Environmental Toxicology and Chemistry—Volume 37, Number 6—pp. 1723–1733, 2018
Received: 11 October 2017 | Revised: 11 December 2017 | Accepted: 24 February 2018

Environmental Toxicology

Adverse Outcome Pathway Networks I: Development and Applications

Dries Knapen,[a,*] Michelle M. Angrish,[b] Marie C. Fortin,[c] Ioanna Katsiadaki,[d] Marc Leonard,[e] Luigi Margiotta-Casaluci,[f] Sharon Munn,[g] Jason M. O'Brien,[h] Nathan Pollesch,[i] L. Cody Smith,[j] Xiaowei Zhang,[k] and Daniel L. Villeneuve[i]

[a]Zebrafishlab, Veterinary Physiology and Biochemistry, University of Antwerp, Wilrijk, Belgium
[b]National Center for Environmental Assessment, US Environmental Protection Agency, Research Triangle Park, North Carolina, USA
[c]Department of Pharmacology and Toxicology, Ernest Mario School of Pharmacy, Rutgers University, Piscataway, New Jersey, USA
[d]Centre for Environment, Fisheries and Aquaculture Science, Weymouth, United Kingdom
[e]L'Oréal Advanced Research, Aulnay-sous-Bois, France
[f]Institute of Environment, Health and Societies, Brunel University London, London, United Kingdom
[g]Joint Research Centre, European Commission, Ispra, Italy
[h]National Wildlife Research Centre, Environment and Climate Change Canada, Ottawa, Ontario, Canada
[i]Mid-Continent Ecology Division, US Environmental Protection Agency, Duluth, Minnesota, USA
[j]Center for Environmental and Human Toxicology, University of Florida, Gainesville, Florida, USA
[k]State Key Laboratory of Pollution Control and Resource Reuse, School of the Environment, Nanjing University, Nanjing, People's Republic of China

Abstract: Based on the results of a Horizon Scanning exercise sponsored by the Society of Environmental Toxicology and Chemistry that focused on advancing the adverse outcome pathway (AOP) framework, the development of guidance related to AOP network development was identified as a critical need. This not only included questions focusing directly on AOP networks, but also on related topics such as mixture toxicity assessment and the implementation of feedback loops within the AOP framework. A set of two articles has been developed to begin exploring these concepts. In the present article (part I), we consider the derivation of AOP networks in the context of how it differs from the development of individual AOPs. We then propose the use of filters and layers to tailor AOP networks to suit the needs of a given research question or application. We briefly introduce a number of analytical approaches that may be used to characterize the structure of AOP networks. These analytical concepts are further described in a dedicated, complementary article (part II). Finally, we present a number of case studies that illustrate concepts underlying the development, analysis, and application of AOP networks. The concepts described in the present article and in its companion article (which focuses on AOP network analytics) are intended to serve as a starting point for further development of the AOP network concept, and also to catalyze AOP network development and application by the different stakeholder communities. *Environ Toxicol Chem* 2018;37:1723–1733. © 2018 The Authors. Environmental Toxicology and Chemistry published by Wiley Periodicals, Inc. on behalf of SETAC.

Keywords: Adverse outcome pathway; Risk assessment; Predictive toxicology; Adverse outcome pathway network; Network development; Network topology

INTRODUCTION

Adverse outcome pathways (AOPs) constitute an important framework that can help support greater and more effective use of mechanistic, or pathway-based, data in risk assessment and regulatory decision-making. While the conceptual underpinnings of AOP frameworks date back to at least the late 1980s (LaLone et al. 2017a); AOPs have rapidly evolved from a conceptual paradigm (Ankley et al. 2010) into a formalized framework for organizing biological and toxicological knowledge according to a set of principles and guidelines that are generally accepted by the scientific and regulatory communities (Organisation for Economic Co-operation and Development 2013a, 2015; Villeneuve et al. 2014b), and for disseminating that knowledge through an internationally harmonized knowledge-base (Society for the Advancement of Adverse Outcome Pathways 2017; Organisation for Economic Co-operation and

This article includes online-only Supplemental Data.

* Address correspondence to dries.knapen@uantwerpen.be
Published online 28 February 2018 in Wiley Online Library (wileyonlinelibrary.com).
DOI: 10.1002/etc.4125

Development 2017). Nonetheless, further development of the framework and the tools, approaches, and concepts surrounding its application is required to fully realize its potential and acceptance by society.

In response to the recognized need to continue advancing the framework, the Society of Environmental Toxicology and Chemistry (SETAC) sponsored a global Horizon Scanning exercise to identify major outstanding topics and challenges related to the AOP framework and its application (LaLone et al. 2017a). Based on a survey of the international stakeholder community, 4 major topics/themes that needed further development were identified: 1) enhancement of communication, outreach, and stakeholder engagement in the development and application of AOP knowledge; 2) enhancement of regulatory use and acceptance of the AOP framework and facilitation of its incorporation into regulatory practices; 3) enhanced use of the framework for quantitative assessments and applications; and 4) development of approaches for deriving, interrogating, and applying networks of AOPs, which is the topic of the present article.

As outlined by Villeneuve et al. (2014b), individual AOPs are viewed as a pragmatic unit of development and evaluation. It is tractable for an individual or a research team to describe and establish, through both biological plausibility and supporting evidence, how a defined perturbation of a biological system can lead, in a causal manner, to a particular adverse outcome. It is far less tractable for that individual or team to describe all the possible adverse effects a given perturbation may cause, or, conversely, all the different perturbations through which stressors may evoke a particular adverse outcome (e.g., reductions in survival, growth, and reproduction; increased risk of disease). It is even more daunting to consider describing those possibilities for all the different taxa, life stages, and sexes (where relevant) that are of interest to a stakeholder. However, at the same time it was recognized that the "one perturbation–one adverse outcome" model that an individual AOP represents is a gross oversimplification of both the complexity of biological systems and the consequences of exposures to the stressors that they face. In most real-world scenarios, exposures are to multiple stressors (i.e., mixtures), not just one stressor at a time. Likewise, even single stressors may induce toxicity by more than one mechanism, via interaction of the chemical with multiple targets in an organism or via interaction with a single target found in multiple compartments (e.g., cell types, tissues, organs, etc.) within a complex organism. Thus, most often, AOPs cannot be considered in isolation. One needs to think about potential interactions among pathways and consider how those interactions may alter the trajectory or intensity of the effects resulting from a chemical exposure.

Recognizing this, one of the core principles of AOP development was that, in contrast to individual AOPs as pragmatic units of development, AOP networks are viewed as the most likely units of prediction (Villeneuve et al. 2014b). In turn, the formalization of the AOP framework, and its implementation via a knowledgebase structure that allowed for sharing of an AOP's modular units (key events and key event relationships, as found in AOP-Wiki; Society for the Advancement of Adverse Outcome Pathways 2017), was conceived and designed to allow for de facto construction of more complex and comprehensive networks from individual AOPs. In this way, a more accurate representation of biological and toxicological complexity that covers more and more of the susceptible taxonomic space and biological contexts (e.g., life stage, sex, impacts in or on different target organs) can be built up gradually through the independent contributions of individuals or groups.

