeavor of AOP development can provide those tools. However, for that to be successful, AOP developers have to keep the principles of modularity, individual AOPs as a pragmatic unit of development and evaluation, and AOP networks as the functional unit of prediction and application in mind as they develop and update AOP descriptions.

AOPs Are Living Documents

The concept of updating AOP descriptions brings us to our last guiding principle of AOP development; AOPs are "living documents." It is important to understand that AOPs are not static. KEs are observable/measurable changes in biological state. As such, the tools, techniques, and assays that we use to measure those biological states and the level of accuracy and precision with which they can be measured can be expected to change over time. At their core, KERs represent weight of evidence supporting the supposition that a measurable change in the upstream KE can lead to a measurable change in the downstream KE that the KER connects. Evidence from the literature either supporting or rejecting a particular KER can also be expected to evolve over time. Thus, AOP descriptions built from linked KE and KER descriptions must also be able to evolve over time. This principle has several important implications for AOP development.

First, there is no objective definition of a "complete" AOP. Rather our understanding of the relative state of an AOP's development must be regarded in the context of a living document that has the potential to develop and evolve over time as

TABLE 3. Three Operationally Defined Stages or Phases of AOP Development

Operationally Defined Stage/Phase	Characteristics
Putative AOP development	Assembly of a hypothesized set of KEs and KERs supported primarily through biological plausibility and/or statistical inference. Assembly of partial AOPs with incomplete linkage between the MIE and AO as a result of known gaps and uncertainties.
Qualitative formal AOP development	Assembly of KEs supported by descriptions of how the KEs can be measured and KERs supported by empirical evidence in addition to plausibility or statistical inference, along with qualitative evaluation of the overall weight of evidence supporting the AOP. Characterized as formal in that the information included in the descriptions is assembled in a manner consistent with the internationally harmonized OECD guidance.
Quantitative AOP development	Assembly of KEs supported by descriptions of how the KEs can be measured and the accuracy and precision with which the measurements are made along with KERs supported by quantitative understanding of what magnitude and/or duration of change in the upstream KE is needed to evoke some magnitude of change in the downstream KE.

additional knowledge becomes available. In practice it has been useful to refer to three phases or stages of AOP development (Table 3); however, it is important to note that categorization in a particular phase is neither entirely objective, nor absolute. Further, although they can be seen to reflect different stages of AOP development, moving from a less mature to more mature AOP description, they may all have their uses, in terms of regulatory decision support. For example, a partial or hypothesis-based AOP, where not all the KEs are known, can be useful in priority setting for further testing and development to fill data gaps. A qualitative (formal) AOP has KEs supported by description of how those KEs are measured and their taxonomic domains of applicability and KERs supported by empirical evidence in addition to biological plausibility or statistical association, but lacks sufficient detail to allow the determination of quantitative relationships between all KEs. Qualitative AOPs may support, for instance, low tiers of risk assessment that apply conservative assumptions and/or do not require a quantitative characterization of uncertainty. Finally, a quantitative AOP has sufficient data and appropriate testing methodologies to establish quantitative linkages between all KEs—from the MIE to the AO. From a practical standpoint, the fact that each stage of AOP development has potential utility and that AOPs can evolve over time toward greater predictive sophistication (or toward obsolescence if rejected by subsequent evidence) means that all levels of AOP development represent a useful contribution. It is not necessary to have a fully formed, quantitative AOP in order to make a significant contribution to the AOP-KB.

Likewise, individual AOPs need not capture all contexts and scenarios. For example, the point has often been raised that a short-duration, high intensity perturbation of the MIE may lead to a different sequence of downstream KEs than a chronic, low intensity perturbation. Both scenarios can be captured via the modular aspects of the AOP network. Although the MIE description would remain the same, separate KERs linking that MIE to the subsequent divergent downstream KEs would describe the differing nature of the perturbation and be reflected as a branching in the AOP network. Outcomes of a common perturbation that diverge in different taxa due to differences in physiology at higher levels of biological organization can be handled in the same manner. Thus, an AOP developer can focus his/her attention on specific scenarios and individual AOPs aligned with his or her expertise and interests and rely on other developers to add branches to the AOP network as they contribute to the AOP-KB.