A vision for AOP networks has just begun to be realized. Following publication of principles and best practices for AOP development (Villeneuve et al. 2014b, 2014c) and public release of the AOP-Wiki (Society for the Advancement of Adverse Outcome Pathways 2017) in 2014, time was needed to allow for an accumulation of a sufficient number of AOPs in the AOP knowledgebase to actually begin exploring their connectivity. Likewise, technical and practical challenges in the development of sharable, modular key event and key event relationship units in the public AOP knowledgebase (e.g., the development of naming conventions, search tools, guidance and training materials, etc.) initially hampered rapid assembly of these de facto networks. Nonetheless, over the last 3 yr, a critical mass of AOP descriptions has started to accumulate, and some of the challenges have been overcome. This has led to the recent realization of some of the first examples of AOP networks (Knapen et al. 2015; Angrish et al. 2016, 2017; Margiotta-Casaluci et al. 2016; LaLone et al. 2017b), as well as opportunities to address key concepts related to the development, analysis, and application of AOP networks.

The present set of 2 articles begins to explore these concepts. In the present article (part I), derivation of AOP networks is considered in the context of how it differs from development of individual AOP descriptions. We then discuss the application of filters and layers to refine and enrich derived AOP networks so that they may be tailored to address specific questions of interest. Modifications to the AOP knowledgebase that may be needed accordingly are also considered. We then briefly introduce a number of analytical and computational approaches that may be used to characterize and analyze the structure of AOP networks to derive information that can guide research and regulatory decision-making. These analytical concepts are further developed and described in part II by Villeneuve et al. (2018), including the use of techniques derived from graph theory (Trudeau 2013) and network science (Lewis 2009), to analyze network topology, the identification of critical paths, and the characterization of interactions among AOPs in a network. Finally, we present a number of application case studies that illustrate concepts underlying the development and analysis of AOP networks, and how those concepts tie in with ultimate application. Although the article is not comprehensive in scope, the intent is to provide an enhanced understanding of AOP network development, AOP network analysis (Villeneuve et al. 2018, part II), and applications of AOPs; as well as to provide perspectives on how some of the challenges identified through the Horizon Scanning exercise (LaLone et al. 2017a) can be addressed.

DEVELOPMENT OF AOP NETWORKS

A first and relevant question is: What exactly is an AOP network? An AOP network is defined as an assembly of 2 or more AOPs that share one or more key events, including specialized key events such as molecular initiating events and adverse outcomes (Table 1). Different AOPs diverging from a single molecular initiating event, or converging to a single adverse outcome, therefore also form AOP networks even if they do not have any other key event in common. Development of individual AOPs can be thought of as the process of 1) graphically defining a sequence of key events that link a molecular initiating event to a defined adverse outcome, 2) describing the change in state that each key event represents and how it is measured, and 3) detailing the weight of evidence that supports inference or extrapolation from one key event to the next in the sequence based on biological plausibility, empirical support, and quantitative understanding (Villeneuve et al. 2014b). The AOP networks can be thought of as emerging from the description of individual AOPs, as soon as key events are described that are shared between 2 or more AOPs. Either the description of networked key events can be an intentional process that is part of the strategy of an AOP developer, or the fact that certain key events are shared among AOPs can be discovered after AOPs have been developed independently. When one is considering different AOP network development processes, it is therefore useful to distinguish between network-guided AOP development and AOP network derivation. Whereas AOP network derivation is defined as a formal AOP network development process based on extracting and linking information that is available in the AOP-Wiki, network-guided AOP development is introduced as a rather broadly defined concept that includes many different AOP network development approaches that do not necessarily rely on database extraction procedures.

Network-guided AOP development

When AOPs are developed in the AOP-Wiki, an AOP network is created by default whenever a key event or key event relationship description is linked to more than one AOP. This is important because it implies that AOP developers are not restricted to describing only linear paths, and can thus intentionally conceive and describe structures that are more complex than the typical one perturbation–one outcome unit. This process could be thought of as network-guided AOP development. The advantage of network-guided AOP development is that it is not conceptually and methodologically different from the development and description of individual AOPs: the same principles, guidance, and practices in terms of description within the AOP-Wiki apply, and no additional tools are required. To develop an AOP network, there is no need to do anything differently than one would for describing a linear AOP, other than to intentionally share key event or key event relationship descriptions (pages) among more than one AOP, a functionality that is currently built into the AOP-Wiki.

Currently, many AOPs are being developed in this network-guided fashion (see Angrish et al. 2016; Nelson et al. 2016; Stinckens et al. 2016; Cavallin et al. 2017; LaLone et al. 2017b). However, it is expected that as the AOP knowledgebase matures, AOP development will increasingly focus on filling data and knowledge gaps in the AOP-Wiki. The AOP network

TABLE 1: Coming to terms with AOP networks

Term	Characteristics
AOP network	An assembly of 2 or more AOPs that share 1 or more key events.
AOP network development	Broad term referring to the description or development of AOP networks, irrespective of the strategy employed.
Network-guided AOP development	AOP network development strategy involving the development of at least 2 individual AOPs containing 1 or more intentionally shared key events.
AOP network derivation	AOP network development by manually or programmatically extracting AOPs relevant for a given application from the AOP-Wiki.
AOP network analytics	Broad term referring to the analysis of AOP networks to reveal, identify, or investigate specific network properties, such as topological features, critical paths, or interactions between AOPs.
AOP network filter	AOP network development tool to refine which key events and key event relationships from a given AOP network are included in downstream applications and analysis based on specified filter criteria.
AOP network layer	Graphical AOP network visualization tool to overlay a given AOP network with additional data such as feedback loops to facilitate interpretation without overly complicating the underlying framework.
AOP network topology	The overall shape and structure of an AOP network, describing the way in which the constituent parts of the network (i.e., key events and key event relationships) are interrelated or arranged.
Convergent topology	Topology in which key events from 2 or more AOPs are directed toward a common key event or adverse outcome, representing a range of possible upstream causes.
Divergent topology	Topology in which 2 or more key event relationships branch off from a single molecular initiating event or key event, representing a range of possible downstream outcomes.
Mixed topology	Topology showing local divergent and convergent regions within the overall network, possibly featuring specific motifs such as bow-tie motifs, which could represent important points of biological integration.
Critical path	The path through an AOP network considered most significant from an investigational, biological, or regulatory standpoint. A critical path does not necessarily correspond to a single AOP described in the AOP knowledgebase.
Interaction between AOPs	One AOP affecting another AOP in such a way that it modulates the adverse outcome compared with the outcome that would be observed had the interaction not taken place.

AOP = adverse outcome pathway.

development thus has the potential to mainly become an exercise in assembling data that already exist in the AOP knowledgebase. The process of developing AOP networks by extracting existing data from the AOP-Wiki and assembling a network based on those AOPs rather than on de novo descriptions of linked AOPs is called AOP network derivation.

AOP network derivation

The first step in network derivation is to extract all AOPs that are relevant for a given application from the AOP-Wiki (Figure 1). The criteria that define which AOPs are relevant will vary and will be defined by the application or stakeholder needs. Theoretically, the AOP knowledgebase can be queried for any property of an AOP, key event, or key event relationship that has been appropriately described and/or structurally annotated. Some examples of extraction criteria include the following: AOPs leading to a single adverse outcome of interest, AOPs known to be induced by a particular stressor or group of stressors, AOPs having key events that map to a particular data set (e.g., a collection of positive high-throughput screening assay responses observed for a particular chemical or mixture of chemicals), AOPs that have a particular species in their applicability domain, AOPs that have key events for a particular tissue type, and so on.

Extraction can be achieved manually, for example by inspecting dedicated pages in the AOP knowledgebase that list all the AOPs that a particular key event links to. However, manual extraction of AOP networks could rapidly become tedious as well as impractical as the AOP knowledgebase grows. Thus, it is important to develop computational tools designed for this purpose, such as the AOPXplorer (http://apps.cytoscape.org/apps/aopxplorer). Using AOPXplorer, any structured annotation field in the AOP knowledgebase can be queried computationally to derive an AOP network. Once such an automated extraction process is complete, the resulting collection of AOPs can be assembled based on their topologies of shared key events and key event relationships into an AOP network that is then called a primary AOP network (Figure 1). In some cases, the resulting primary network will be directly suitable for a certain application. In others, it may be desirable to refine (simplify and/or enrich) the network using a series of filters and data layer options, or to more deeply interrogate and statistically analyze the network, as discussed in the next section.

Refining AOP networks using filters

The structural complexity of AOP networks will depend on various factors. Ideally, AOP network derivation tools should include ways to focus and refine the network to fit the needs of a given application and enhance the information content conveyed from the overall network diagram. For example, risk assessment of individual chemicals or mixtures might be focused on a particular effect (e.g., impaired reproduction) in a specific class of organisms. In such a scenario, one might want to remove AOPs that relate to nonreproductive endpoints, as well as AOPs that are relevant to other taxa. On the other hand, efforts targeting mode of action identification could benefit from examination of highly branched networks encompassing many different MIEs and their associated pathways. Thus, it was conceived that one should be able not only to construct a primary network based on extraction criteria, but also to filter that network based on additional annotation terms that would allow one to focus on the pathway(s) of greatest interest.

It is envisioned that AOP network filters can be used to further define which key events and key event relationships from the primary AOP network would be included in downstream applications and analysis (Figure 1). For example, the structured key event and key event relationship domain of applicability terms selected in the AOP knowledgebase could be used to restrict a network to only those key events and key event relationships that are relevant to a given life stage, thereby simplifying the overall network. Alternatively, one might want to filter an AOP network to only those key events measured at a defined biological level of organization; to select appropriate endpoints, one might measure in a specific cell line or tissue. A

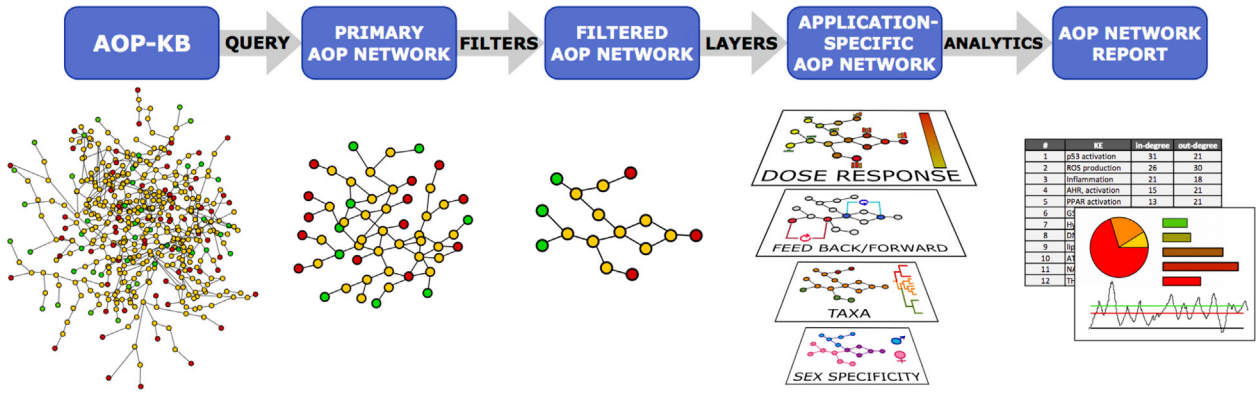

FIGURE 1: Graphical representation of the adverse outcome pathway (AOP) network derivation–refinement–analysis workflow. A primary AOP network is constructed by querying the AOP knowledgebase (AOP-KB). Filters are then applied to derive a filtered network containing AOPs of interest for a given application or research question. Layers can be added in a next step to add data relevant to the application. Finally, the AOP network can be analyzed to produce metrics related to the topology and other properties of the network.

range of different filters, based on either structured ontology terms that are part of the AOP descriptions (e.g., taxonomic applicability) or on network metrics (e.g., how strongly connected key events are to the network) could be envisioned. Supplemental Data Table S1, provides a list of possible filters that could be envisioned, including filters for taxonomic, life stage, or sex applicability, network metrics, and critical paths. Each could be used to help tailor an AOP network to a given problem formulation or research question. Finally, we propose a confidence assessment filter that can be used to filter AOP networks based on various weight-of-evidence, biological plausibility, essentiality, and other assessments of the constituent AOPs.

Visualizing AOP network data using layers

A simplified representation of a set of key events and key event relationships (i.e., an AOP network) can easily be visualized graphically, where each unique key event is represented by a single node, and the key event relationships are represented by edges (Figure 1). Although such a simple graphical representation can depict the general structure of an AOP network, it is not a practical means of displaying and interrogating all the complex information captured within each of its key event and key event relationship descriptions. In addition, one may wish to supplement a network with additional data that are external to the AOP knowledgebase (e.g., experimental data), which can further convolute the information associated with an AOP network. To aid in the visualization and interpretation of the complex information in AOP networks, we propose a mechanism to visually superimpose this information, as needed, as layers on top of an AOP network image (Figure 1). These AOP network layers can be viewed as analogous to the data layers employed in geographic information systems: information relevant to interpretation or application of an AOP network can be laid over the filtered AOP network, much like traffic or public transportation information is laid over a city map. Ideally these layers could capture data derived from structured annotation fields within the AOP-Wiki, and they could also incorporate other types of data that are not necessarily part of formal AOP descriptions.

There has been resistance to the explicit representation of additional data such as feedback loops as additional types of nodes and edges in an AOP network, because they may overly complicate network interpretation for many applications. On the other hand, for some applications, such additional levels of detail may yield insights that may allow for more accurately predicting biologically relevant outcomes. Layers add information to an AOP network without modifying or influencing the network's overall properties, structure, and topology, and they are viewed as a way to address competing desires for greater information richness and detail on the one hand versus clear-cut interpretive simplicity on the other hand. The consideration of feedback loops and modulating factors within AOPs and AOP networks provides a useful example of this. At present, events associated with a feedback loop may be included as key events in the AOP when a feedback response is causally linked to the adverse outcome and is measurable. In other cases, however—for example, when an understanding of the feedback loop may aid in predicting how severely a particular key event must be perturbed to progress further along the pathway—knowledge of the feedback loop can be included in the "quantitative understanding of the linkage" section of the relevant key event relationship pages (see Q&A 13 in LaLone et al. 2017a). Therefore, feedback, feedforward, or other types of signaling motifs or loops are not specifically annotated as such in AOP descriptions and are thus very difficult to identify automatically. Likewise, modulating factors that are extrinsic to the AOP network (i.e., are not driven by interactions among existing key events found in the network), such as dietary factors, genetic susceptibility or resistance, disease states, environmental factors, and so on, are currently only captured in the free-text descriptions of quantitative understanding of the key event relationships. Whereas potential intrinsic modulating factors are captured de facto in the structure of the network because they arise from a shared key event or key event relationship and, therefore, do not need explicit annotation, extrinsic modulating factors require separate descriptions and anchoring to the AOP network.