The modularity of KEs and KERs in the AOP-KB and AOP networks has an important practical significance relative to the principle of AOPs as living documents. Notably, it creates efficiency in updating and constructing AOPs. For KE or KER descriptions created in the AOP-KB (using the AOP-Wiki or Effectopedia) any updates made to those KE or KER descriptions (e.g., adding additional supporting evidence as new studies are published) are automatically updated for all AOPs that included that KE or KER. There is no need to go in and manually update a dozen different AOP descriptions that may contain that KE or KER. Likewise, as an AOP developer adds a new AOP description s/he can simply contribute to or update existing KE or KER descriptions rather that generating new descriptions from scratch. This can be both an incredible time-saver and can lead to more robust AOPs overall, as different developers with different backgrounds and expertise and corresponding familiarity and access to different supporting data, incrementally contribute to and evaluate shared KE and KER descriptions. Thus, the concepts of AOP modularity, as well as AOPs as living documents are closely intertwined and represent key principles that inform and guide the process of AOP development.

CONCLUSIONS

As a conceptual framework, AOPs have promise as a tool to help support so called twenty-first century approaches to regulatory toxicology (Krewski *et al.*, 2010). However, in order to realize that promise, it is important for the toxicology research community to develop AOP descriptions in a consistent, scientifically rigorous, and transparent manner. The principles and strategies described here support those objectives. Definition of relevant toxicological response motifs in the form of a network of AOPs that cover the diversity of biological and exposure contexts in which they occur is a daunting scientific challenge. Arguably, the only tractable approach is to collectively assemble such an AOP network from modular components composed of individual AOPs as a practical unit of development and evaluation, KERs as a functional unit of inference/extrapolation that lend AOPs their predictive utility, and KEs which can be used to track or verify progression of toxicity along a particular AOP or within an AOP network. Keeping this core set of guiding principles in mind during AOP development can help harness the collective expertise of the toxicology community in a coordinated manner that yields a useful AOP-KB. A set of best practices based on these principles are provided in a companion paper (Villeneuve *et al.* 2014).

Practical application of the non-chemical-specific toxicological response motifs represented by AOPs for chemical-specific predictive risk assessment still requires suitable measurements or estimates of potency, along with complementary toxicokinetic modeling tools particularly if biological effects data are to be extrapolated from simplified *in vitro* systems to predicted *in vivo* outcomes. Additionally, AOP development is just one step in a broader process. Other products of the international expert workshop on "*Advancing Adverse Outcome Pathways (AOP) for Integrated Toxicology and Regulatory Applications*" (March 2–7, 2014, Somma Lombardo, Italy) address other critical aspects concerning regulatory application of the AOP framework including: weight of evidence evaluation (Becker *et al.*—https://aopkb.org/saop/workshops/somma.html), regulatory acceptance (Perkins *et al.*—https://aopkb.org/saop/workshops/somma.html), use of AOPs to support integrated approaches to testing and assessment (Tollefsen *et al.*, 2014), and near term AOP development priorities (Groh *et al.*—https://aopkb.org/saop/workshops/somma.html). Together, these efforts are aimed at realizing the potential of the AOP framework for supporting more predictive approaches to regulatory toxicology.

FUNDING

The American Chemistry Council, BioDetection Systems, European Centre for Ecotoxicology and Toxicology of Chemicals, Environment Canada, European Commission Directorate General Joint Research Center, Human Toxicology Project Consortium, International Life Sciences Institute—Health and Environmental Science Institute, The Research Council of Norway (Grant no. 221455), Society for Environmental Toxicology and Chemistry, US Army Engineer Research and Development Center, and the US Environmental Protection Agency contributed support for the workshop.

ACKNOWLEDGMENTS

The content of this paper is the result of an international expert workshop on *Advancing Adverse Outcome Pathways (AOPs) for Integrated Toxicology and Regulatory Applications*. The authors acknowledge the workshop organizing committee, and all workshop participants for both plenary and informal discussions that influenced the content presented here. Special thanks to Dr Michael Hornung for reviewing an earlier draft of this article. The views expressed are those of the authors, and do not necessarily represent the views of the organizations the authors are affiliated with or the sponsors. Mention of trade names or commercial products does not constitute endorsement or recommendation for use.