Operationally (i.e., from the perspective of further development of the AOP knowledgebase), the implementation of certain types of layers would involve the introduction of additional structured annotation fields (Ives et al. 2017) in the key event and key event relationship descriptions of the AOP knowledgebase. In the case of known modulating factors, this could, for example, involve the introduction of an optional "modulating factor" field to key event relationship descriptions, whereby users could define a modulating factor and provide additional text description and supporting references. An advanced implementation of feedback loop layers could allow future key events also affecting the feedback loop to reveal interactions between AOPs that are not necessarily evident from individual key events. However, even at the most basic level, the ability to apply a layer that identifies those key event relationships for which feedback or modulating factors are known to influence response–response relationships could be very informative and could signal a user to explore the additional details provided in the AOP description to determine whether they are relevant to the application in question. Although these capabilities have not yet been implemented as computational features of the AOP knowledgebase, the concepts and features outlined in the present article have been communicated to the AOP knowledgebase development team to inform ongoing software development aimed at enhancing the utility of AOP networks.

In addition to feedback loop layers and modulating factor layers, a number of other data layers were identified that could reflect taxonomic, life stage, and sex applicability domains, genetic heterogeneity, tissue specificity, and temporality, as well as quantitative response data (Supplemental Data, Table S2). We propose that in combination, the use of filters and layers will help to achieve a network representation that is suited for the intended application and will make the AOP knowledgebase more user friendly and useful for other intended audiences (such as risk assessors) in addition to research scientists. Importantly,

by overlaying certain data types on the key event relationships within an AOP network, the network representation can be transformed into a mathematical construct allowing for different types of analyses to be applied (Figure 1; Villeneuve et al. 2018, part II).

Analyzing AOP networks

An AOP network organizes sets of biological perturbations that may interact and influence one another in such a way that a significant understanding of the biology may be derived through examination and analysis of the structure of the network. Although visual examination of the network graph is compelling, the use of techniques from graph theory (Trudeau 2013) and network science (Lewis 2009) facilitates an encompassing review of the network, especially when networks become larger and more complex. Villeneuve et al. (2018, part II) address several aspects of AOP network analytics, building on the basic AOP network concepts described in the present article. They specifically focus on 3 key elements: 1) AOP network topology analysis, 2) critical path identification, and 3) characterization of interactions among AOPs in a network. In the present article we provide a few topical examples of analytical procedures that may be applied to AOP networks to give the reader a brief introduction to some of the concepts involved. The companion article (Villeneuve et al. 2018, part II) gives a complete overview and in-depth discussion of AOP analytics.

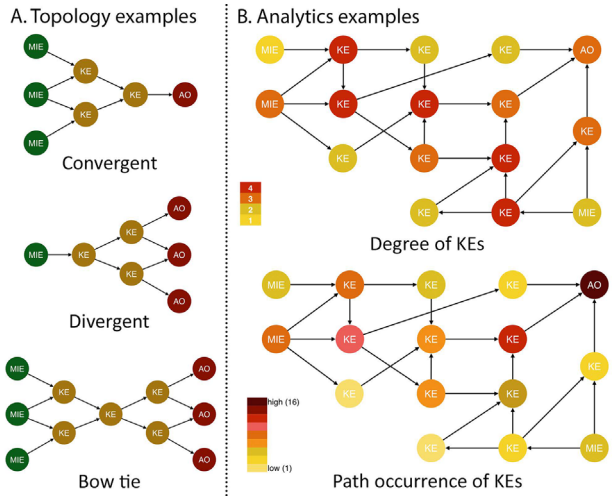

FIGURE 2: Examples of adverse outcome pathway (AOP) network analysis concepts and approaches. (**A**) Network topology analysis can reveal converging, diverging, or mixed patterns. A mixed pattern can take the shape of a bow-tie motif. (**B**) Two different examples of network metrics calculated for the same hypothetical AOP network. The degree of a node (key event [KE]) in the network is equal to the number of edges (key event relationships [KERs]) connecting the node to the network and is one way of expressing how connected that node is to the network. The path occurrence is the number of times a node (KE) occurs in a path connecting a molecular initiation event (MIE) to an adverse outcome (AO) after evaluating all possible paths between the MIEs and AOs of the network. The path occurrence may be an indication of the relative importance of a node within the overall network.

In AOP network topology analysis, a large variety of metrics can be calculated that describe the overall shape and structure of the network or identify specific nodes in the network that may be of particular interest. For example, one of the first topological properties of interest is comprised of points of convergence and divergence within a given network (Figure 2A). In a convergent topology, AOPs are directed toward a common key event or adverse outcome, whereas a divergent topology involves AOPs branching off from a common molecular intiating event or key event. Conceptually, the degree of convergence or divergence of a network may affect the intensity of the adverse outcomes, and analysis of convergence/divergence of AOP networks may inform on the existence of potential additive, synergistic, or antagonistic effects and interactions, or may, for example, be used to develop assays that would capture a broad range of molecular intiating events, versus assays predictive of a group of related adverse outcomes, versus assays predictive of only a very specific adverse outcome. Most real-life AOP networks will likely be mixed networks (i.e., have local divergent and convergent regions within the overall network). This could lead to specific motifs, such as a node that is a local site of convergence and divergence simultaneously, a mixed structure that would create a bow-tie motif (Figure 2A) and could represent important integrative biological signals. Computationally, a large number of metrics can be calculated to describe network topologies, each providing a specific view of the network and complementary opportunities for identifying network nodes of interest. A few examples of such metrics are given in Figure 2B.

A second and highly relevant characteristic of AOP networks is that they provide a framework for the description of the overall landscape of potential adverse outcomes resulting from particular biological perturbations. This can enable strategic identification of paths that have the greatest biological likelihood and/or relevance for risk assessment. Within an AOP network, the most significant path from an investigational or biological standpoint is termed the critical path. In the present article we distinguish "path" from "pathway" to recognize that the critical path may not necessarily follow an entire AOP, and may in fact emerge only through the assembly and consideration of the interactions between multiple AOPs. The interpretation of what constitutes a critical path can vary widely depending on the context and perspective of the AOP developed or end user. Critical paths may be representative of a specific research question, or of the strongest weight of evidence for certain elements of the network. They may also represent the most toxicologically relevant path that may have great importance in the application of AOP networks for risk assessment. This can in turn aid identification of endpoints or assays that can serve as useful alternatives to the direct measurement of apical adverse outcomes (Organisation for Economic Co-operation and Development 2016a). Also, AOP network–based critical path delineation efforts may be useful for identifying data gaps that are required to achieve a complete critical path description in scenarios in which the AOP network includes poorly supported AOPs. Even though critical paths currently remain a relatively loosely defined concept and quantitative approaches (i.e., quantitative AOP development) may be required to formulate

a more stringent definition, Villeneuve et al. (2018, part II) recognize the need for different critical path identification strategies, and distinguish among problem formulation, weight of evidence, and biologically–toxicologically defined critical paths, as well as the pure empirical identification of critical paths.

A third, and probably the most challenging, aspect of AOP network analysis is the identification and characterization of potential interactions between AOPs. The AOP interactions describe how one or more components of a pathway may affect another pathway in such a way that it modulates the adverse outcome in terms of its biological properties, intensity, probability, rate, and so on, compared with the outcome that would be observed had the interaction not taken place. Interactions between AOPs may be described as cross-talk between AOPs, but because the concept of cross-talk is typically associated with specific and rather strictly defined molecular processes such as signal transduction cascades, "interactions" is preferred as the descriptor. From a procedural perspective, because nodes in AOPs represent directional changes in the state of biological components (e.g., increased vs decreased testosterone concentrations are 2 different key events) rather than the biological components themselves (e.g., testosterone), it is recognized that tools to automatically map key events occurring on the same components during AOP network extraction and analysis will be required before the full potential of interaction analysis is achieved. Nevertheless, interactions are anticipated to result in additive, synergistic, or antagonistic responses (Vert and Chory 2011), and their analysis may provide the opportunity to guide a more rational assessment of mixture toxicity, for example (Villeneuve et al. 2018, part II).

AOP NETWORK APPLICATION: CASE STUDIES

As described by Villeneuve et al. (2014b), AOP networks (compared with single AOPs) were envisioned to be a more realistic representation of the complex biological interactions that would, for example, occur in response to exposures to chemical mixtures or single toxicants exhibiting multiple biological activities. The development and analysis of AOP networks have the 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. In addition, analysis of the intersections (shared key events and key event relationships) among AOPs that make up an AOP network can reveal unexpected or underappreciated biological connections. Consequently, it is anticipated that AOP networks will ultimately be more informative than individual AOPs in a decision-making context. For example, when one is mapping the landscape of AOPs for a particular adverse effect, the network will indicate the points of convergence of different pathways, which may indicate the most promising key event for development of in vitro assays that can be tailored to capture all the pathways upstream from that key event. This approach may be very useful for informing the construction of integrated approaches to testing and assessment to cover the relevant biology for a wide range of potential adverse outcomes (Tollefsen et al. 2014). The AOP networks may also offer insights into approaches for evaluating the toxicity of mixtures to understand how a chemical acting via one AOP may be impacted by another chemical acting via another AOP in a relevant mixture.

Although some of the most prominent potential applications of AOP networks have been noted, other applications may undoubtedly emerge. For example, AOP networks could help speed the design of new drugs or chemicals by providing early warnings of potential side effects or toxicological events that could possibly end up in adverse effects. Likewise, mapping layers of information on modulating factors onto an AOP network could help to identify vulnerable subcategories of people or wildlife whose susceptibility may be increased or decreased as a function of health status, microelement deficiencies, environmental stresses, and so on. These could either exacerbate the adverse effect of a chemical, or (equally undesirable) undermine or counteract the effect of a drug. Given the broad range of applications, it is impractical to illustrate them all. Thus, in the context of the present article, we highlight just a few application case studies that illustrate some of the concepts of AOP network development and analysis described previously and also show how those processes can be applied to help address questions related to chemical safety assessment.

Case study 1: AOP network for metabolic disorders mediated by hepatic steatosis

The need to develop AOP networks to effectively evaluate complex diseases was recently highlighted in the development of mechanistic toxicity tests based on an AOP network for hepatic steatosis, leveraging a large amount of publicly available mechanistic, phenotypic, and toxicological liver data (Angrish et al. 2016, 2017; Bell et al. 2016; Oki et al. 2016). Steatosis, also known as fatty liver disease, is a regulatory endpoint and pathologic condition in which energy metabolism is disrupted and fat accumulates in the liver. Energy homeostasis is dependent on the balance between energy intake and expenditure, a process regulated by endocrine and cellular communication among the brain, gut, and metabolic tissues such as adipose, striated (skeletal and cardiac) muscle, pancreas, and liver. At the molecular level, metabolism is coordinated by broad chemical signals, including nutrients, hormones, and environmental chemical signals that control systemic energy homeostasis by binding to cognate cell surface, cytosolic, and nuclear receptors. Chemical contact at any point along this neuro–endocrine–organ network can impact complex signal transduction, gene expression, protein activation cascades, and so on to coordinate the energy demands of a biological system. The challenge is that, because these receptors and signaling pathways cross-talk, it is difficult to adapt existing assay data (e.g., data from current ToxCast^{TM} and Tox21 assays) to strategies predictive of a steatotic outcome, possibly because the events these assays represent are too far upstream of the adverse outcome to allow for facilitation of reliable prediction of

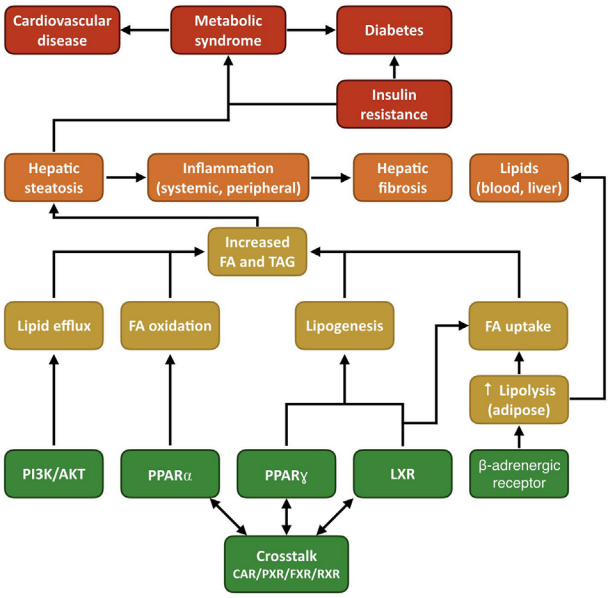

FIGURE 3: Adverse outcome pathway (AOP) network for metabolic disorders mediated by hepatic steatosis. The high level of cross-talk between the different receptors and associated signaling pathways complicates the use of existing high-throughput screening data as predictors of a steatotic outcome. This challenge was overcome by identifying a network topology converging into 4 key events (i.e., lipogenesis, and fatty acid uptake, efflux, and oxidation) that were viewed as critical paths leading to steatosis. Assays measuring these points of convergence integrate the complex interplay of upstream events and translate them into measures that are more directly related to the adverse outcome. FA = fatty acid; TAG = triacylglycerol; PI3K = phosphatidylinositol-3-kinase; AKT = protein kinase B; PPAR = peroxisome proliferator-activated receptor; LXR = liver X receptor; CAR = constitutive androstane receptor; PXR = pregnane X receptor; FXR = farnesoid X receptor; RXR = retinoid X receptor.

outcomes. Effectively, the interactions that occur in between are too complex to model practically or reliably.

In the steatosis AOP network, this challenge was overcome by identifying a network topology converging into 4 key events that were viewed as critical paths leading to steatosis (i.e., fatty acid uptake, efflux, synthesis, and oxidation; Figure 3). The assumption was that assays measuring these points of convergence would integrate the complex interplay of upstream events and translate them into key event measures or points of departure that are more proximally located relative to the adverse outcome. It is conceivable that such an approach would have the power to capture not only single chemical exposures, but also mixture effects, as long as the effects were upstream of the convergent key events.

Once the convergent key events were identified and the corresponding assays were developed, a second step was to utilize data from those assays to predict steatotic outcomes as well as their severity. A challenge is that the compensatory actions of these 4 key events collectively balance liver lipid levels. Consequently, progression toward a steatosis adverse outcome depends on the combination and magnitude of the change in key events and the interaction among all 4 key events and their associated AOPs. Although in some cases, only 1 of those 4 key events may be impacted and that 1 alone could be sufficient to elicit the adverse outcome, in most cases it is likely that more than 1 of the convergent key events will be affected. This can be expected to yield consequences different from those that might be predicted based on impacts on any one of those key events alone. For example, an exposure that increases lipid uptake may be sufficient to cause steatosis, whereas an alternative exposure that also activates lipid efflux may compensate for increased uptake and restore balance such that no adverse outcome is observed. This is a salient example of why the consideration of AOP networks has been viewed as critical to the use of the AOP framework for predictive toxicology. As such critical paths and points of convergence are identified, AOP network analyses can inform the development of complementary, biologically based mathematical models that facilitate an alternatives-based (e.g., cell-based assays) chemical evaluation workflow.