REFERENCES

Ankley, G. T., Bencic, D., Breen, M., Collette, T. W., Connolly, R., Denslow, N. D., Edwards, S., Ekman, D. R., Jensen, K. M., Lazorchak, J., *et al.* (2009). Endocrine disrupting chemicals in fish: developing exposure indicators and predictive models of effects based on mechanisms of action. *Aquat. Toxicol.* **92**, 168–178.

Ankley, G. T., Bennett, R. S., Erickson, R. J., Hoff, D. J., Hornung, M. W., Johnson, R. D., Mount, D. R., Nichols, J. W., Russom, C. L., Schmieder, P. K., *et al.* (2010). Adverse Outcome Pathways: a conceptual framework to support ecotoxicology research and risk assessment. *Environ. Toxicol. Chem.* **29**, 730–741.

Boobis, A. R., Cohen, S. M., Dellarco, V., McGregor, D., Meek, M. E., Vickers, C., Willcocks, D., and Farland, W. (2006). IPCS framework for analyzing the relevance of a cancer mode of action for humans. *Crit. Rev. Toxicol.* **36**, 781–792.

Boobis, A. R., Doe, J. E., Heinrich-Hirsch, B., Meek, M. E., Munn, S., Ruchirawat, M., Schlatter, J., Seed, J., and Vickers, C. (2008). IPCS framework for analyzing the relevance of a noncancer mode of action for humans. *Crit. Rev. Toxicol.* **38**, 87–96.

Bradbury, S. P., Feijtel, T. C., and Van Leeuwen, C. J. (2004). Meeting the scientific needs of ecological risk assessment in a regulatory context. *Environ. Sci. Technol.* **38**, 463A–470A.

Crofton, K. M. (2008). Thyroid disrupting chemicals: mechanisms and mixtures. *Int. J. Androl.* **31**, 209–223.

Csete, M., and Doyle, J. (2004). Bow ties, metabolism and disease. *Trends Biotechnol.* **22**, 446–450.

Farmahin, R., Wu, D., Crump, D., Herve, J. C., Jones, S. P., Hahn, M. E., Karchner, S. I., Giesy, J. P., Bursian, S. J., Zwiernik, M. J., *et al.* (2012). Sequence and in vitro function of chicken, ring-necked pheasant, and Japanese quail AHR1 predict in vivo sensitivity to dioxins. *Environ. Sci. Technol.* **46**, 2967–2975.

Kimber, I., Dearman, R. J., Basketter, D. A., and Boverhof, D. R. (2014). Chemical respiratory allergy: reverse engineering an adverse outcome pathway. *Toxicology* **318C**, 32–39.

Knudsen, T. B., and Kleinstreuer, N. C. (2011). Disruption of embryonic vascular development in predictive toxicology. *Birth Defects Res. Part C Embryo Today Rev.* **93**, 312–323.

Krewski, D., Acosta, D., Jr, Andersen, M., Anderson, H., Bailar, J. C., 3rd, Boekelheide, K., Brent, R., Charnley, G., Cheung, V. G., Green, S., Jr, *et al.* (2010). Toxicity testing in the 21st century: a vision and a strategy. *J. Toxicol. Environ. Health B Crit. Rev.* **13**(2–4), 51–138.

Lalone, C. A., Villeneuve, D. L., Garber, K. V., Russom, C., Etterson, M. A., Bennett, R. S., and Ankley, G. T. (2013). Development and evaluation of adverse outcome pathways predicting adverse effects of conazole fungicides on avian species. SETAC North America, 34th Annual Meeting, Nashville, TN, USA, November 17–21, 2013.

Meek, M. E., Boobis, A., Cote, I., Dellarco, V., Fotakis, G., Munn, S., Seed, J., and Vickers, C. (2014). New developments in the evolution and application of the WHO/IPCS framework on mode of action/species concordance analysis. *J. Appl. Toxicol.* **34**, 1–18.