Case study 2: Decreased serum thyroid hormone AOP network for alternative assay development

An example of network-guided AOP development that has led to de facto construction of an AOP network in the AOP-Wiki is centered around circulating thyroid hormone concentrations. Two major points of convergence/divergence (i.e., key events resembling the knot of a bow-tie motif; see the *Visualizing AOP network data using layers* section) in this multitaxon AOP network are decreased serum T4 (thyroxine) and decreased serum T3 (triiodothyronine, see Figure 4A).

This thyroid hormone disruption AOP network has been employed to support the development and application of guideline toxicity tests and, subsequently, alternatives to those same whole-animal test guidelines. For example, the amphibian metamorphosis assay (Organisation for Economic Co-operation and Development [OECD] test guideline 231; Organisation for Economic Co-operation and Development 2009) was developed for the purpose of screening chemicals for their ability to disrupt the thyroid hormone signaling axis in vertebrates. The branches in the AOP network provide the scientifically plausible and evidence-based foundation for linking the shared key event of decreased serum T4 to impaired amphibian metamorphosis as an indicator of thyroid axis disruption. Adverse neurodevelopmental outcomes in rodents build the case for the relevance of the amphibian metamorphosis assay for screening thyroid-disrupting chemicals that can be adverse to humans (Figure 4A). Given the time- and resource-intensive nature of the amphibian metamorphosis assay, it was desirable to replace it with in vitro assays that could be used to screen large libraries of chemicals for their ability to disrupt the thyroid axis. Based on the AOP network, assays for thyroid peroxidase activity, the sodium iodide symporter, iodothyronine deiodinase, and iodotyrosine deiodinase activities were developed to assess the potential mechanisms through which chemicals could alter circulating T4 and/or tissue T3 concentrations (Figure 4A). Not all these targets have been covered in existing high-throughput screening programs (e.g., ToxCast, Tox21), so the AOP network helps to inform the development of a more comprehensive screening battery for this important mode of endocrine disruption.

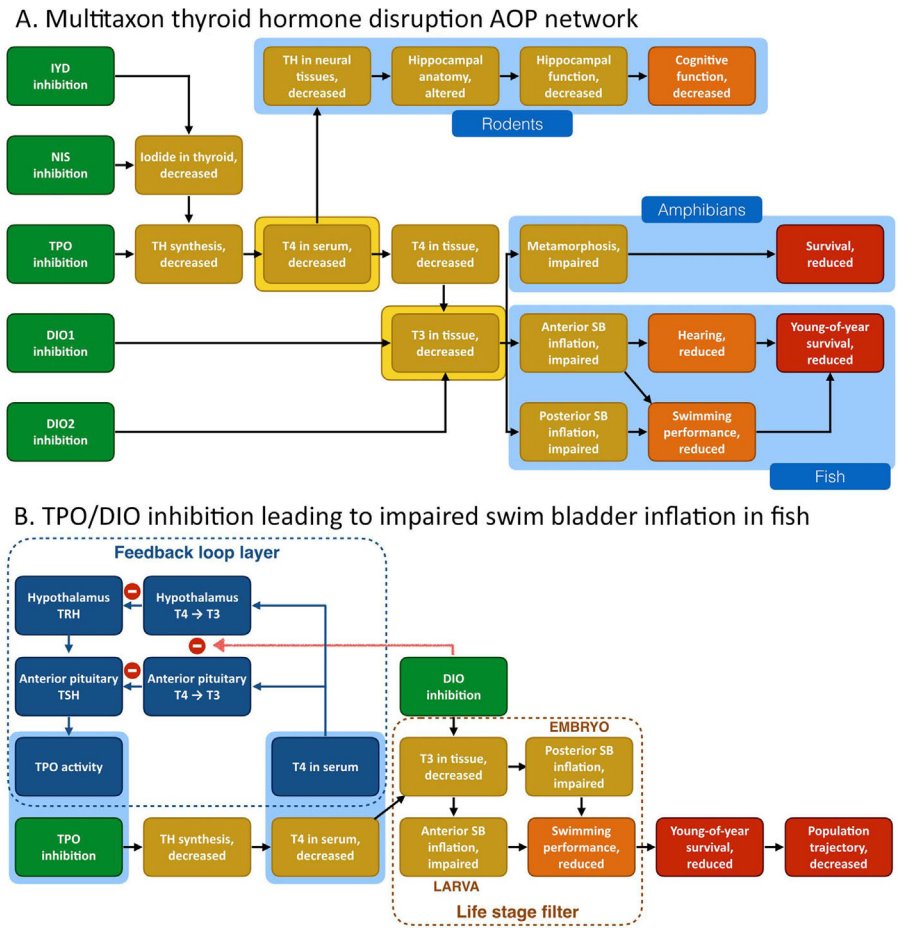

FIGURE 4: Adverse outcome pathway (AOP) networks related to disruption of the thyroid axis. (**A**) Multitaxon thyroid hormone disruption AOP network including mammalian, amphibian, and teleost endpoints. The blue regions illustrate how a taxonomic applicability layer may be used to add relevant data to the primary network representation. The key events highlighted in yellow indicate 2 major points of convergence/divergence in the network, resembling the knot of a bow-tie motif. (**B**) Filtered thyroid AOP network only containing key events that are relevant to fish. The dashed brown area illustrates how additional filtering might be used to further refine the network (e.g., to only include key events that are relevant to specific life stages). The blue area illustrates the use of a layer to indicate the presence of a feedback loop acting on an AOP in the network, and the interaction between the feedback loop and one of the molecular initiating events in the network. Red negative sign = inhibition processes. Red arrow = DIO inhibition, which decreases conversion of T4 into T3, thereby inhibiting the feedback inhibition of T3 on TRH and TSH synthesis. IYD = iodotyrosine deiodinase; NIS = sodium-iodide symporter; TPO = thyroperoxidase; DIO = iodothyronine deiodinase; TH = thyroid hormone; T4 = thyroxine; T3 = triiodothyronine; TRH = thyrotropin-releasing hormone; TSH = thyroid-stimulating hormone; thyrotropin; SB = swim bladder.

As part of another alternative testing development effort, a question was posed as to how the fish early life stage test (OECD test guideline 210; Organisation for Economic Co-operation and Development 2013b) might be replaced by more rapid and cost-effective alternatives (Villeneuve et al. 2014a). Although a modified fish embryo test (OECD test guideline 236; Organisation for Economic Co-operation and Development 2013c) had been proposed as an alternative that could cover much of the toxicological space encompassed by the fish early life stage test, it was recognized that certain developmental events occurring after hatch, during the larval to juvenile transition, could be missed. One example was swim bladder inflation, which in common laboratory-model cyprinids such as zebrafish and fathead minnow occurs in 2 stages: inflation of the posterior chamber shortly after hatch, followed by inflation of the anterior chamber several days to weeks later (Villeneuve et al. 2014a; Nelson et al. 2016; Stinckens et al. 2016; Cavallin et al. 2017). Although a range of biological perturbations may disrupt this event, decreases in circulating T4 and/or deiodination of T4 to T3 have been defined, through development of an AOP network, as a means through which chemicals could impact swim bladder inflation in fish, a key event that has been linked to reduced young-of-year survival (Czesny et al. 2005; Woolley and Qin 2010). The AOP network that focused on swim bladder inflation in fish was subsequently integrated with the broader amphibian/mammalian AOP network described in the previous paragraph, resulting in a multitaxon thyroid AOP network (Figure 4). Consequently, the same battery of in vitro assays that can plausibly screen for thyroid-disrupting chemicals in amphibian and mammalian models could also cover toxicological space that might be missed if a fish embryo test were employed as the only alternative to a fish early life stage test.

From a network development perspective, the thyroid AOP network demonstrates how some of the proposed filters and layers might be applied (Figure 4B). For example, application of a life-stage filter would show that the AOP mediated via

inhibition of the thyroid peroxidase enzyme is only relevant to larval fish. If the exposure was during the embryo stage only and the focus was inflation of the posterior chamber, then iodothyronine deiodinase enzyme inhibition would represent the critical path in the network. Alternatively, if the exposure occurred or was sustained until after hatch, both thyroid peroxidase and iodothyronine deiodinase inhibition would be inferred to be contributing to reduced anterior swim bladder inflation, suggesting that the outcome may be more severe than that triggered by a chemical exhibiting only one of the 2 bioactivities. Furthermore, invoking the feedback loop layer in the AOP network visualization could unveil additional detail relevant to predicting the interactive effect of these 2 AOPs, because the molecular initiating event of iodothyronine deiodinase inhibition also impacts the negative feedback loop mechanism itself. Adding the quantitative properties of this feedback mechanism to the response–response relationship of the key event relationship linking decreased T4 levels to reduced anterior swim bladder inflation might provide for a more accurate prediction of the joint effect of the 2 AOPs than the basic AOP network alone would provide.

Additional case studies

Two additional, fully described case studies are given in the Supplemental Data to provide the interested reader with further examples illustrating AOP network development and application in more advanced scenarios. The first case study illustrates the application of AOP networks to support the assessment of complex mixtures. A water sample extract of a metropolitan wastewater treatment plant was tested using a number of ToxCast assays to evaluate the ability of the sample to activate different nuclear receptors and transcription factor promoter–regulated reporter sequences. Assay activity was mapped to molecular initiating events described in the AOP-Wiki, and the resulting AOP network was filtered to focus on key events that were directly relevant to the observed bioactivities. The resulting set of AOP networks was further filtered to exclude AOPs that did not terminate at adverse outcomes that would be considered relevant to ecological risk assessment. Focusing on the remaining AOPs, known potential hazards to aquatic vertebrate wildlife associated with this mixture could be identified. The second case study provides an example of how an AOP network approach was used to explore the polypharmacological profile of the pharmaceutical beclomethasone dipropionate using the fathead minnow. Because of its ability to modulate the glucocorticoid receptor, beclomethasone dipropionate is used to treat chronic inflammatory conditions, but the drug also has the ability to modulate the androgen and progesterone receptors. Data generated during drug development were used to identify the cascades of key events likely to be triggered, and this information was organized within an AOP network. Chronic in vivo exposures to beclomethasone dipropionate were then carried out to generate a quantitative AOP network, which provided evidence that the polypharmacology profile of the beclomethasone dipropionate was indeed critically important to interpret and accurately predict the toxicological profile of the drug (Margiotta-Casaluci et al. 2016).

SUMMARY AND CONCLUSIONS

Based on the results of a SETAC-sponsored Horizon Scanning exercise focused on advancing the AOP framework, the development of guidance and best practices related to AOP network derivation and application was identified as a critical need. This not only included questions and concerns focusing directly on AOP networks, but also on different related topics such as mixture toxicity assessment, the implementation and graphical representation of feedback loops within the AOP framework, the characterization of interactions among pathways, the ability to include information on extrinsic modulating factors, and so on. Although the concept of constructing networks has always been deliberately, but possibly rather implicitly, built into the AOP framework (Villeneuve et al. 2014b, 2014c), the number of available AOPs has only recently reached a level sufficient to begin developing AOP networks. Recognizing different needs and strategies for developing AOP networks, we distinguish between network-guided AOP development and AOP network derivation based on the AOP knowledgebase. We then propose the use of filters and layers to simplify visualization and interpretation of AOP networks, and to tailor them to suit the needs of a given research question or application. The AOP networks can subsequently be analyzed in a variety of ways to extract useful information, including topological analyses, critical path identification, and characterization of interactions among AOPs within a network. The concepts described in the present article, and in its companion article focused on AOP network analytics, are intended to serve as a starting point for further development of the AOP network concept and of the AOP knowledgebase to increase its capabilities for managing and analyzing AOP networks, but also to catalyze AOP network development and application by the different stakeholder communities. Along with other manuscripts produced as a result of the April 2017 SETAC Pellston Workshop on Advancing the Adverse Outcome Pathway Framework (LaLone et al. 2017a), we hope to serve the ongoing development of the AOP framework in general as a critical concept to support 21st century approaches to toxicological research and regulation.

Supplemental Data—The Supplemental Data are available on the Wiley Online Library at DOI: 10.1002/etc.4125.

Acknowledgment—We gratefully acknowledge the Society of Environmental Toxicology and Chemistry (SETAC) North America staff, in particular G. Schiefer, N. Mayo, and T. Schlekat, who provided support to the workshop co-chairs, steering committee, and participants before, during, and after the workshop. We appreciate funding support from SETAC, the US Environmental Protection Agency, the American Cleaning Institute, the European Chemical Industry Council Long-Range Research Initiative, Chevron Environmental, the European Center for Ecotoxicology and Toxicology of Chemicals, the European Commission Joint Research Centre, European Crop Protection Association, ExxonMobil, Humane Society International, The Humane Society of the United States, the Human Toxicology Project Consortium, Syngenta, and Unilever. In addition, we thank the groups from academia, industry, and

government who supported participants' travel. The authors thank the workshop co-chairs, C. LaLone and M. Hecker, for their coordination, organization, and guidance. We acknowledge the other workshop participants for their stimulating discussions and feedback, and the respondents to the Horizon Scanning effort for the charge questions and themes that informed our discussion.

Disclaimer—The contents of this article represent the personal opinions of the authors and neither constitute, nor necessarily reflect, the policies or viewpoints of their employers or institutes.

Data Availability—Data, associated metadata, and calculation tools are available from the corresponding author (dries.knapen@uantwerpen.be).

REFERENCES

Angrish MM, Kaiser JP, McQueen CA, Chorley BN. 2016. Tipping the balance: Hepatotoxicity and the 4 apical key events of hepatic steatosis. *Toxicol Sci* 150:261–268.

Angrish MM, McQueen CA, Cohen-Hubal E, Rooney JP, Bruno M, Ge Y, Chorley BN. 2017. Mechanistic toxicity tests based on an adverse outcome pathway network for hepatic steatosis. *Toxicol Sci* 159:159–169.

Ankley GT, Bennett RS, Erickson RJ, Hoff DJ, Hornung MW, Johnson RD, Mount DR, Nichols JW, Russom CL, Schmieder PK, Serrrano JA, Tietge JE, Villeneuve DL. 2010. Adverse outcome pathways: A conceptual framework to support ecotoxicology research and risk assessment. *Environ Toxicol Chem* 29:730–741.