OECD. (2011). Report of the workshop on using mechanistic information in forming chemical categories. Demcember 8–10, 2010, Crystal City, VA, USA. In *Series on Testing and Assessment No. 138*, Vol. ENV/JM/MONO(2011)8, p. 176. Organisation for Economic Cooperation and Devleopment, Environment Directorate Paris, France.

OECD. (2013). Guidance document on developing and assessing adverse outcome pathways. In *Series on Testing and Assessment, No. 184*, Vol. ENV/JM/MONO(2013)6, p. 45. Organisation for Economic Cooperation and Devleopment, Environment Directorate Paris, France.

Pavlopoulos, G. A., Secrier, M., Moschopoulos, C. N., Soldatos, T. G., Kossida, S., Aerts, J., Schneider, R., and Bagos, P. G. (2011). Using graph theory to analyze biological networks. *BioData Min.* **4**, 10.

Perkins, E. J., Chipman, J. K., Edwards, S., Habib, T., Falciani, F., Taylor, R., Van Aggelen, G., Vulpe, C., Antczak, P., and

Loguinov, A. (2011). Reverse engineering adverse outcome pathways. *Environ. Toxicol. Chem.* **30**, 22–38.

Russom, C. L., LaLone, C. A., Villeneuve, D. L., and Ankley, G. T. (2014). Development of an adverse outcome pathway for acetylcholinesterase inhibition leading to acute mortality. *Environ. Toxicol. Chem.* **33**, 2157–2169.

Schmieder, P. K., Tapper, M. A., Denny, J. S., Kolanczyk, R. C., Sheedy, B. R., Henry, T. R., and Veith, G. D. (2004). Use of trout liver slices to enhance mechanistic interpretation of estrogen receptor binding for cost-effective prioritization of chemicals within large inventories. *Environ. Sci. Technol.* **38**, 6333–6342.

Shiizaki, K., Yoshikawa, T., Takada, E., Hirose, S., Ito-Harashima, S., Kawanishi, M., and Yagi, T. (2014). Development of yeast reporter assay for screening specific ligands of retinoic acid and retinoid X receptor subtypes. *J. Pharmacol. Toxicol. Methods* **69**, 245–252.

Tollefsen, K. E., Scholz, S., Cronin, M. T., Edwards, S. W., de Knecht, J., Crofton, K., Garcia-Reyero, N., Hartung, T., Worth, A., and Patlewicz, G.. (2014). Applying adverse outcome pathways (AOPs) to support integrated approaches to testing and assessment (IATA). *Reg. Toxicol. Pharmacol.* (in press).

Villeneuve, D. L., Crump, D., Garcia-Reyero, N., Hecker, M., Hutchinson, T. H., LaLone, C. A., Landesmann, B., Lettieri, T., Munn, S., Nepelska, M., *et al.* (2014). Adverse outcome pathway (AOP) development II: best practices. *Toxicol. Sci.* (this issue).

Villeneuve, D., Volz, D. C., Embry, M. R., Ankley, G. T., Belanger, S. E., Leonard, M., Schirmer, K., Tanguay, R., Truong, L., and Wehmas, L. (2013). Investigating alternatives to the fish early life-stage test: a strategy for discovering and annotating adverse outcome pathways for early fish development. *Environ. Toxicol. Chem.* **33**, 158–169.

Volz, D. C., Belanger, S., Embry, M., Padilla, S., Sanderson, H., Schirmer, K., Scholz, S., and Villeneuve, D. (2011). Adverse outcome pathways during early fish development: a conceptual framework for identification of chemical screening and prioritization strategies. *Toxicol. Sci.* **123**, 349–358.

Zhang, L., Sedykh, A., Tripathi, A., Zhu, H., Afantitis, A., Mouchlis, V. D., Melagraki, G., Rusyn, I., and Tropsha, A. (2013). Identification of putative estrogen receptor-mediated endocrine disrupting chemicals using QSAR- and structure-based virtual screening approaches. *Toxicol. Appl. Pharmacol.* **272**, 67–76.