Bell SM, Angrish MM, Wood CE, Edwards SW. 2016. Integrating publicly available data to generate computationally predicted adverse outcome pathways for fatty liver. *Toxicol Sci* 150:510–520.

Cavallin JE, Ankley GT, Blackwell BR, Blanksma CA, Fay KA, Jensen KM, Kahl MD, Knapen D, Kosian PA, Poole ST, Randolph EC, Schroeder AL, Vergauwen L, Villeneuve DL. 2017. Impaired swim bladder inflation in early life stage fathead minnows exposed to a deiodinase inhibitor, iopanoic acid. *Environ Toxicol Chem* 36:2942–2952.

Czesny SJ, Graeb BDS, Dettmers JM. 2005. Ecological consequences of swim bladder noninflation for larval yellow perch. *Trans Am Fish Soc* 134:1011–1020.

Cytoscape. 2017. AOPXplorer. Available from: http://apps.cytoscape.org/apps/aopxplorer

Ives C, Campia I, Wang RL, Wittwehr C, Edwards SW. 2017. Creating a structured AOP knowledgebase via ontology-based annotations. *Appl* In Vitro *Toxicol* 3:298–311.

Knapen D, Vergauwen L, Villeneuve DL, Ankley GT. 2015. The potential of AOP networks for reproductive and developmental toxicity assay development. *Reprod Toxicol* 56:52–55.

LaLone CA, Ankley GT, Belanger SE, Embry MR, Hodges G, Knapen D, Munn S, Perkins EJ, Rudd MA, Villeneuve DL, Whelan M, Willett C, Zhang X, Hecker M. 2017a. Advancing the adverse outcome pathway framework—An International Horizon Scanning approach. *Environ Toxicol Chem* 36:1411–1421.

LaLone CA, Villeneuve DL, Wu-Smart J, Milsk RY, Sappington K, Garber KV, Housenger J, Ankley GT. 2017b. Weight of evidence evaluation of a network of adverse outcome pathways linking activation of the nicotinic acetylcholine receptor in honey bees to colony death. *Sci Total Environ* 584:751–775.

Lewis TG. 2009. *Network Science: Theory and Applications*. John Wiley & Sons, Hoboken, NJ, USA.

Margiotta-Casaluci L, Owen SF, Huerta B, Rodriguez-Mozaz S, Kugathas S, Barcelo D, Rand-Weaver M, Sumpter JP. 2016. Internal exposure dynamics drive the Adverse Outcome Pathways of synthetic glucocorticoids in fish. *Sci Rep* 6:6:21978.

Nelson KR, Schroeder AL, Ankley GT, Blackwell BR, Blanksma C, Degitz SJ, Flynn KM, Jensen KM, Johnson RD, Kahl MD, Knapen D, Kosian PA, Milsk RY, Randolph EC, Saari T, Stinckens E, Vergauwen L, Villeneuve DL. 2016. Impaired anterior swim bladder inflation following exposure to the thyroid peroxidase inhibitor 2-mercaptobenzothiazole part I: Fathead minnow. *Aquat Toxicol* 173:192–203.

Organisation for Economic Co-operation and Development. 2009. Test No. 231: Amphibian metamorphosis assay. *OECD Guidelines for the Testing of Chemicals*. Paris, France.

Organisation for Economic Co-operation and Development. 2013a. Guidance on developing and assessing adverse outcome pathways. Series on Testing and Assessment, No. 184. ENV/JM/MONO(2013)6. Paris, France.

Organisation for Economic Co-operation and Development. 2013b. Test No. 210: Fish, early-life stage toxicity test. *OECD Guidelines for the Testing of Chemicals*. Paris, France.

Organisation for Economic Co-operation and Development. 2013c. Test No. 236: Fish embryo acute toxicity (FET) test. *OECD Guidelines for the Testing of Chemicals*. Paris, France.

Organisation for Economic Co-operation and Development. 2016a. Users' handbook supplement to the guidance document for developing and assessing AOPs. ENV/JM/MONO(2016)6. Paris, France. [cited 2017 April 7]. Available from: https://doi.org/10.1787/5jlv1m9d1g32-en

Organisation for Economic Co-operation and Development. 2016b. Guidance document for the use of adverse outcome pathways in developing integrated approaches to testing and assessment (IATA). Series on Testing and Assessment, No. 260. ENV/JM/MONO(2016)67. Paris, France.

Organisation for Economic Co-operation and Development. 2013. Adverse outcome pathways knowledge base. Paris, France. [cited 2017 April 7]. Available from: aopkb.oecd.org/.

Oki NO, Nelms MD, Bell SM, Mortensen HM, Edwards SW. 2016. Accelerating adverse outcome pathway development using publicly available data sources. *Curr Environ Health Rep* 3:53–63.

Society for the Advancement of Adverse Outcome Pathways. 2013. AOP-Wiki. [cited 2017 April 7]. Available from: http://aopwiki.org/.

Stinckens E, Vergauwen L, Schroeder AL, Maho W, Blackwell BR, Witters H, Blust R, Ankley GT, Covaci A, Villeneuve DL, Knapen D. 2016. Impaired anterior swim bladder inflation following exposure to the thyroid peroxidase inhibitor 2-mercaptobenzothiazole part II: Zebrafish. *Aquat Toxicol* 173:204–217.

Tollefsen KE, Scholz S, Cronin MT, Edwards SW, de Knecht J, Crofton K, Garcia-Reyero N, Hartung T, Worth A, Patlewicz G. 2014. Applying adverse outcome pathways (AOPs) to support integrated approaches to testing and assessment (IATA). *Regul Toxicol Pharmacol* 70:629–640.

Trudeau RJ. 2013. *Introduction to Graph Theory*. Dover, New York, NY, USA.

Vert G, Chory J. 2011. Crosstalk in cellular signaling: Background noise or the real thing? *Dev Cell* 21:985–991.

Villeneuve D, Volz DC, Embry MR, Ankley GT, Belanger SE, Leonard M, Schirmer K, Tanguay R, Truong L, Wehmas L. 2014a. 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.

Villeneuve DL, Crump D, Garcia-Reyero N, Hecker M, Hutchinson TH, LaLone CA, Landesmann B, Lettieri T, Munn S, Nepelska M, Ottinger MA, Vergauwen L, Whelan M. 2014b. Adverse outcome pathway (AOP) development I: Strategies and principles. *Toxicol Sci* 142:312–320.

Villeneuve DL, Crump D, Garcia-Reyero N, Hecker M, Hutchinson TH, LaLone CA, Landesmann B, Lettieri T, Munn S, Nepelska M, Ottinger MA, Vergauwen L, Whelan M. 2014c. Adverse outcome pathway development II: Best practices. *Toxicol Sci* 142:321–330.

Villeneuve DL, Angrish MM, Fortin MC, Katsiadaki I, Leonard M, Margiotta-Casaluci L, Munn S, O'Brien JM, Pollesch NL, Smith LC, Zhang X, Knapen D. 2018. Adverse outcome pathway networks II: Network analytics. *Environ Toxicol Chem* 37:1734–1748, (*this issue*).

Woolley LD, Qin JG. 2010. Swimbladder inflation and its implication to the culture of marine finfish larvae. *Rev Aquacult* 2:181–190.

